

# TeamForge 17.4



# Contents

<b>TeamForge 17.4.....</b>	<b>12</b>
How to use TeamForge 17.4.....	12
Get started.....	12
Contribute to a CollabNet TeamForge project.....	32
Lead a project team with CollabNet TeamForge.....	135
Support project teams with CollabNet TeamForge.....	314
Reference information about CollabNet TeamForge 17.4.....	377
Reference information about your CollabNet TeamForge site.....	377
Reference information for TeamForge project administrators.....	387
Reference information about CollabNet TeamForge site administration.....	410
CollabNet TeamForge full-text search.....	434
<b>Install / Upgrade.....</b>	<b>438</b>
Plan your installation/upgrade.....	438
TeamForge installation requirements.....	441
Set up networking for your TeamForge server.....	444
Which ports are required to be open for a TeamForge site?.....	444
Install or upgrade TeamForge.....	447
Install TeamForge.....	448
Upgrade TeamForge.....	494
Back up and restore TeamForge and EventQ.....	566
<b>EventQ.....</b>	<b>571</b>
Get started.....	571
Overview.....	571
Architecture.....	572
User authentication and authorization.....	575
Activity Sources.....	576
Manage build source.....	576
Manage review sources.....	577
Manage commit sources.....	578
Manage work item sources.....	579
Extensible Data Source (XDS) Overview.....	580
Associations.....	581
TeamForge Associations.....	582
<b>Set up Review Board.....</b>	<b>583</b>
Install Review Board.....	583
Upgrade Review Board.....	585
Integrate Review Board with your TeamForge site.....	587
Add Review Board to a TeamForge project.....	588
TeamForge - Review Board integration: FAQs.....	589
TeamForge - Review Board integration: Install FAQs.....	589
TeamForge - Review Board integration: General usage FAQ.....	589
How do I control user access for Review Board?.....	590

<b>Set up Git</b> .....	<b>591</b>
TeamForge—Git integration: an overview.....	591
Install Git on a separate server.....	592
Set up a Git replica server.....	594
Replicate repositories with Git replica servers.....	597
Set up Large File Storage (LFS) for TeamForge repository.....	601
Enable LFS for a repository.....	602
Set up LFS client and work with large files.....	603
Add Git as a linked application.....	604
Control the code review policy for Git repositories.....	605
Mandatory code reviews for Git repositories.....	606
Optional code review for Git repositories.....	606
Default code review for Git repositories.....	607
Custom code review for Git repositories.....	608
Pull request reviews for Git repositories.....	609
User defined reviews for Git repositories.....	610
TeamForge-Git integration: History protection.....	611
TeamForge-Git integration: Enable history protection.....	611
History protection reports.....	612
TeamForge-Git integration: History protection FAQ.....	613
TeamForge-Git integration: FAQ.....	615
TeamForge-Git integration: General usage FAQ.....	615
TeamForge-Git integration: Technical concepts FAQ.....	619
TeamForge Git integration reference.....	621
Mappings between TeamForge and Gerrit.....	621
Gerrit directory structure, connectivity and more.....	624
Access rights in Gerrit 2.13.x.....	625
Gerrit configuration properties.....	628
Illustrations on TeamForge-Gerrit communication.....	632
<b>Integrate TeamForge 17.4 with other tools</b> .....	<b>634</b>
Adapters: Integration using TeamForge EventQ.....	634
Source association keys.....	634
JIRA adapter overview.....	635
JIRA Cloud integration overview.....	643
Jenkins adapter overview.....	647
TeamCity adapter overview.....	652
External Git repository integration overview.....	654
External Subversion repository integration overview.....	656
Crucible adapter overview.....	660
Review Board adapter overview.....	662
Chef adapter overview.....	664
Install Nexus.....	669
Install the TeamForge-Nexus integration plugin.....	670
Upgrade the TeamForge-Nexus integration plugin.....	671
TeamForge-Nexus release notes.....	672
Install Artifactory Pro.....	674
Install the TeamForge-Artifactory integration plugin.....	675
TeamForge EventQ JIRA Integration: an overview.....	675
TeamForge-TestLink integration.....	676
TeamForge-TestLink integration: an overview.....	676
Installation and configuration.....	679
Define requirements, create test plans and test cases, and execute test cases.....	683

Set up Pebble.....	686
Install Pebble.....	686
Integrate Pebble into your TeamForge site.....	687
Add Pebble to a TeamForge project.....	688
Authenticate your integrated application with TeamForge.....	689
Internationalize your integrated application.....	690

## **Maintain your TeamForge 17.4 site..... 691**

TeamForge System Administration.....	691
Install project templates manually.....	691
Install project templates using the installer.....	691
Set up an Oracle database.....	691
Change the scmviewer password.....	694
Install Memcached.....	694
Synchronize TeamForge source control integrations.....	695
Upgrade CLI reports to the latest version on Red Hat or CentOS.....	696
Customize your Apache configuration.....	696
Customize your PostgreSQL configuration.....	697
EventQ System Administration.....	698
Monitoring TeamForge EventQ with New Relic.....	700
EventQ Troubleshooting.....	700
Authenticate users with LDAP.....	701
Set up LDAP integration for the TeamForge site.....	701
Turn off LDAP authentication.....	702
Authenticate users with LDAP using Auth Manager.....	702
Maintain database.....	712
Schedule data extraction for reporting.....	712
Raise the logging visibility of selected database requests.....	713
Optimize PostgreSQL with vacuum.....	713
Move the datamart (dedicated database server mode).....	714
Protect your CollabNet TeamForge site.....	715
Encrypt communication between the application and database servers.....	715
Set up SELinux.....	716
Protect your TeamForge site with SSL.....	717
Logs.....	720
Change the location of a log file.....	720
Change the logging level on your site.....	720
Search index optimization.....	721
Rebuild TeamForge search indexes.....	721
Permit big file uploads.....	721
Allow searching of archive files.....	722
SCM.....	722
Provide more than one source control server.....	722
Email.....	723
Removing users from monitoring objects.....	723
Limit the size of message attachments.....	723
Limit the size of document attachments.....	724
Discussions.....	724
Who can post to discussions by email?.....	724
Who can monitor discussions?.....	725
Reduce discussion spam.....	725
General.....	726
Change your site's domain name.....	726
Let users see what's in a project template.....	727

<b>Reference information about TeamForge.....</b>	<b>728</b>
Configuration files in TeamForge.....	728
site-options.conf.....	728
pebble-app.xml.....	776
pebble-dep.xml.....	778
How is an integrated application described?.....	778
install.conf - Pebble.....	782
install.conf - Review Board.....	782
iptables.....	783
login-config.xml.....	783
Scripts installed with TeamForge.....	785
TeamForge script.....	785
add_eventq_to_projects.py.....	789
backup-data.py.....	790
<b>backup-rb-data.py</b> .....	790
db.py.....	791
domain_change_db.py.....	791
etl-client.py.....	792
fix_data_permission.sh.....	793
install.sh.....	793
pbl.py.....	794
password_util.sh.....	795
post-install.py.....	796
psql-wrapper.....	797
psql-reporting-wrapper.....	797
restore-data.py.....	797
<b>restore-rb-data.py</b> .....	798
SearchReindex.py.....	798
set_auth_key.py.....	800
sync-orchestrate.py.....	800
Log files in TeamForge.....	801
JBoss logs.....	801
Oracle logging.....	801
SCM (CVS and Subversion) logs.....	801
Email logs.....	802
Search logs.....	803
Project Build Library audit log.....	803
Profile audit log.....	803
User Audit Log.....	804
Host audit log.....	804
Project audit log.....	804
etl.log.....	805
Reference information about integrated applications.....	805
IntegratedAppSupport.....	805
SOAP calls for integrated applications.....	806
Pebble.....	809
 <b>Extend TeamForge.....</b>	 <b>821</b>
Building on the TeamForge platform.....	821
Event handlers: Create custom work flows in TeamForge.....	821
API guide: Use TeamForge features via SOAP.....	831
TeamForge EventQ API documentation.....	884

## Frequently asked questions about CollabNet TeamForge 17.4..... 941

TeamForge installation/upgrade/configuration/administration FAQs.....	941
How many servers do I need to run a TeamForge site?.....	941
Should I upgrade to TeamForge 17.4 on a new server?.....	942
Should I move my TeamForge database to its own server?.....	942
Should I move my source control application to its own server?.....	943
What are the right PostgreSQL settings for my site?.....	943
Is it possible to change artifact prefixes in TeamForge?.....	944
I have made changes to the James config, how can I enable them without restarting TeamForge?.....	944
Why am I getting an error about NoEmailShouldEverBeSentToNoBody@localhost?.....	944
How does TeamForge handle third-party applications?.....	944
When do I run the initial load job?.....	945
When do I run the Tracker initial load job?.....	946
Why do I get an error during tracker incremental load job?.....	947
Can I run other java applications in the same JBoss instance as CollabNet TeamForge?.....	947
Users invoked via su command in TeamForge.....	947
Why does the the Yum installer display a warning message on CentOS 6?.....	948
Why am I getting a Yum repository filename conflict?.....	948
Why is a user restricted and cannot be changed?.....	948
How do TeamForge licenses work?.....	948
What happens when my site license expires?.....	951
How to reinstall a deleted installation directory?.....	951
How do I configure the user session timeout?.....	952
Where do I configure my client proxy settings?.....	952
How do I configure the timeout for Apache in TeamForge?.....	952
How do I make TeamForge work the same when the IP address of the server changes?.....	952
How do I change the JBoss log level in TeamForge 17.4?.....	952
Upgrade to TeamForge 8.0 and later fails if the field 'team' already exists. How do I resolve this?....	952
Troubleshooting TeamForge errors.....	953
I deleted the email that verifies my CVS password. How can I retrieve my password?.....	953
Not all my team members appear in the "assigned to" field. Why is that?.....	954
Why don't the branding repo changes get rendered into UI?.....	954
Why am I getting "Could not connect" status for my email and search server?.....	954
Why am I not able to see the charts for tracker metrics?.....	954
Post install fails for Git. What should I do?.....	954
Why am I getting a 'Not running' message when the Datamart service is stopped?.....	954
Why are the dynamic images that TeamForge creates broken?.....	954
Due to firewall restrictions I cannot send email from James. How can I resolve this?.....	955
Why is my email taking a long time to arrive?.....	955
How do I enable the PostgreSQL log file archiving when the services are not started using the CollabNet startup script?.....	956
Why am I not getting any error messages when executing the Subversion upgrade script?.....	956
Why do we have errors creating or altering repositories and adding or removing users?.....	956
Why am I getting an email specifying that the ETL job has failed?.....	956
My custom event handler is not working! What is wrong?.....	957
Why do I get a JBoss error - "failed to start in 240 seconds, giving up now" - while installing TeamForge?.....	957
JBoss crashed with out of memory error, how do I prevent this?.....	957
JBoss status is in 'starting' for a long time. How to have JBoss started successfully?.....	958
Why would some users not get email?.....	958
Why do I get 'Name is Invalid' error when trying to create a project using createProject () method via the SOAP API?.....	958
Why is the password and login shell changed for users on my cvs/svn server?.....	958
Why am I not able to see the status of the Postgres in the collabnet startup script?.....	959
What does the "psql: could not connect to the server: No such file or directory" error message mean?.....	959
Why do I get a proxy timeout when I try to view certain SCM pages?.....	959
Why do search and email server show "Could not connect"?.....	959

Why do I get a server status error when I perform a search?.....	959
Why can't TeamForge send my outbound mail?.....	960
Why does the SOAP services show "could not connect" on the Server Status page when everything else appears to work?.....	960
Why does startup fail or produce errors?.....	960
Why do I get a TeamForge system error in the project template creation page?.....	960
How can a user who deleted a document get it back?.....	960
Why do I get a URL "not found" or "moved permanently" error after applying a patch/upgrade?.....	961
Unable to embed Jenkins page into an iframe. How to fix it?.....	962
The TeamForge-Binary integration is down after enabling SSL. What should I do?.....	962
Why do ETL jobs fail post TeamForge upgrade?.....	962
Why are statistics charts broken on a wiki page?.....	963
Error while modifying custom reports via branding repository. How to fix?.....	963
Manage security.....	964
After switching to ADS authentication, why did the Create button disappear from the user admin section?.....	964
Can I block project data from public search engines?.....	964
How can I enforce strong passwords?.....	964
How do I configure Subversion to authenticate against multiple LDAP domains?.....	964
How do I authenticate multiple LDAP via Apache?.....	965
How does TeamForge authenticate CVS users?.....	966
Can the users be forced to change their passwords at first login?.....	967
Does TeamForge work with LDAP?.....	967
Why do I get the "Invalid command 'AuthLDAPAuthoritative'" error when I try to set LDAP for SVN users?.....	968
How does TeamForge handle multiple redundant LDAP servers?.....	969
What user activities are tracked?.....	969
What is a CERT advisory?.....	969
J2EE Architecture and security.....	969
With self-signed certificates, what is the recommended protocol (SSH/HTTPS) while cloning a repository from TeamForge?.....	970
Can I run TeamForge on RHEL?.....	970
Roles and permissions.....	970
Can I block a user's access to source code?.....	970
Can I assign a role to all users of the site at once?.....	970
Can I request a role?.....	970
Why don't I have access to the "Reported in Release" information in my artifact?.....	971
How do user permissions in CollabNet Enterprise Edition map to permissions in TeamForge 17.4?.....	971
Who can access an application?.....	972
Can I delete an item if I have 'Administer' permission?.....	973
Can I disable creation of user accounts?.....	973
Why do I see the project home page though I lack the necessary permission?.....	974
Why do you need the authentication and authorization plugin for Hudson?.....	974
Can I control user access to an integrated application?.....	974
How does inheritance work?.....	975
How does inheritance work for project groups?.....	975
Can I create new user accounts as "unrestricted"?.....	975
Why can't Oracle connect to my TeamForge installation?.....	976
Why can't I move these artifacts into this planning folder?.....	976
Who can access a project?.....	976
Who should I allow into my project?.....	976
How do I give read-only anonymous access to cvs/svn repository?.....	977
Who can access source code?.....	978
Can I ensure that only site admins can create projects?.....	980
How do user roles work?.....	981
How does CollabNet TeamForge help protect data access?.....	982
How does CollabNet TeamForge help protect my data?.....	983
How do site administrator roles work?.....	983
Which role is assigned to me?.....	985
Who can access project planning information?.....	985
Who can see a project page?.....	986

Who can see a project page component?.....	986
Who is an owner of the sub component?.....	986
As a project admin, why don't I have permissions to the wiki?.....	986
Concepts and terms in TeamForge.....	986
What is an association?.....	987
What is change management?.....	987
What is a document in CollabNet TeamForge?.....	987
What is a feature tree?.....	987
What is a discussion forum?.....	988
What is an integrated application?.....	988
What is an iteration?.....	989
What is a linked application?.....	989
What is the look project?.....	989
What is a parent project?.....	990
What is a patch?.....	990
What is a planning folder?.....	990
What are Planning, Task and Kanban boards?.....	991
What is a product owner?.....	992
What is a project administrator?.....	993
What is a project?.....	993
What is an agile project team?.....	993
What is in a project template?.....	994
What is a project template?.....	995
What is a release?.....	995
What is a story point?.....	995
What is a task dependency?.....	996
What is a tracker?.....	996
What is tracker work flow?.....	996
What is a user group?.....	997
Velocity and average velocity.....	997
What is an activity table component?.....	998
What is a documents component?.....	999
What is a global project role?.....	999
What is a project page?.....	999
What is a project page component?.....	999
What is an Available upon Request role?.....	999
What is a news component?.....	1000
What is a project page component for?.....	1000
What is a Tracker Search Results component?.....	1000
What is a project statistics component?.....	1001
What is TeamForge?.....	1001
What is a text component?.....	1002
What is a wiki component?.....	1002
What is the complete wiki syntax that TeamForge supports?.....	1002
What is Team?.....	1002
What is a team owner?.....	1002
General FAQs.....	1002
What's the best way to search the CollabNet User Information Center?.....	1003
What is a Publishing repository? How does it work?.....	1003
Is the CollabNet TeamForge SOAP API backward-compatible?.....	1004
What can I do with the SOAP API?.....	1004
How should my API client store user passwords?.....	1005
How does an application interact with TeamForge SOAP services?.....	1005
How do I require approval for new user accounts?.....	1006
How does TeamForge automatically sum up effort estimates?.....	1006
How does TeamForge use Velocity templates?.....	1007



How do URLs in CollabNet TeamForge 17.4 compare with URLs in CollabNet Enterprise Edition?.....	1009
Moving to a CollabNet TeamForge project from a CollabNet Enterprise Edition project.....	1011
Do I have to use the administrator-given password always?.....	1012
What software configuration management tools are available in CollabNet TeamForge ?.....	1012
How do I determine which version of TeamForge is currently running?.....	1012
Why can't I reply when someone comments on my review?.....	1012
How does TeamForge support dynamic planning?.....	1012
How do you measure "effort"?.....	1013
Why is TeamForge Explorer forgetting my settings and not logging me in?.....	1013
What is the file size limit on uploads or attachments in TeamForge?.....	1014
How can I import tasks into TeamForge Task Manager?.....	1014
Can I associate objects of different projects?.....	1014
How can I find the number of files in a repository without checking it out?.....	1014
What happens when log files get too big?.....	1014
How can I make the project pages invisible for any user who is authorized to see the project?.....	1014
How do I make the monitoring messages be sent from Forge Administrator?.....	1015
How can I access the native install version of TeamForge?.....	1015
How do I enable/disable path-based permissions via SOAP?.....	1015
How do I enable post-commit logging?.....	1015
How do I put a notice to my users on the project home page?.....	1016
What is a private IP address and what are the private IP ranges?.....	1016
What is the suggested log configuration for a production system?.....	1016
What is the difference between a stagger and normal header in query result heading settings?.....	1016
Can the query result be listed without the issue id?.....	1017
How do I remove a news item?.....	1017
How do I remove the OCN link from my TeamForge instance?.....	1017
How do I send email from a specific sender address instead of the member address?.....	1017
Where can I download the TeamForge DTD?.....	1017
How do I remove the build and test link from TeamForge pages?.....	1017
Can I query the CollabNet TeamForge database?.....	1018
Advantages of using the Apache TIKa parser library for indexing.....	1018
How do I connect to the Teamforge Postgres database?.....	1019
How do I capture the output of "top" command?.....	1019
How do I use the TeamForge updater to manage backups of old versions of TeamForge?.....	1019
How do I resolve timeouts when calling web services?.....	1019
What resources can be searched on a TeamForge site?.....	1020
Locales and Languages in TeamForge.....	1020
Customization.....	1020
What elements of a site can I customize?.....	1021
How does TeamForge use stylesheets?.....	1021
Can I substitute my own images for the default TeamForge images?.....	1021
Can I use my own custom JavaScript on my site?.....	1021
Can I customize the web interface?.....	1021
Integrated applications.....	1021
How is an integrated application described?.....	1022
How does an integrated application interact with other TeamForge tools?.....	1025
How can I import tasks into TeamForge Task Manager?.....	1027
SCM commits in TeamForge with EventQ integration.....	1027
What is in the integrated application's configuration file?.....	1027
Lab Management.....	1028
Can I get more systems for my project?.....	1028
What happens when I move a machine between projects?.....	1028
What is an audit log?.....	1028
What is the correct procedure for modifying a hosted Lab Manager profile?.....	1029
What is port forwarding?.....	1029
What is involved in administrating profiles?.....	1029

Who controls which profiles can be used in a project?.....	1031
How are my project systems being utilized?.....	1032
What is host URL mapping?.....	1033
Why doesn't URL mapping work for me?.....	1033
How does host URL mapping compare with port forwarding?.....	1034
How are virtual guests different from physical machines?.....	1034
What is involved in migrating a virtual guest?.....	1036
What's a good way to read log information?.....	1036
Email.....	1036
Can I customize my site's email notifications?.....	1036
Can I specify an alternate email address?.....	1037
How can I check if port 25 is open?.....	1037
When a discussion forum is set up, do all members receive a notification mail?.....	1037
Does the project owner receive a mail when a discussion forum is started in his project?.....	1037
I am unable to edit a specific artifact via email, but I can via the web UI. Why is this?.....	1037
How do I set up a local alias via James?.....	1038
Does TeamForge support using /etc/aliases for local mail delivery?.....	1038
How can I stay informed about events on my CollabNet TeamForge site?.....	1038
How can I reduce the number of emails from monitoring a folder?.....	1039
Can TeamForge accept email for more than one domain?.....	1039
How do I configure TeamForge to send mail on a specific network adapter in a multi-NIC configuration?.....	1039
Reports.....	1040
Why do I get performance issues when retrieving flex field information from Datamart?.....	1040
What kind of objects can I create reports on?.....	1041
How do I connect to the Datamart?.....	1041
How do I change the time to run the ETL jobs?.....	1041
How can I check the status of ETL?.....	1042
What can I learn from a burndown chart?.....	1042
What can I learn from a capacity chart?.....	1043
How do I export reports?.....	1043
What can I learn from an Open by Priority chart?.....	1043
What can I learn from an Open vs. Closed chart?.....	1044
How does TeamForge deliver activity reports?.....	1044
Why do I get duplicate records in a tracker report?.....	1044
Subversion and CVS.....	1045
What's the difference between a 'managed' and 'unmanaged' CVS server?.....	1045
Subversion replication in TeamForge.....	1045
How do I move an existing CVS repository into TeamForge?.....	1046
How do I move an existing SVN repository into TeamForge?.....	1047
How do I enable Neon debugging in my Subversion client?.....	1047
How do I create a local backup of a remote SVN repository?.....	1047
Trackers.....	1048
Can I assign an artifact or task to a group of users?.....	1048
How do I create an artifact through email?.....	1048
I don't need the priority and assigned-to fields, how do I get rid of them?.....	1048
Why do the open tracker counts differ from when I filter on 'Open'?.....	1048
How do I edit the predefined data in the various fields in the issue tracker?.....	1048
Can I limit the number of characters for fields in a tracker?.....	1048
Is it possible to move an artifact from one tracker to another?.....	1048
How can I change the display order of the fields in my tracker?.....	1049
How do I create a parent or a child artifact?.....	1049
What does the status of a planning folder mean?.....	1049
How is a project template structured?.....	1050
How do I create new values in the 'Reported in Release' field?.....	1050
What fields can I use in a tracker?.....	1050
What happens when a changed value makes a dependent field invalid?.....	1051

Discussions.....	1052
Is the discussion forum creator subscribed by default?.....	1052
How can I monitor a forum or mailing list?.....	1052
How do I create a new forum topic?.....	1053
How do I find the email address for a forum?.....	1053
How do I remove a user from a discussion?.....	1053
How do I post a reply to a forum topic?.....	1053
I can see the message posted to a discussion in the web UI, but I didn't receive any of the responses through email. Why is this?.....	1053
Why are some of the discussions threaded?.....	1054
What happens when I post to a moderated discussion forum?.....	1054
Who can be a moderator?.....	1054
Can I move a discussion to another project?.....	1054
Do project owners get automatically subscribed to the discussion forum started by another member?.....	1054
Can I subscribe to a TeamForge discussion forum's mailing list through an email?.....	1054
Why would I want to make a discussion forum moderated?.....	1055
Documents.....	1055
How does TeamForge support documents?.....	1055
Why can't I edit a document when it opens in my browser?.....	1056
Why are some uploaded documents missing icons when displayed in TeamForge?.....	1056
Can I link to documents outside of TeamForge?.....	1057
Can I lock a document in TeamForge?.....	1057
Does TeamForge automatically resolve conflicts in documents made by multiple concurrent editors?.....	1057
What document types are supported in TeamForge?.....	1057
Why doesn't an open review automatically close when a new version of the document is uploaded?.....	1057
Who can work with documents?.....	1058
Can I set permissions so that users can move documents but not delete them?.....	1058
Why am I not able to access a folder in TeamForge documents?.....	1058
Are role-based permissions allowed for subfolders in the TeamForge Documents?.....	1058
Can I set permissions for documents alone?.....	1058
Wiki.....	1058
How do I generate a wiki table of contents?.....	1058
How do I delete an attachment from a wiki page?.....	1059
How can I detect orphan wiki pages?.....	1059
How do I edit the wiki home page?.....	1059
How do I make the version comment required for wiki updates?.....	1059
I set my project to 'use wiki homepage' so why isn't my wiki showing up?.....	1060
Why would I attach things to a wiki?.....	1060
<b>TeamForge 17.4 release notes.....</b>	<b>1061</b>
New features in TeamForge 17.4.....	1061
Issues resolved in TeamForge 17.4.....	1068
Known issues in TeamForge 17.4.....	1068
<b>CollabNet TeamForge 17.4 Update 1 release notes.....</b>	<b>1070</b>
<b>CollabNet TeamForge 17.4 Update 2 release notes.....</b>	<b>1071</b>

# TeamForge 17.4

---

CollabNet TeamForge<sup>®</sup> is the industry's only open and extensible collaborative software delivery platform for distributed teams. It enables organizations to improve collaboration, agility and governance by connecting teams, processes, IP and tools via a centralized, secure, web-based system. With TeamForge, your entire organization can collaborate and become more agile at scale. And it scales to tens of thousands of users.

## How to use TeamForge 17.4

---


TeamForge makes it easy and fun to manage all phases of the software lifecycle, from requirements through release. You get software configuration management (SCM), continuous build and test, issue tracking, project management, lab management, and collaboration tools, all in a single tightly-integrated, cloud-based platform that empowers distributed teams to deliver great software.

For highlights of the many new features that TeamForge 17.4 offers, take a look at the [release notes](#).

For downloading this Help as a PDF file, see [Downloadable help](#).

## Get started

Take these steps to get started quickly and efficiently on your TeamForge 17.4 site.


 **Note:** If you are coming to TeamForge from a CollabNet Enterprise Edition site, see [Moving to a CollabNet TeamForge project from a CollabNet Enterprise Edition project](#) on page 1011 for guidance on how to keep doing what you were doing on your old site.

### Quick start: Working on a TeamForge project

Jump right in! Join a project, check out the code, meet your collaborators.

#### Join a TeamForge site

To collaborate with others on a CollabNet TeamForge site, start by getting a user account.


 **Note:** You need a license to use a TeamForge site. Your site administrator may have already assigned you a license. If not, you'll be asked to choose the type of license you need when you create your account.

- An ALM license enables you to use the full set of tools on the site, including Source Code, Trackers, Documents, Tasks and Reports.
- An SCM license enables you to use the site for basic source code management. You can browse repositories and participate in Wiki conversations.
- A DevOps license enables you to use Package Management ( Application & Environment ) and FRS.
- A Version Control license enables source code management , FRS and Review tool.

1. If your site authenticates through LDAP, follow these steps.

If your site does not authenticate through LDAP, skip to the next set of steps.

a) In the **Log Into CollabNet TeamForge** section of the site home page, enter your corporate LDAP user name and password.

 **Tip:** In most cases, your user name and password are the user name and password with which you log in to your corporate network.

b) Click **Log In**.

c) On the **Create Account** page, re-enter your LDAP password.

d) Enter your full name and email address and click **Create**.

2. If your site does not authenticate through LDAP, follow these steps.

a) Click **Create an Account** in the **New Users** section of the TeamForge home page.

b) On the **Create New Account** page, enter a user name for your account.



**Note:** Your user name must meet these criteria:

- User name is case sensitive.
- Minimum number of characters as specified in the `site-options.conf` file.
- No spaces.
- Should have at least one letter.
- The first character is a letter.

c) Enter and confirm a password.

d) Fill in the rest of the fields and click **Create**.

3. Check your email at the address that you provided.

4. Follow the link in the email to the **Log In** page.

5. Enter your user name and password, and click **Log In**.

Your TeamForge account is now activated.

## My Workspace

Your **My Workspace** is a personal workspace that offers a handful of configurable widgets such as "My Recent Projects", "My Recent Repositories", "Git Code Reviews", "Project News" and so. You can use these widgets to view recent projects, recent repositories, recent commits, items assigned to you (TeamForge artifacts and document reviews), Git code reviews, project news and reports.

## Getting there

After logging into TeamForge, you are taken to your **My Workspace**.

## My Workspace dashboard

You can create one or more dashboards to categorize and view data in many different views. For instance, you may decide to create three dashboards: Dashboard A, B and C. With these three dashboards, you may choose to configure Dashboard A and B to view events from projects A and B respectively and configure Dashboard C to view just the code reviews.

My Workspace | My Page | Jump to ID | Ex: artf2345

Dashboard C

**My Recent Projects**

- RTFM
- TeamForge Helpdesk
- Customer Issues

**My Recent Repositories**

- output

**My Recent History**

- cmmt370188: Commit by Meenakshi Kamaraj (kmeenakshi)
- cmmt370187: Commit by Meenakshi Kamaraj (kmeenakshi)
- cmmt370184: Commit by Meenakshi Kamaraj (kmeenakshi)
- cmmt370183: Commit by Meenakshi Kamaraj (kmeenakshi)
- cmmt370132: Commit by Meenakshi Kamaraj (kmeenakshi)

**My Artifacts**

ID	TITLE	STATUS	CREATED BY
artf289835	Doc task for artf236905 and artf289686	Open	Arun Pand
artf289235	Request to document tracker artifact limit in HCN	Open	Prem Kum
artf288732	Docs: Add step to remove any add-ons prior to upgrade	Open	Richie Jar

**New My Workspace Dashboard**

Your My Workspace dashboard can be customized to suit your own needs. Click the customize icon (🔧) at the top of the page to get started.

The dashboard consists of three main things:

1. A name, which you can update via the text box at the top of the page.
2. A layout, which you can choose from a number of page layouts, which give you options in how to arrange your dashboard.
3. Widgets, which provide the functionality you want to expose on your dashboard. To add a widget just click the **Add widget** button at the top of your dashboard. A variety of widgets have been provided, including this one which is a Markdown widget that allows you to add links and images to your dashboard. Most widgets have criteria or content that controls what information they display. These settings are controlled by clicking the widget configuration icon (⚙️) on the widget.

Arrange widgets by clicking on the widget location icon (📍) and dragging the widget to the place in your layout that you want to see it.

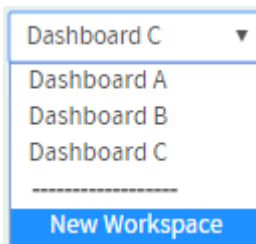
When you are satisfied with your changes click the **Save** button.


There is a **Cancel** button if you do not want to keep your changes, and the **Use default** button will return the dashboard to the default settings.

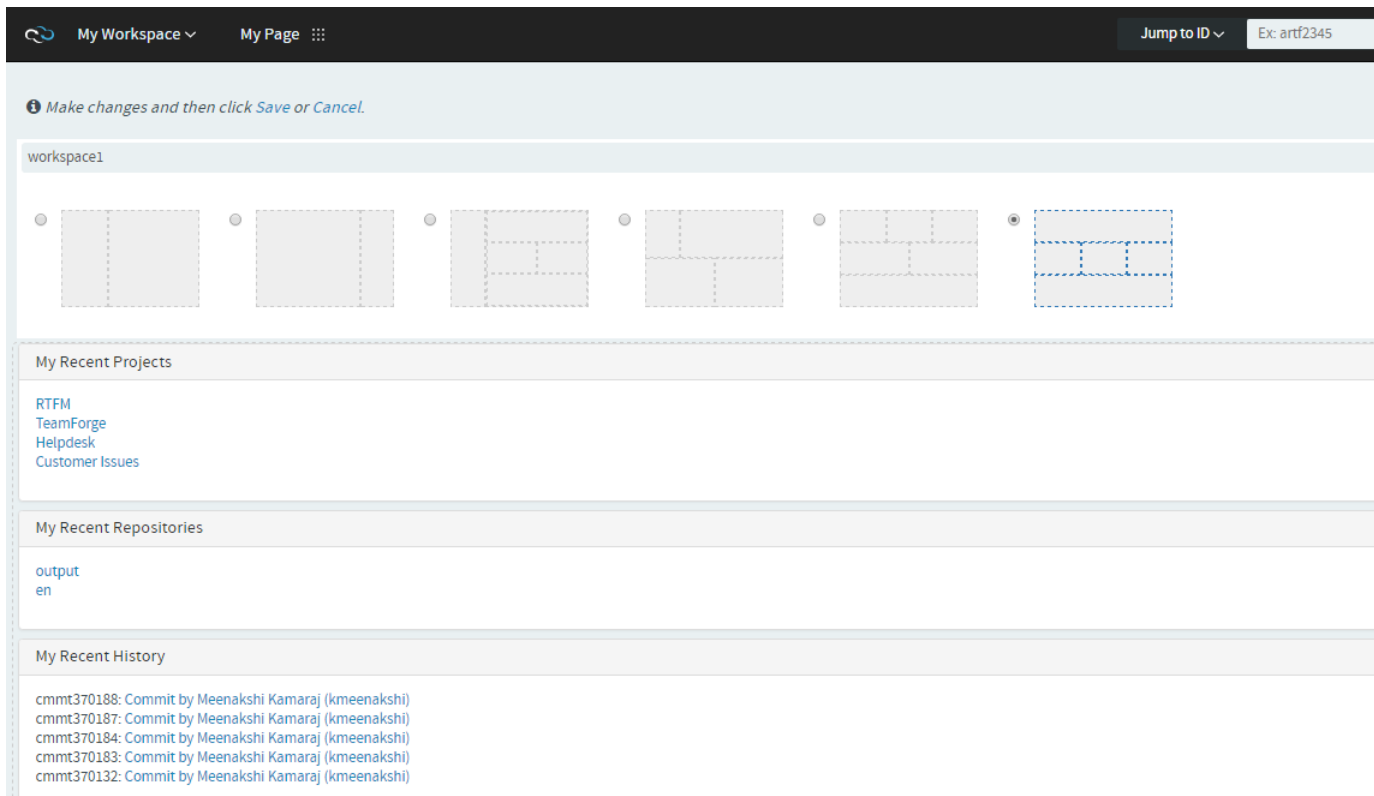
Feel free to delete or edit the content of this widget once you understand the basic concepts.


## Create and Configure Dashboards: Step-by-Step

1. To create a new dashboard, click New Workspace from the workspace drop-down on your dashboard.





2. Select a dashboard from the drop-down list and click the customize () button.



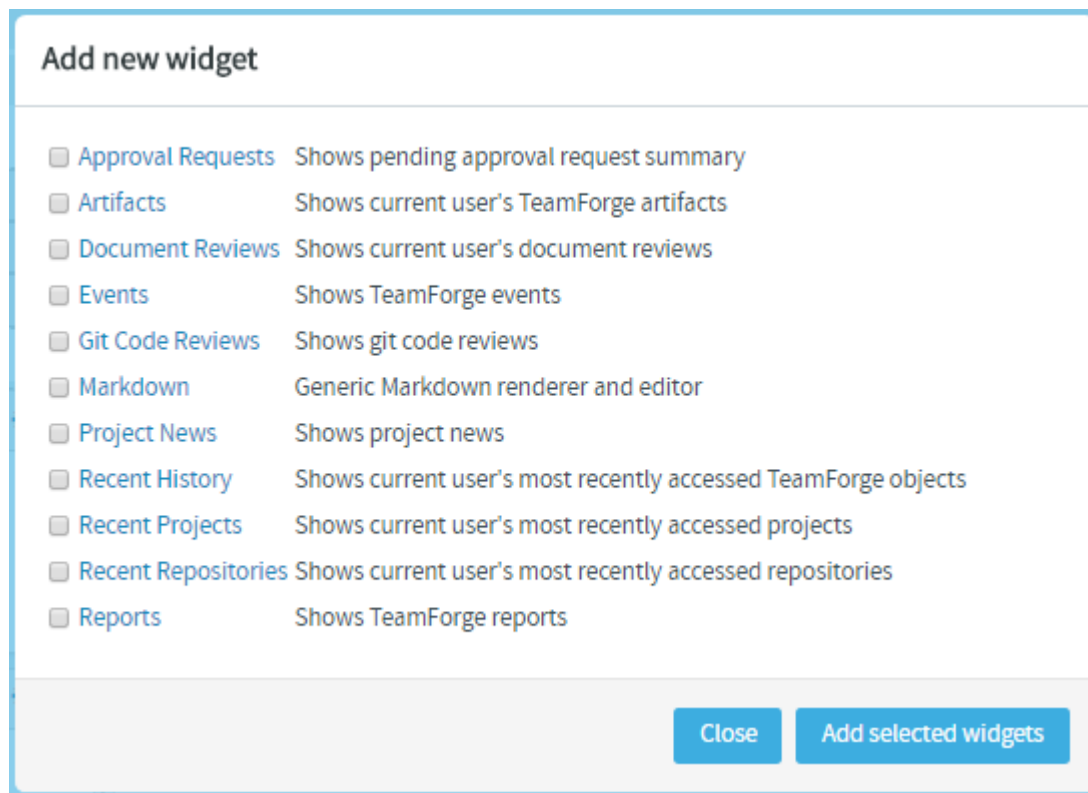
3. Enter the workspace name.
4. Select a suitable layout that should accommodate your widgets appropriately.
5. Click **Add widget** and select a widget. For details about each widget, see .
6. Click **Add selected widgets** to get the selected widgets added to your dashboard. Click **Close** to get back to your dashboard without adding widgets.
  -  **Note:** You can add multiple instances of the same widget to your **My Workspace**. For example, you can add two instances of the **Artifacts** widget, one each to show the "Open" and "Closed" artifacts assigned to you.
7. All widgets have criteria or content that controls what information they display. These settings can be configured by clicking the **Edit widget configuration** icon on the widget after adding them to the **My Workspace**.
8. Arrange widgets by clicking the **Change widget location** icon and dragging the widget to the place in your layout where you want to see it.
9. Use the **Collapse widget** and **Remove widget** icons to collapse and remove widgets.
10. Repeat steps 5 through 8 to add and configure more widgets.
11. When you are satisfied with your changes, click **Save** . There is a **Cancel** button if you do not want to keep your changes and the **Use default** button will return the dashboard to the default settings.

Create as many dashboards as required.

You can select a dashboard from the drop-down list and click the  button to delete the dashboard.



## More about widgets



A variety of widgets have been provided, which include the following:


- **Approval Requests:** Shows summary of pending approval requests
- **Artifacts:** Shows artifacts assigned to the user
- **Document Reviews:** Shows document reviews assigned to the user
- **Events:** Shows TeamForge events
- **Git Code Reviews:** Shows Git Code Reviews
- **Markdown:** Generic Markdown renderer and editor
- **Project News:** Shows project news
- **Recent History:** Shows objects most recently accessed by the user
- **Recent Projects:** Shows projects most recently accessed by the user
- **Recent Repositories:** Shows repositories most recently accessed by the user
- **Reports:** Shows TeamForge reports

## Join a project

To create, store, and share work on a CollabNet TeamForge site, first join a project.

1. Log into TeamForge .
2. From your **My Page** menu, click **Projects**.  
Now you're looking at the projects of which you are already a member, if any.
3. Click the **All Projects** tab to see all TeamForge projects, based on each project's access settings.

4. Select the project you want to join and click **Request Membership**.
5. On the **Join Project** page, explain why you want to join the project in the **REQUEST COMMENT** text box, and click **Submit**.  
Your request is submitted to the project administrator for approval. You will receive an email notification when your request is either approved or denied.


 **Tip:** You can also request project membership by clicking **Join this Project** on the home page of the project you want to join.

### Identify yourself


To help get started on collaborating with other project members, you can provide some information about yourself.

This information appears whenever someone clicks your name anywhere on the site. For example, when an artifact is assigned to you, your name appears as a link that other users can click.

1. Click **My Settings** from the **My Page** menu.
2. In the **User Details** section, click **Edit** and provide some information to help potential coworkers get to know you.
  - a) Click **Choose File** to upload a 100-by-100-pixel picture of yourself, or of something that suggests who you are.
  - b) Under **Detail**, provide a summary of your interests, skills, or other characteristics.
3. Select the language you prefer to use. You can choose to receive emails from your TeamForge site in English, Chinese, Japanese, or Korean.

 **Note:** If your browser is set to use one of the available languages, you should already be seeing TeamForge in that language in your web browser.

4. If necessary, review and change your password, official name, and email address.

 **Note:** This is the email address where you receive alerts about changes to items you are involved with, such as a discussion forum you are monitoring or a code commit associated with an artifact you created.

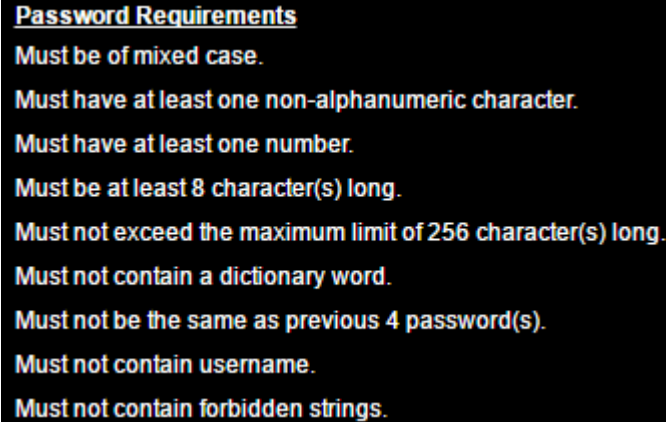
5. Click **Save**.

### Change your password

To protect the security of your TeamForge account, change the password regularly.

1. Select **My Settings** from the **My Page** menu.
2. In the **User Details** section, click **Change Password**.
3. In the **My Workspace / Change Password** page, enter your current password.
4. Enter your new password and confirm it.

The password requirements depend on the [security policies enforced on your site](#). Here's an example:



**Password Requirements**

- Must be of mixed case.
- Must have at least one non-alphanumeric character.
- Must have at least one number.
- Must be at least 8 character(s) long.
- Must not exceed the maximum limit of 256 character(s) long.
- Must not contain a dictionary word.
- Must not be the same as previous 4 password(s).
- Must not contain username.
- Must not contain forbidden strings.


5. Click **Update**.  
You will see a message that your password was changed successfully.


### Get the code

Browse TeamForge to find the code you want to work on, then check out the code.



You can view the contents of each file in a repository, plus additional information about each file such as revision history, comments, date and time of submission, and branch and tag information. You can also view differences (diffs) between any two files. You can also search for code in a repository using TeamForge Code Search.

 **Restriction:** You can see only those paths in the repository that the repository administrator has given you access to.

 **Note:** If you're getting code from a Subversion replica repository, the TeamForge account used while setting up the replica determines what's available to be checked out. This account could have been provided total access to the master repository, or restricted access using path-based permissions.


1. Click **SOURCE CODE** from the **Project Home** menu.
2. On the list of project repositories, click the name of the repository in which you want to view code. For each file, the revision number, time since check-in, author, and last log entry appear in the **Repository Browser**.
  - To view a file, or to view the diffs between two files, click the file name.
  - To view a specific version of the file, click **Download**.
  - To view the differences between two files, do either of these:
    - Click **[select for diffs]** next to each of the two files that you want to compare.
    - Enter the file revision numbers in the Diffs between boxes at the bottom of the page.
3. If you need to diff files, choose a display from the **Type of Diff** menu, then click **Get Diffs**. The differences between the two files are displayed.
4. Use your source control client to check out the code to your local machine.

#### Internal code browser

For Subversion and Git repositories, you have the option to use the TeamForge code browser which is turned on by default while integrating the source code server. For more information, see [Integrate a source code server](#).

5. On the list of project repositories, click the name of a Subversion or a Git repository in which you want to view code.

On the top right of code browser, you can select the branch/tag (for Git) or specify the revision (for SVN) you want to browse.

- **View:** This tab allows you to do the following:
  - Browse through the folder hierarchy of the repository and view the content of specific files. For any folder or file you are viewing within a branch (Git) or revision (SVN), you can obtain the commit information pertaining to its last update.
  - While viewing a single specific commit or a file, you can see the paths that were modified in that commit, the associations including JIRA such as builds, code reviews and so on (from EventQ events) for the specific commit and the differences between files in that commit.
    -  **Important:** To view the associations, you must have installed EventQ and must have RBAC (role-based access control) permissions to use the "EventQ READ" or "Reporting API" of TeamForge EventQ. If either of these requirements is not met, this section will not show up at all.
  - While viewing a folder, if there is a file named `readme`, `readme.txt` or `readme.md` that file will automatically be rendered beneath the list of files in the folder. If the file contains markdown formatting, it will be rendered as rich text.
- **Changes:** This tab lets you view all of the commits that touched a specific path you are browsing within a branch or revision. Click a commit to view its details.
- **Graph:** This tab provides a graphical representation of the changes made including branching and merging of repositories.
- **Branches** (for Git): This allows you to see all of the branches in the repository in their relation to the default (master) one. Using **Compare Branch** you can see the commits in the branch that do not exist in the default branch.
- **Tags:** This tab lets you create Git tags and tag specific points in history as being important. Typically you can use this functionality to mark release points (v1.0, and so on) with an option to add Release Notes for the

tagged revision. Once you create a tag, you can use it to download source code as a zip/tar file and view the tag information in **Changes** and **Graph** tabs.

The first screenshot shows the 'Create Tag' form with fields for 'TAG NAME:', 'INITIAL REVISION:', and 'NOTES:'. A 'Create Tag' button is at the bottom right.

The second screenshot shows the 'TAGS' tab with a table of tags:

Tag Name	Time Ago	SHA-1	Actions
build/17.1.522	1 day ago	b92f308	zip tar.gz
build/17.1.516	2 days ago	6730baa	zip tar.gz
build/17.1.514	2 days ago	4266727	zip tar.gz
build/17.1.513	2 days ago	9f46266	zip tar.gz
build/17.1.512	3 days ago	3e360e3	zip tar.gz
build/17.1.511	3 days ago	59c7869	zip tar.gz

The third screenshot shows the 'Commits on 01/27/2017' section with two commits:

- Commit b92f308: Add validation to moveDocument3() SOAP API. Authored 01/27/2017 4:40 am PST (2 days ago) by Subhashini Marimuthu. Includes tags #8329, master, and build/17.1.522.
- Commit 0322dae: [artf286666] Planning Folder REST API is not working. Authored 01/26/2017 10:51 pm PST (2 days ago) by Annapoorni Dakshinamurthy. Includes tag #8328.

The fourth screenshot shows the 'GRAPH' tab with a commit graph and a list of commits:

- b92f308 2 days ago by Subhashini Marimuthu: Add validation to moveDocument3() SOAP API. Includes tags #8329, master, and build/17.1.522.
- 0322dae 2 days ago by Annapoorni Dakshinamurthy: [artf286666] Planning Folder REST API is not working. Includes tag #8328.
- a98573a 2 days ago by Rajesh Vadivelu: Merge "[artf286653] : Column Configurations are Not saved,"
- 7b81c82 2 days ago by Rajalakshmi Madasamy: [artf286585] : Permission denied to view documents when logged in user has permissions
- 53448d8 2 days ago by Rajesh Vadivelu: [artf286653] : Column Configurations are Not saved, Includes tag #8323.

- **Reviews:** This tab lists all the Open, Merged and Abandoned reviews, both Pull Requests and Gerrit single-commit reviews. Pull requests allow developers to collaborate with each other on a code change before merging it into another branch on a GIT repository. You can access this tab only when the repository owner has enabled this feature. For more information, see [Pull request step-by-step](#) on page 100.
- **Support for both Pull Requests and single-commit Gerrit Reviews:** Supports all types of code review policies, which include Pull Requests and single commit Gerrit Reviews.

VIEW CHANGES GRAPH BRANCHES TAGS **REVIEWS** SEARCH SETTINGS

Open Merged Abandoned Includes both Gerrit and Pull Request reviews

ID	Title	Branch	Reviewer
#5061	[artf255076] updated title :)	master	Miikka Ander
#5050	[artf260665] Add test cases for server controller	master	Stephen Else
#5057	[artf271495] testing testing	master	Miikka Ander
#5055	Use rest api to delete associations	master	Baybora Akse
#5042	[artf262520]: Deprecate *_REPOSITORY_BASE and move default rep...	master	vetrivel vaith
#5025	[artf275684]: Basic changes to list comments [artf255243]: Paginatio...	master	Rakesh Kuma
#5012	[artf274066]: backend changes for Wiki component	master	Aishwarya Ra
#5043	[WIP] [artf260017] Upgrade to Wildfly 10.0	master (wildfly10-upgrade)	Johannes Pa

Gerrit code review

Pull Request review

**Auto refresh when a Pull Request changes:** When a pull request changes, the page is automatically refreshed to reflect the changes.

- **Comments to support @mentions:** Inline comments are parsed for @mentions and users called out via @mentions are added as reviewers.
- **Open your Gerrit Reviews in Gerrit's user interface:** A new button has been added to let you open your Gerrit Reviews in Gerrit's user interface.

VIEW CHANGES GRAPH BR

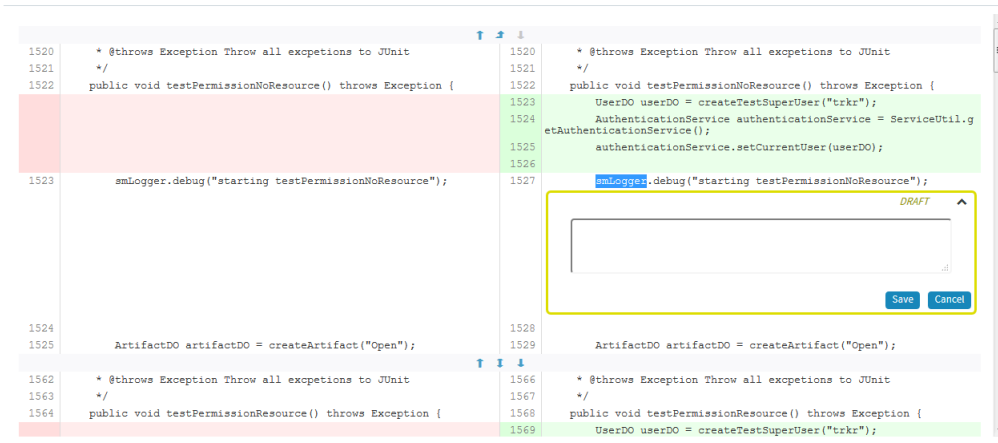
**K** #4456 - [artf263918]- Package Compon

OPEN refs/changes/56/4456/1 → m

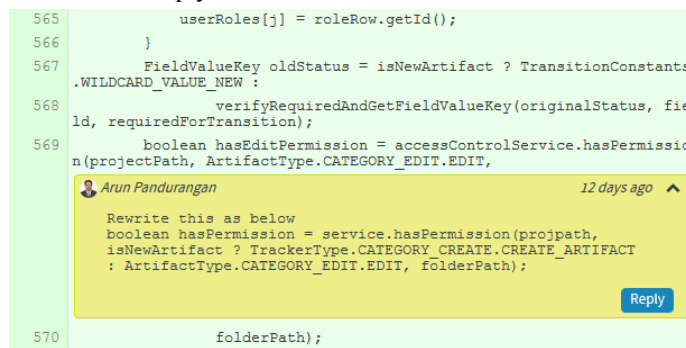
Messages 3 Files 5

View this code review in Gerrit

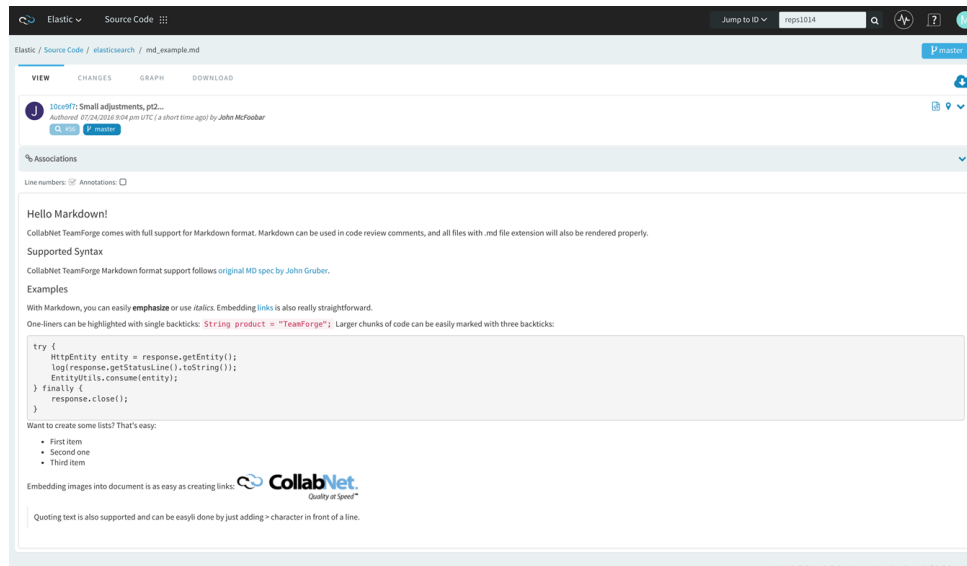
- **Code commenting:** During code reviews, you can now add line comments in context while looking at the files in diff view. You can double-click to block a line/text and add a comment.



You can also reply to line comments.



- **Ability to diff the change against the Base or a previous Patch Set:** As part of the Gerrit review workflow, you now have the ability to diff the change against the Base or a previous Patch Set.
- **Markdown support:** Markdown support for all .MD files: Render Markdown files when viewed through Code Browser



- **Mass delete/resurrect options:** Mass delete/resurrect options in History Protect tab:

Elastic / Source Code / elasticsearch

VIEW CHANGES GRAPH **BRANCHES** REVIEWS SETTINGS

Default branch: master

93cc3e4: testing pr  
Authored 06/13/2016 8:22 pm UTC (1 month ago) by John McFoobar

Active Merged All Deleted/Rewritten

	TYPE	REFERENCE NAME	USER	REWRITE DATE	ORIGINAL COMMIT ID	NEW COMMIT ID
<input checked="" type="checkbox"/>		build_update	john	03/07/2016 4:45 pm UTC	c310cc3	
<input type="checkbox"/>		conflict_branch	mike	03/07/2016 8:40 pm UTC	265edc5	
<input checked="" type="checkbox"/>		license_update	john	03/07/2016 7:11 pm UTC	2e42c1c	
<input type="checkbox"/>		master	admin	03/07/2016 2:13 pm UTC	5194590	050bfc3

Resurrect Delete Permanently

© 2016 CollabNet. CollabNet is a registered trademark of CollabNet, Inc.

Elastic / Source Code / elasticsearch

VIEW CHANGES GRAPH **BRANCHES** REVIEWS

Default branch: master

93cc3e4: testing pr  
Authored 06/13/2016 8:22 pm UTC (1 month ago) by John McFoobar

Active Merged All Deleted/Rewritten

**Resurrect A Rewritten Reference**

Do you want to resurrect 2 selected references to their original name?

Cancel Resurrect

	TYPE	REFERENCE NAME	USER	REWRITE DATE	ORIGINAL COMMIT ID	NEW COMMIT ID
<input checked="" type="checkbox"/>		build_update	john	03/07/2016 4:45 pm UTC	c310cc3	
<input type="checkbox"/>		conflict_branch	mike	03/07/2016 8:40 pm UTC	265edc5	
<input checked="" type="checkbox"/>		license_update	john	03/07/2016 7:11 pm UTC	2e42c1c	
<input type="checkbox"/>		master	admin	03/07/2016 2:13 pm UTC	5194590	050bfc3

Resurrect Delete Permanently

© 2016 CollabNet. CollabNet is a registered trademark of CollabNet, Inc.

- **Inline editing of files:** Quick changes to files, if required only to few files, can be done using the inline edit feature from within the code browser without having to clone an entire repository. Browse the repository, locate and open the file in the **View** tab, click **Edit** to open the file in the **File Editor**, make your changes, **Create code review** and **Publish** your changes for review.

The screenshot shows the TeamForge interface. At the top, there are navigation tabs: VIEW, CHANGES, GRAPH, BRANCHES, TAGS, REVIEWS, SEARCH, SETTINGS, and DOWNLOAD. Below this, a commit is displayed with the message "4266727: [artf285315] Updated documentation to reflect HAProxy changes". The commit is authored by Johannes Passing. Below the commit, there are buttons for "#8236" and "build/17.1.514". Underneath, there is an "Associations" section. At the bottom of the commit view, there is a "Show diff against:" dropdown menu with "No commit selected" and an "Edit" button. A red box highlights the "Edit" button. Below the commit view, a "File Editor" window is open for the file "README.md". The editor shows the following content:

```

11 - ### Cloning the repository
12
13 Go to [the repository viewer of the TF Git repo]
    (https://forge.collab.net/ctf/code/git/projects.teamforge/scm.TeamForge_C
  
```

At the bottom of the file editor, there are three buttons: "Preview", "Cancel", and "Create code review". The "Create code review" button is highlighted with a red box. A red line connects the "Edit" button from the commit view to the "Create code review" button in the file editor.

- **Submit whole topic:** You can now bundle related changes (code reviews) by topic and submit the whole topic for review instead of just submitting changes one-by-one. Just open a review, click the **Set Topic** link and enter the topic name.

The screenshot shows the TeamForge interface with a code review open. The review title is "#7969 - [artf284935] Enabled FindBugs for 'core' module". The review is in "OPEN" status. Below the title, there is a link for "refs/changes/69/7969/10" and a link for "master" with a "(Set topic) (no topic)" link next to it. The "(Set topic) (no topic)" link is highlighted with a red box. Below the review, there is a list of changes, including "Enabled plugin" and "Fixed code/added suppressions". At the bottom of the review, there is an input field for the topic name, which is highlighted with a red box. A red line connects the "(Set topic) (no topic)" link to the input field.

- **Search:** This tab lets you search for code via TeamForge Code Search powered by *Elasticsearch*. You can search all files in a repository or narrow your scope to specific file types such as C, C++, C# and so on. Type your search keyword, select a file extension (optional) and click **Search**. For more information, see [Search code](#) on page 107.

The screenshot shows the TeamForge interface with the "SEARCH" tab selected. At the top, there are navigation tabs: VIEW, CHANGES, GRAPH, BRANCHES, TAGS, REVIEWS, SEARCH, and SETTINGS. Below this, there is a search input field with the placeholder text "Search this repository". To the right of the input field, there is a dropdown menu for "Extensions..." and a "Search" button. The "Search" button is highlighted with a red box. Below the search input field, there is a small text indicating the last indexed revision: "Last indexed revision: 9cbfcbf 04/18/2016 3:27 am PDT (9 months ago)".

- **Settings:** This tab lets you configure the repository settings.
- **Download:** This tab allows you to download a copy of the required file.

## Find a forum

Read through a project's discussion forums to find out about the project's plans, accomplishments, and challenges.

1. Click **Discussions** from the **Project Home** menu.  
The **Forum Summary** page displays a list of all project forums, and summary information about each forum's activity.
2. Click a forum name.  
The **Topic Summary** page displays a list of all forum topics and summary information about each forum topic's activity.
3. Click a topic name.  
The **Forum Topic** page displays all messages that have been posted in response to the original forum topic.

You can configure your view in the following ways:

- To set the number of messages that display on each page, choose a value from the **View** menu, then click **Update**.
- To display threaded messages, choose one of the following options, then click **Update**.
  - **Nested - by Reply** - Displays threaded messages in a hierarchical, nested format of up to 20 layers. This is the default display.
  - **Flat - Chronological** - Displays all messages in chronological order with no nesting.

## Catch up on project news


Watch project news updates to stay up to date on events that might affect your participation in the project.



News items appear on the project home page, on the **News** tab of each user's **My page**, and on the CollabNet TeamForge home page.

1. Go to the project home page.
  - From within the project, click **Project Home**.
  - From anywhere else in CollabNet TeamForge, choose the project from the list of projects available under **My Workspace > My Projects**.
2. Look for news items at the bottom of the project home page.  
Each news item lists its author and the date it was posted.

## Find project resources

When you want to search for TeamForge objects such as tracker artifacts or documents, you can quickly search using a unique identifier or a keyword or you can perform an advanced search.

-  **Tip:** In addition to the **Jump to ID** search, you can select an object category from the drop-down list and search within the selected object category. You can also select **Advanced Search** from the drop-down list to perform an advanced search.

1. Log on to TeamForge.
  -  **Note:** If you are not logged on, you can search only projects and items that have been designated public.
2. If you know the unique identifier of an object and want to quickly go to the object, type the unique identifier in the **Jump to ID** text box and click the search icon.
  -  **Note:** The default quick search option is **Jump to ID**.
3. If you want to do a keyword search of a specific object type (such as documents or discussions), type the keyword in the text box, select an object type from the drop-down list and click the search icon.


The following table lists the search-able object types you can select from the drop-down list.

Searchable object types	Description
<b>Discussions</b>	Select this option to search in discussion forums.
<b>Documents</b>	Select this option to search for documents.

Searchable object types	Description
<b>File Releases</b>	Select this option to search in file releases.
<b>News</b>	Select this option to search project news.
<b>Project Pages (Visible)</b>	Select this option to search project pages.
<b>Projects</b>	Select this option to search projects.
<b>Source Code</b>	Select this option to search the source code. For more information, see <a href="#">How to search for Source Code</a> on page 108.
<b>Tasks</b>	Select this option to search for tasks.
<b>Trackers</b>	Select this option to search for tracker artifacts.
<b>Users</b>	Select this option to search for users.
<b>Wiki</b>	Select this option to search in Wiki pages.

- If you want to perform an advanced search, select **Advanced Search** from the drop-down list.
- On the **Search Criteria** page, enter the keywords to search for.
- Select the **Search Attachments** check box and the **Search Comments** check box if you want to search attachments and comments respectively.  
Attachments refer to tracker artifact attachments. Comments refer to tracker artifact comments and task comments.
- Select one of the two options, **Search Active Versions Only** or **Search All Versions**, to specify whether you want to search active document versions only or all document versions respectively.  
Searching only active document versions allows you to eliminate search results for outdated documents.
- Click **Search**.

Your search results are organized by TeamForge application. The search score indicates the relevance of each result to your search criteria.

 **Note:** You can see only those items that your project membership and permissions allow you to see.

### Filtering tables

Filtering lets you view a subset of a list of items that match your filter criteria.

The TeamForge application uses tables for presenting list of items such as tracker artifacts, documents, tasks, and so on. You can filter a subset of a list of items that meets your filter criteria on all the tables with filter fields.

- Specify the filter criteria in one or more filter fields (at the top of each column) and click **FILTER**.
  - You can find a filter field at the top of each column in most of the tables in the TeamForge application.
  - The filter field could be a text box or a drop-down list with multi-select check boxes.




Stories Page 1 of 40 (600 Items) FILTER  COLUMN

PRIORITY	ARTIFACT ID : TITLE	ASSIGNED TO	TEAM	SUBMITTED BY
Any	Filter by	Filter by	Any	Filter by
3 MEDIUM	artf284437 : Create UI for viewing all diffs for entire review on one screen	C. Michael Pilato	SCM	Mark Phippard
3 MEDIUM	artf277197 : Workspace widget for artifacts	Baybora Aksoy	SCM	Mark Phippard
4 LOW	artf246564 : Reports - write protractor tests to cover e2e tests.	None	Reports (Datamart)	Muhammed Ashik
2 HIGH	artf287743 : Include 'Artifact Tags' in Project Templates	Ramkumar Rathinasamy Viveganandan	TF Core > Avengers - Tags	Aishwarya Rajan
2 HIGH	artf287742 : Tags in Edit Artifact Page	Ramkumar Rathinasamy Viveganandan	TF Core > Avengers - Tags	Aishwarya Rajan
2 HIGH	artf283822 : Synchronize SVN Indexer execution	akil roy	Devops/Releng	Johannes Passing
2 HIGH	artf278265 : Cleanup the old EventService	Nithya Pattabiraman	TF Core > Jaguars - Architecture	Sriram
3 MEDIUM	artf288250 : EventQ meta RPM	Sinhuja Harikrishnan	Devops/Releng	Sinhuja Harikrishnan
3 MEDIUM	artf287737 : Update the default/trial license to include DevOps (and perhaps new license types)	Gowthaman Ravindran	Mavericks - CRUS	Mark Phippard
3 MEDIUM	artf287193 : SVN load balancing	akil roy	Devops/Releng	Johannes Passing
3 MEDIUM	artf286210 : Monit Snapshot before restart	h.karthik	Devops/Releng	Johannes Passing

1 2 3 4 5 6 7 8 9 10 >

Monitor Import Export Cut Edit Edit Inline Mass Update

- You can type your filter criteria in the text boxes. The search text is case-insensitive.
- You can also select the filter values from one or more drop-down lists. By default, you can only select up to 10 filter values in a drop-down list. However, you can set a value that suits your requirement for the `FILTER_DROPDOWN_MAX_SELECTION` token in the `site-options.conf` file to increase or decrease the count.
- Filter-as-you-type:** You can find the **Enter keywords** text box in all filter drop-down lists. As you type your filter keyword, instant search results are shown in the drop-down list. For example, in the following illustration, typing "R" instantly shows all statuses having the alphabet "R". The search text is case-insensitive.



- Some search filters may not appear if your site administrator has not enabled them.
- After filtering, if you want to save a filter for future use:
    - Click **FILTER** and select **Save** from the drop-down list. The **Save Filter As** window appears.
    - Type a name for the filter in the **FILTER NAME** text box.
    - Click **Save**. The filter is saved. You can view and select the saved filters at a later point in time by clicking the **Filter** drop-down list.

**Note:** You can save filters only in specific contexts. This feature may not be available in all the tables where you can filter table list items.
  - After filtering, if you want to clear the filters, click **FILTER** and select **Clear** from the drop-down list.
  - To delete saved filters:
    - Click **FILTER** and select **Delete** from the drop-down list. The **Select Filters To Be Deleted** window appears.

b) Select one or more filters to delete.

 **Tip:** Press and hold the **Ctrl** key to select more than one filter.

c) Click **Delete**.

A message such as `2 tracker filter(s) have been deleted successfully.` is displayed if the process was successful.

### Quick start: Managing a TeamForge project

Jump right in! Set up your project site, add your project members, plan out your work.

#### Create a TeamForge project

Create a new project when you have identified work to be done that has its own distinct character, dependencies or schedule.

What constitutes a project depends on your organization. Some organizations favor a small number of big, centralized projects. Others prefer a larger number of smaller, specialized projects. Your site administrator can help you decide if your work should be part of a larger project or a project of its own.

1. Access the **Projects** page through either of the following ways:

- Go to **My Workspace > View All Projects**
- Go to **My Page > Projects**.

2. Click **Create New Project**.

3. On the **Create Project** page, give the project a name and a brief description.

- The name will appear in project lists and on the project's home page.
- A terse description is recommended. There will be unlimited room to discuss the project's aims and methods in detail on the project pages themselves.

4. Provide a URL name for the project, if you want the URL for the project to be different from the internal project name.

If you do not enter a URL name, the project URL will be the same as the project name.

5. If your site administrator has provided project templates, select the appropriate one for your new project.

Project templates give you ready-made artifact types, work flow support, user roles and other start-up content appropriate to the kind of project you are creating.


6. Click **Create**.

The project is submitted to the CollabNet TeamForge site administrator for approval. You will receive an email notification when the site administrator approves or rejects your project. When your project is approved, you can get to it from your **My Projects** page or from the **Projects** menu in your navigation bar.

#### Add users to a project

Before a person can work on a project, you have to make him a member of the project.

You can make any registered user on your CollabNet TeamForge site a project member. You can assign roles to the user at the same time.

 **Note:** Project members, roles and the associated permissions can be inherited via project hierarchy and reused in subprojects. If your project is a subproject of any other project, you may have inherited some roles, project members or user groups.


1. Click **PROJECT ADMIN** from the **Project Home** menu.

2. On the **Project Admin** menu, click **User Membership**.

3. Click **Add**.


4. On the **Add User** page, find the users you want by one of these methods:

- Under **Search for Users**, filter the list of site users eligible to join this project. You can filter by full or partial name or user name.

 **Note:** Search text is not case-sensitive.


- Browse the list of registered users on the site. Sort them by name, user name, email address or membership status.

The inherited project members continue to hold the inherited roles with corresponding permissions, as specified in the source projects.

 **Note:** If a site has a great many users, you must filter them first to narrow down the list. This helps avoid slowing down the system.

5. Select the users you want to add.

6. Under **Assign Roles**, select the roles you want the users to have.

 **Tip:** If you prefer, you can skip this step and assign roles later on the **Project Admin > Permissions** page. Note that using the **Assigned Project Members** page, you can assign roles only to the direct project members.


7. Save your changes.

- Click **Save** to return to the **Project Membership** page.
- Click **Save and Add More** to keep adding users.

### Create a project page

To provide information and functionality to people viewing your project, build one or more project pages.

1. Go to the page to which your new page will belong.

 **Tip:** Any project page can have sub-pages belonging to it. A page that belongs directly to the project home page is called a "top-level page."

2. Click **Configure: On**.

3. Choose where your new page will fit in your project's structure.

- To create a page just under the project home page, click **Add top-level page**.


A top-level page's title is always visible in the navigation tree at left.

- To create a page under the page you are on right now, click **Add sub-page**.

4. Give your new project page a title.

Keep the title brief and descriptive.

5. Choose who can see this page.

 **Note:** Your choice will apply to all subpages that you create under this page.

- To show this page to anyone with the necessary permissions, select **Visible**. For example, if you have defined a group of users who have access to your project, your new project page is visible only to those users. If your project is open to the public, anyone in the world can see it. Use this option when the information on this page is ready for a wide audience.
- To show this page only to users with the project administrator role, select **Hidden**. Use this option if you are drafting content that you aren't ready to share yet, or want to share only with other project managers.

6. Click **Save**.

7. Click **Configure: Off**.

Now you are ready to build functionality into your project page with components such as text, news or tracker queries.

### Create a tracker


Create a tracker whenever you need to report and track bugs, feature requests, support requests, or any other type of issue where ownership, status, and activity must be managed.

Individual tracker entries are referred to as tracker artifacts. A tracker is a set of tracker items with a common purpose, such as bug reports, feature requests, or tasks.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Tracker Settings**.

3. Click **Create**.

4. On the **Create Tracker** page, provide a name and description for the tracker.

 **Tip:** Descriptions help users learn how best to provide the information you want from them. To maximize your chances of getting useful data, make your description as informative as you can.


5. Select an icon that suggests the type of work the tracker is handling.

This icon will appear with any artifact in this tracker, wherever it is viewed on the site. For example, if someone brings an artifact from this tracker into a planning folder, users of the planning folder can glance at the artifact's icon to see where it comes from.

6. Select the relevant unit from the **DISPLAY EFFORT IN** field. The units displayed here are configured based on the size of the artifacts in the tracker. Eg. Select the unit as **HOURS** for a tracker of small defects, **DAYS** for a tracker of tasks, and **WEEKS** for a tracker of epics.

 **Note:** Configure the units at the project level and not at the planning folder level.

7. Select **INCLUDE FOREIGN CHILDREN** to include points and efforts from children artifacts across the projects in TeamForge.

 **Note:** In a parent artifact, enabling **CALCULATE POINTS** field sums and rolls up the points from all its children artifacts within the project. In this total, if you want to include children artifacts from other projects across TeamForge, have the **INCLUDE FOREIGN CHILDREN** option enabled.

8. Click **Create**.

The new tracker appears at the bottom of your list of trackers.

9. If necessary, drag the tracker to a place in your tracker list that makes sense.

The order you set here controls the order of every tracker list in your project.

10. You'll probably need some custom fields to capture information that's specific to your project.

See [Create your own tracker fields](#) on page 54.

11. To speed up the team's work, you may want to set up some rules for automatically reassigning artifacts when their contents change.

See [Create a tracker workflow](#) on page 55.

### Create a planning folder

A planning folder is a way to organize work into feasible chunks and monitor its progress. As a project admin or as a user with the appropriate permissions, create and populate all the planning folders you need to capture the work you are planning.


When you've thought through your plan, express it in one or more planning folders.

 **Tip:** It often makes sense to set up planning folders *after* you have outlined and analyzed the features you plan to deliver. See [Define the scope of your project](#) on page 171.

For your agile projects, you have the option to create planning folders specific to iterations and releases, or you can create generic planning folders which you may customize later. When you've thought about the general categories the work falls into, you are ready to create planning folders that reflect those ideas.

A planning folder can represent:

- A set of tasks, such as Iteration 3, or "Initial infrastructure development."
- A period of time, such as "April," or "Q2-2010."
- A phase of development, such as "Testing" or "Deployment."
- A component of the product, such as "Chapter 12" or "Rear stabilizer."





1. Click **TRACKERS** from the **Project Home** menu.
2. Under **PLANNING FOLDERS**, click the folder in which you want to create a new planning folder. Click **NEW**. An appropriate planning folder creation page is displayed.
3. Select the type of planning folder you want to create: **Release**, **Iteration** or **Folder**.

While "Release" and "Iteration" are specific planning folder types, "Folder" is your standard, generic planning folder. Though having multiple planning folder types help you logically organize your planning folders, there's no hierarchical restrictions on these different planning folder types. For example, as a user, you have the liberty to create a "Release" planning folder as a child of an "Iteration" planning folder, which may not be the ideal use case in an Agile project.

4. Enter a brief and descriptive name for your planning folder.  
For example:

- In an agile project, depending on your requirement, you can create two or more iteration folders called "Iteration 1", "Iteration 2" and so on within the Release planning folder.
- In a phased, waterfall-style project, you might name your first planning folder "Design," the next "Build," and so on.

 **Tip:** Don't worry if you don't have anything of interest to put into your planning folders yet. The parallel process of filling out the feature tree will provide plenty of material for this.


 **Note:** This field is mandatory.

5. Use the **DESCRIPTION** to briefly signal the kind of work that will be contained in this planning folder.

Include enough information to help people get up to speed quickly when they join your project. But save most of the detail for the individual tracker artifacts, where project members will spend most of their time.


6. For **STATUS**, (available since TeamForge 5.4), select a value that communicates where the planning folder is in its life cycle.

For a new planning folder, you'll probably want to select Not started.


 **Tip:** You can create more values to choose from in your Project Settings.

7. If you have set up a file release in the File Releases tool to deliver the work you are tracking here, you can identify it in the **FILE RELEASE** field (available since TeamForge 5.4).


Any artifact you add to this planning folder will also appear in the Planned Tracker Artifacts tab when you look at that file release.

 **Note:** If you are creating an Iteration folder, this field displays the default file release value set at its parent folder level ("Release" folder) which you can modify.

8. For **CAPACITY**, (available since TeamForge 5.4), provide a number that expresses how much work your team thinks it can do in the period represented by this planning folder. This field is available only for "Iteration" and "Folder" types.

 **Tip:** When you create a new planning folder, it can be a good idea to set its capacity to zero until you get a feel for the artifacts and resources you are working with.

9. For **POINTS CAPACITY**, (available since TeamForge 7.1), provide a number that represents the amount of work (expressed as number of story points) that you think can be handled in that specific planning folder. This field is available only for "Release" and "Folder" types.

 **Note:** To assist project managers gauge the planning folder's points capacity, when a release or standard planning folder is selected in one of the *planning board* swim lanes, this information is displayed as a ratio of planned number of story points versus the actual number of story points. This information does not apply to Iteration folders and therefore does not appear on the planning board swimlanes when you select an Iteration planning folder. For more information, see *Use the Planning Board* on page 182.

10. Select or enter a start and end date for the work that this planning folder will contain.

You can change these dates at any time. However, it's important to specify dates even if they're not firm, because useful visual aids and reports depend on them.

 **Note:**

- The start and end dates of the child planning folder should be within that of its parent folder.
- These dates are mandatory for an "Iteration" planning folder because the burndown chart for Iteration planning folders is enabled only when you specify the start and end dates of an Iteration planning folder. Also, these dates are not mandatory for a Release planning folder because the burndown chart of a Release planning folder takes into account the start and end dates of its child planning folders (Iterations) and not its own.


11. Select the unit that is suitable for the planning folder from the **DISPLAY EFFORT IN** field. Eg. if the planning folder represents a sprint, then select the unit as **Hours** and if it represents a collection of sprints, then select the unit as **Points**.

 **Note:** Units are configured at the project level and not at the planning folder level.


12. Select either **Effort** or **Points** from the **BURNDOWN VIEW BY** drop-down list to set the View Mode in the burndown charts (in the planning folder **List Artifacts** page). This field is available only for "Folder" type.

Depending on the planning folder type, the burndown chart is displayed.


- Burndown chart for the standard planning folder type which includes the View Mode:



- Burndown chart for "Iteration" planning folder type:



- Iteration-wise burndown chart for a "Release" planning folder type:



13. Click **Create**. The planning folder is created.

**Set up hardware for your team to use**

When you set up Lab Management, your team members can use TeamForge to access their own virtual machines for developing and testing.

1. Go to your Lab Management project.

2. Select **Project Admin** from the **Project Home** menu.
3. On the **Project Admin** menu, click **Tools**.
4. Click **Add Tool**.
5. Select **Other** from **Select Tool Type** drop-down.
6. Enter the display name as **Lab Mgmt**.
7. Select the **Show in Toolbar** checkbox.
8. Enter the server location or URL of your Lab Management server.
9. If you are a TeamForge administrator, select whether you want to use single sign-on for the linked application.
  - If you use single sign-on, TeamForge manages authentication for Lab Management, and users don't have to log into Lab Management after they have logged into TeamForge.
  - If you do not use single sign-on, users must log into Lab Management using its native authentication system.
10. Click **Browse** and select this icon for the linked application:



This icon will appear with the name **Lab Mgmt** in the project navigation bar.

11. Click **Save**.

A **Lab Mgmt** button is added to your **Project Home** menu.. Clicking it launches the application in the main TeamForge project window.

## Contribute to a CollabNet TeamForge project

CollabNet TeamForge is a set of integrated tools designed to make your collaborative experience more efficient, productive and fun.

### Plan and track anything

Use the issue tracking system to track and manage everything from requirements, user stories, and feature requests to defects, build issues, and customer support cases.

Project administrators can model any desired lifecycle process by customizing tracker fields to correspond to the different states required, defining a work flow that governs the transitions between those states, and specifying which users can make those transitions.

### Estimate effort for agile development

Once the product features have been sketched out, it falls to the project team to figure out what it will take to deliver them.

Effort estimates are almost always more accurate when they address small, discrete chunks of work. But you need to forecast effort for a large, complex undertaking. TeamForge helps bridge this gap by dynamically calculating effort estimates at multiple levels.

There are two ways to determine the effort involved in an artifact:

- Break the work down into discrete chunks, then use your experience and judgment to come up with a real-world estimate for this artifact.
- Have the effort estimates for all the artifact's children (if any) automatically added and shown in this artifact.

You will be using a combination of these techniques as you work through the estimating process.

Analyzing the work is an iterative, interactive process. As you go through it, you'll find your effort estimates changing shape in response to feedback from scoping the product and planning out the work.

1. With your team, review the artifacts the product owner has created.

How much effort will each one require? Do any of them require special skills or training? Record your observations about each artifact in its **Comment** field.



**Note:** This is a rough cut. Don't try to come up with a final effort estimate yet.



- For each user story, identify the specific tasks that will be required. Create an artifact to contain each of these tasks.

Declare each of these task artifacts a "child" of the initial artifact.

Now you can attempt to estimate the work involved in each task artifact. This will be the raw material for your higher-level effort estimates and progress tracking.

- Back in the "parent" artifact, select **Sum effort from children**.

Now this artifact's effort field shows the total of the effort units assigned to each of its "children."



**Note:** If an artifact has children that belong to another project, the effort value for those child artifacts is not counted.

- As you proceed, you will find the overall picture taking on definition.

Relationships among artifacts will emerge, and many of them will seem to fall into one logical group or another. Let these groupings determine the names and purposes of your planning folders.

### Find tracker artifacts

The tracker contains a comprehensive search system that enables you to find a specific artifact or set of artifacts quickly.

#### *View the right tracker columns*

When you're looking at the artifacts list in a tracker, a planning folder or a team, you can select the columns you want to see, either for this session or permanently.

You set your column preferences for each tracker, planning folder or team independently. If your project administrator has set default columns for the entire project, your individual column choices override those settings.

- Click **TRACKERS** from the **Project Home** menu.
- Select a tracker, planning folder or team and click **COLUMNS > Configure**.
  - If you've already saved a column configuration, click it and skip the rest of these steps.
  - To go back to the default column configuration, click **System (default)** and skip the rest of these steps.
  - To set up a new configuration, click **Configure**.
- Choose your columns.
  - Move the columns you want from **Available Columns** to **Selected Columns**.  
**Artifact ID : Title, Priority and Status** are required columns.
    - Note:** Selecting more columns can increase the time required to load the listing page.
  - Remove any columns you don't need from **Selected Columns**.
  - Use the move up and move down arrows to change the display order of the columns.
- Apply your choices to your view of the tracker.
  - To use this arrangement this time only, click **Apply**. The next time you log in, you'll start with the default view again.
  - To save your column layout for repeated use, click **Apply and Save**, then give your arrangement a name. The next time you log in, you'll see the column arrangement you just selected. (If you've sorted the records in your view, that sort order is saved too.)



**Tip:** If you are editing a column configuration that already exists, you can rename it by saving it under a new name.

- To make the same set of columns appear every time you come to this tracker, planning folder or team, click **COLUMNS > Save** and from **Save Column Configuration** page, select **Make this my default view**.

#### *View the right tracker artifacts*


When a large number of tracker artifacts makes it difficult to find the one you want, filter and sort the list to narrow down the possibilities.

- Click **TRACKERS** from the **Project Home** menu.
- On the **Tracker Summary** page, click the title of the tracker in which you want to look at artifacts.
- Specify the filter criteria in one or more filter fields (at the top of each column) and click **FILTER**.

- You can find a filter field at the top of each column in most of the tables in the TeamForge application.
- The filter field could be a text box or a drop-down list with multi-select check boxes.

The screenshot shows a table titled "Stories" with "Page 1 of 40 (600 Items)". The table has columns for PRIORITY, ARTIFACT ID : TITLE, ASSIGNED TO, TEAM, and SUBMITTED BY. Each column has a filter field at the top. The filter fields are: PRIORITY (Any), ARTIFACT ID : TITLE (Filter by), ASSIGNED TO (Filter by), TEAM (Any), and SUBMITTED BY (Filter by). The table contains 10 rows of data, each with a priority level (e.g., MEDIUM, LOW, HIGH), an artifact ID, a title, an assigned person, a team, and a submitted by person. At the bottom of the table, there are buttons for Monitor, Import, Export, Cut, Edit, Edit Inline, and Mass Update.

- You can type your filter criteria in the text boxes. The search text is case-insensitive.
- You can also select the filter values from one or more drop-down lists. By default, you can only select up to 10 filter values in a drop-down list. However, you can set a value that suits your requirement for the `FILTER_DROPDOWN_MAX_SELECTION` token in the `site-options.conf` file to increase or decrease the count.
- **Filter-as-you-type:** You can find the **Enter keywords** text box in all filter drop-down lists. As you type your filter keyword, instant search results are shown in the drop-down list. For example, in the following illustration, typing "R" instantly shows all statuses having the alphabet "R". The search text is case-insensitive.



- Some search filters may not appear if your site administrator has not enabled them.

#### 4. After filtering, if you want to save a filter for future use:

- a) Click **FILTER** and select **Save** from the drop-down list.

The **Save Filter As** window appears.

- b) Type a name for the filter in the **FILTER NAME** text box.

- c) Click **Save**.

The filter is saved. You can view and select the saved filters at a later point in time by clicking the **Filter** drop-down list.

- 👉 **Note:** You can save filters only in specific contexts. This feature may not be available in all the tables where you can filter table list items.

#### 5. To delete saved filters:

- a) Click **FILTER** and select **Delete** from the drop-down list.  
The **Select Filters To Be Deleted** window appears.
- b) Select one or more filters to delete.

 **Tip:** Press and hold the **Ctrl** key to select more than one filter.

- c) Click **Delete**.  
A message such as `2 tracker filter(s) have been deleted successfully.` is displayed if the process was successful.
6. After filtering, if you want to clear the filters, click **FILTER** and select **Clear** from the drop-down list.
  7. Use the up-down arrow at the top of any column to sort your list by that column.
    - Your primary sort column is identified by a superscript **1** next to the up-down arrow, and your secondary and third-level sort columns, if any, are likewise marked.
    - Click the up-down arrow again to reverse the sort order.

#### *View the right planning folder artifacts*

When a large number of artifacts in a planning folder makes it difficult to find the one you want, you can filter tracker artifacts by their overall status or you can filter and sort the tracker artifact list column-wise in the Planning Folders list view.

- In the Planning Folder list view, you can filter only by **Priority, Artifact ID, Title, Assigned To** and **Team** columns.
  - Though the artifacts are listed in a tree view (parent artifact with its child artifacts), the filter is applicable only for the parent artifacts and not their children.
  - The filter is available only in the Sort mode and not in the Rank mode.
  - The filter that you set is retained even after you navigate to other pages and return to this page.
1. Click **TRACKERS** from the **Project Home** menu.
  2. Click **Planning Folders**.
  3. On the Summary page, click the planning folder in which you want to look at artifacts.
  4. Specify the filter criteria in one or more filter fields (at the top of the filterable columns) and click **FILTER**.
    - The filter field could be a text box or a drop-down list with multi-select check boxes.
    - You can type your filter criteria in the text boxes. The search text is case-insensitive.
    - You can also select the filter values from one or more drop-down lists. By default, you can only select up to 10 filter values in a drop-down list. However, you can set a value that suits your requirement for the [FILTER\\_DROPDOWN\\_MAX\\_SELECTION](#) token in the `site-options.conf` file to increase or decrease the count.
    - **Filter-as-you-type:** You can find the **Enter keywords** text box in all filter drop-down lists. As you type your filter keyword, instant search results are shown in the drop-down list. The search text is case-insensitive.
  5. After filtering, if you want to clear the filters, click **FILTER** and select **Clear** from the drop-down list.
  6. Use the up-down arrow at the top of any column to sort your list by that column.
    - Your primary sort column is identified by a superscript **1** next to the up-down arrow, and your secondary and third-level sort columns, if any, are likewise marked.
    - Click the up-down arrow again to reverse the sort order.

#### **Filter by artifact status**

In addition to the column-wise filter, you can also filter and view artifacts based on their status alone using the status drop-down list.

Release Template		Open only		FILTER		SORT		RANK		COLUMNS	
PRIORITY	ARTIFACT ID	TITLE	ASSIGNED TO	TEAM	STATUS	GROUP	CATEGORY	POINTS			
HIGHEST	artf255409	Prepare product pipeline for new version	None	Devops/Releng	Open			0			
HIGHEST	artf255407	Prepare system pipeline for new version	None	Devops/Releng	Open			0			
HIGH	artf263978	Validate SSL cypher suites	None	Devops/Releng	Open			0			
MEDIUM	artf256259	Update thirdparty components	None	Devops/Releng	Open			0			
MEDIUM	artf270970	Create data dump with release-specific data	None	Devops/Releng	Open			0			
MEDIUM	artf264726	Update Agile Baseline project template	None		Open			0			

7. Select **All**, **Open only** or **Closed** from the drop-down list and all the relevant artifacts (including child artifacts) are displayed. For example, select **Open only** and all the artifacts with associated status values such as 'Under Construction', 'Open' and 'In Progress' are displayed.



**Note:** The statuses **All**, **Open only** and **Closed** are associated with user-defined status values while configuring a tracker field on the **Edit Tracker Field** page (**Project Admin > Tracker Settings**). For more information, see [Configure tracker "Select" field values](#) on page 54.

### View the right team artifacts

You can filter artifacts by their overall status or you can filter and sort the tracker artifact list column-wise in the Teams list view.

- In the Teams list view, you can filter only by **Priority**, **Artifact ID**, **Title**, **Assigned To** and **Planned For** columns.
- The filter that you set is retained even after you navigate to other pages and return to this page.



**Note:** The Sort and Rank modes are not available on the Teams list view.

### Filter by columns

1. Click **TRACKERS** from the **Project Home** menu.
2. Click **Teams**.
3. On the Summary page, click the team in which you want to look at artifacts.
4. Specify the filter criteria in one or more filter fields (at the top of the filterable columns) and click **FILTER**.
  - The filter field could be a text box or a drop-down list with multi-select check boxes.
  - You can type your filter criteria in the text boxes. The search text is case-insensitive.
  - You can also select the filter values from one or more drop-down lists. By default, you can only select up to 10 filter values in a drop-down list. However, you can set a value that suits your requirement for the [FILTER\\_DROPDOWN\\_MAX\\_SELECTION](#) token in the `site-options.conf` file to increase or decrease the count.
  - **Filter-as-you-type:** You can find the **Enter keywords** text box in all filter drop-down lists. As you type your filter keyword, instant search results are shown in the drop-down list. The search text is case-insensitive.
5. After filtering, if you want to clear the filters, click **FILTER** and select **Clear** from the drop-down list.
6. Use the up-down arrow at the top of any column to sort your list by that column.
  - Your primary sort column is identified by a superscript **1** next to the up-down arrow, and your secondary and third-level sort columns, if any, are likewise marked.
  - Click the up-down arrow again to reverse the sort order.


### Filter by artifact status

In addition to the column-wise filter, you can also filter and view artifacts based on their status alone using the status drop-down list.

QE - Core Automation		Open only	FILTER		COLUMNS			
PRIORITY	ARTIFACT ID	TITLE	ASSIGNED TO	STATUS	PLANNED FOR	POINTS	GROUP	CATEGORY
HIGHEST	artf286416	Run AppScan against 17.1.460 (IT 3) build	Thiru Kandasamy	In Progress	Backlog > Release 17.1 > Hardening			
HIGH	artf286828	Run AppScan against 17.1.493 build	Thiru Kandasamy	In Progress	Backlog > Release 17.1 > Hardening			
MEDIUM	artf287474	Automation Regression test suite needs to be run on Migrated Database	None	Open		0		
HIGH	artf288405	Run AppScan against IT 2 (17.4.254) build	Murali Yalamandala	Open	Backlog > Release 17.4 > Iteration 2			

Monitor Create Artifact Defects

7. Select **All**, **Open only** or **Closed** from the drop-down list and all the relevant artifacts (including child artifacts) are displayed. For example, select **Open only** and all the artifacts with associated status values such as 'Under Construction', 'Open' and 'In Progress' are displayed.

 **Note:** The statuses **All**, **Open only** and **Closed** are associated with user-defined status values while configuring a tracker field on the **Edit Tracker Field** page (**Project Admin > Tracker Settings**). For more information, see [Configure tracker "Select" field values](#) on page 54.

### Search a tracker

If the filter returns too many results or not enough, try the search facility.

#### Search for tracker artifacts

Find tracker artifacts by supplying a keyword, an artifact ID, a date range or some other value.


You can also search comments, attachments, and user-defined fields.

1. Click **TRACKERS** from the **Project Home** menu.
2. Set the scope of your search.
  - To search all the trackers in your project, click **Search Trackers**.
  - To search a specific tracker, click the tracker you want to search, then click **Search Tracker** in that tracker's list view.
3. In the **Tracker Search Criteria** section, enter the keywords to search for.

 **Note:** Wildcards are allowed.

4. Select the elements of the artifact to search.  
For example, to search for a keyword in the title and description fields, select the **Title** and **Description** check boxes. If you want to search all the text fields (including **Title**, **Description** and user-defined text fields), select the **All Text Fields** check box.
5. To search for a keyword in the comments field, select **Include Comments**.
6. If you know the artifact's ID, enter that.
7. Click the calendar icon to select dates, if appropriate. You can also specify relative dates such as "Within the last 7 days".
8. Use as many of the remaining tracker search criteria as seems useful.
9. Click **Search**.

All tracker artifacts matching your search criteria are displayed.

 **Note:** If your search included multiple trackers, the icon next to each artifact in your search results can help you identify which tracker that artifact belongs to.

#### Find your own artifacts

To narrow your scope, try searching only for artifacts assigned to or submitted by you.

1. Click **TRACKERS** from the **Project Home** menu.
2. On the list of project trackers, click the title of a tracker.

3. In the tracker artifact list view, click **Search Tracker**.
4. In the **Tracker Search Criteria** section, select the **User running search** option in the **Assigned To** or **Submitted By** fields.
5. Click **Search**.

All tracker artifacts matching your search criteria are displayed.

#### Save a search for tracker artifacts

To reuse a search that you devised for tracker artifacts, save the search criteria.

To be able to save your search, you must first select a tracker, and create and run your search.

1. In the **Tracker Search Results** section, click **Save Search from Results**.
2. In the **Save Search As** page, enter a descriptive name for the search, and click **Save**.

The search appears under **My Saved Tracker Searches**. The name of the search, the tracker for which it was run, the criteria used, and a system-assigned ID are displayed.

#### Share a saved tracker search

To enable other users to use a search that you have devised, save it as a shared search.

To share your search with other users, you must be a tracker administrator.

1. In the **Tracker Search Results** section, click **Save Search from Results**.
2. In the **Save Search As** page, enter a descriptive name for the search.
3. Select the **Shared Search** option.

The search appears under **Shared Tracker Searches**. The name of the search, the tracker for which it was run, the criteria used, the user who created the search, and a system-assigned ID are displayed.

#### Run a saved search

To find artifacts in the current tracker or a different one in the project, run a saved search in the appropriate context.

1. Click **TRACKERS** from the **Project Home** menu.
2. On the list of project trackers, click the title of a tracker.
3. In the tracker artifact list view, click **Search Tracker**.
4. Select the **Show My Searches Across Trackers** and **Show Shared Searches Across Trackers** options.
5. Expand the **My Saved Tracker Searches** and **Shared Tracker Searches** sections.  
The saved searches for all trackers in the project are displayed.
6. To run a search for the current tracker, click the **Search** or **Search Here** link below the search name in the required search.  
All artifacts in the current tracker, that match the search criteria, are displayed.
7. To run a search for a different tracker, click the **Search There** link below the tracker name in the required search.  
All artifacts that match the search criteria in the selected tracker, are displayed.

#### Refine a saved search

Refine your saved searches and keep them updated.

To be able to refine (edit) a previously saved search, you must first run the search and click **Refine Search**.


1. Click **TRACKERS** from the **Project Home** menu.
2. Click **Search Trackers**.
3. Click **My Saved Tracker Searches**.
4. Click the **Search** link under the saved search (in **Saved Search Name** column) you want to refine.  
The saved search is run and results appear.
5. Click **Refine Search**.  
The Tracker Search Criteria for the saved search shows up.
6. Modify the search criteria and click **Save Search**.  
The **Save Search As** dialog box appears.
7. Type a name and click **Save**.

You may save the search with the same name, particularly if you have already shared the search with others. Otherwise, you can save the search with a new name too.

The search is saved with your new search criteria.

#### Remove a saved search

To remove a saved search for tracker artifacts, delete it from the list of personal or shared searches.

 **Note:** Only tracker administrators can remove a shared search.

1. Click **TRACKERS** from the **Project Home** menu.
2. On the list of project trackers, click the title of the tracker in which you want to delete a saved search.
3. In the tracker artifact list view, click **Search Tracker**.
4. Expand **My Saved Tracker Searches** or **Shared Tracker Searches** so that the available searches are displayed.
5. Select a search and click **Delete**.
6. Accept the confirmation message.

The search is removed from the table where you deleted it.

#### Create a tracker artifact


Create a tracker artifact whenever you need to report and track a bug, feature request, support request, or other type of issue.

Individual tracker entries are referred to as tracker artifacts, or just artifacts.


1. In any tracker, planning folder or teams view, click **Create Artifact** and select the tracker in which you want to create your artifact.

By default, your new artifact is created in the tracker, planning folder or team you are currently looking at.


2. Answer the questions posed by the required fields.

 **Note:** Different trackers will have different combinations of fields to fill out, depending on what the tracker administrator has set up.


- a) Provide a **Title** and **Description** that summarize the issue or work item in a few words.


 **Tip:** Descriptions help users learn how best to provide the information you want from them. To maximize your chances of getting useful data, make your description as informative as you can.

**Artifacts support @mentions:** Artifact description and comments now support @mentions and users called out via @mentions are added to the monitoring list. Include usernames with "@" as prefix (for example, @mhippard) to add users to the monitoring list.

 **Note:** Users called out via @mentions must have "Artifact View" permission to be added to the monitoring list.

- b) Now project members with CREATE/EDIT permissions can also create new tags or add existing tags, if required from **Create Artifact** page. Tags creation from **Create Artifact** page enables you to create tags on the go and overrides the limitation of creating tags only from the **Tags** page. However, you cannot rename or

delete a tag from **Create Artifact** page. Click the  button next to **Tags** field to get the list of tags mapped to your project. You can add up to a maximum of 10 tags to any artifact and a message is displayed if you try to add more than 10 tags. If the entered tag name is not available already, a context menu **Create a new tag** shows up for you to create a new tag with the desired tag name.

TAGS:  [Add a tag](#)

GROUP:

STATUS:\*

CATEGORY:

CUSTOMER:

PRIORITY:

TEAM:

ASSIGNED TO:

NRB

installer

must

508 compliance testing

BLOCKER


PSR

QA-17.1 results


CRI

license generator - FCN

REG

 **Note:** Wherever the tag widget is not applicable, the associated tags are displayed as read only tags. For example in **View Artifact** page.

- c) For **Priority**, select a value that expresses the importance or urgency of the work you are describing.
- d) Assign the artifact to a team by selecting a name from the **Team** list.
- e) If you want to assign it to a specific project member, choose a name from the **Assigned To** list. This lists displays the names of all the project members, irrespective of the team you may have selected in the previous step.

 **Note:** If your project administrator has configured the tracker to automatically assign artifacts to project members, you can skip this step.

Reassigning artifacts can now be done in no time. Use the links under the "Assigned To" field to quickly unassign the artifact to "None", reassign the artifact to yourself or to the previous assignee. You can also click the "Re-assign" icon to search and reassign the artifact to any other user.

ASSIGNED TO: [Kevin Maples](#) 

[Me](#) | [Unassign](#)

ASSIGNED TO: [Karan Garewal](#) 

[Me](#) | [Unassign](#) | Previous Assignee : [Cole Miller](#)

- f) Select the planning folder that the work belongs to from the **Planning Folder** list.
3. Record any other information that may be appropriate.  
For example, if your project is using a Scrum-based methodology, your project manager may have provided a **Points** field to track estimates of relative effort.
  4. Add a file attachment, if appropriate.
  5. Save your changes.


### Create a tracker artifact by email

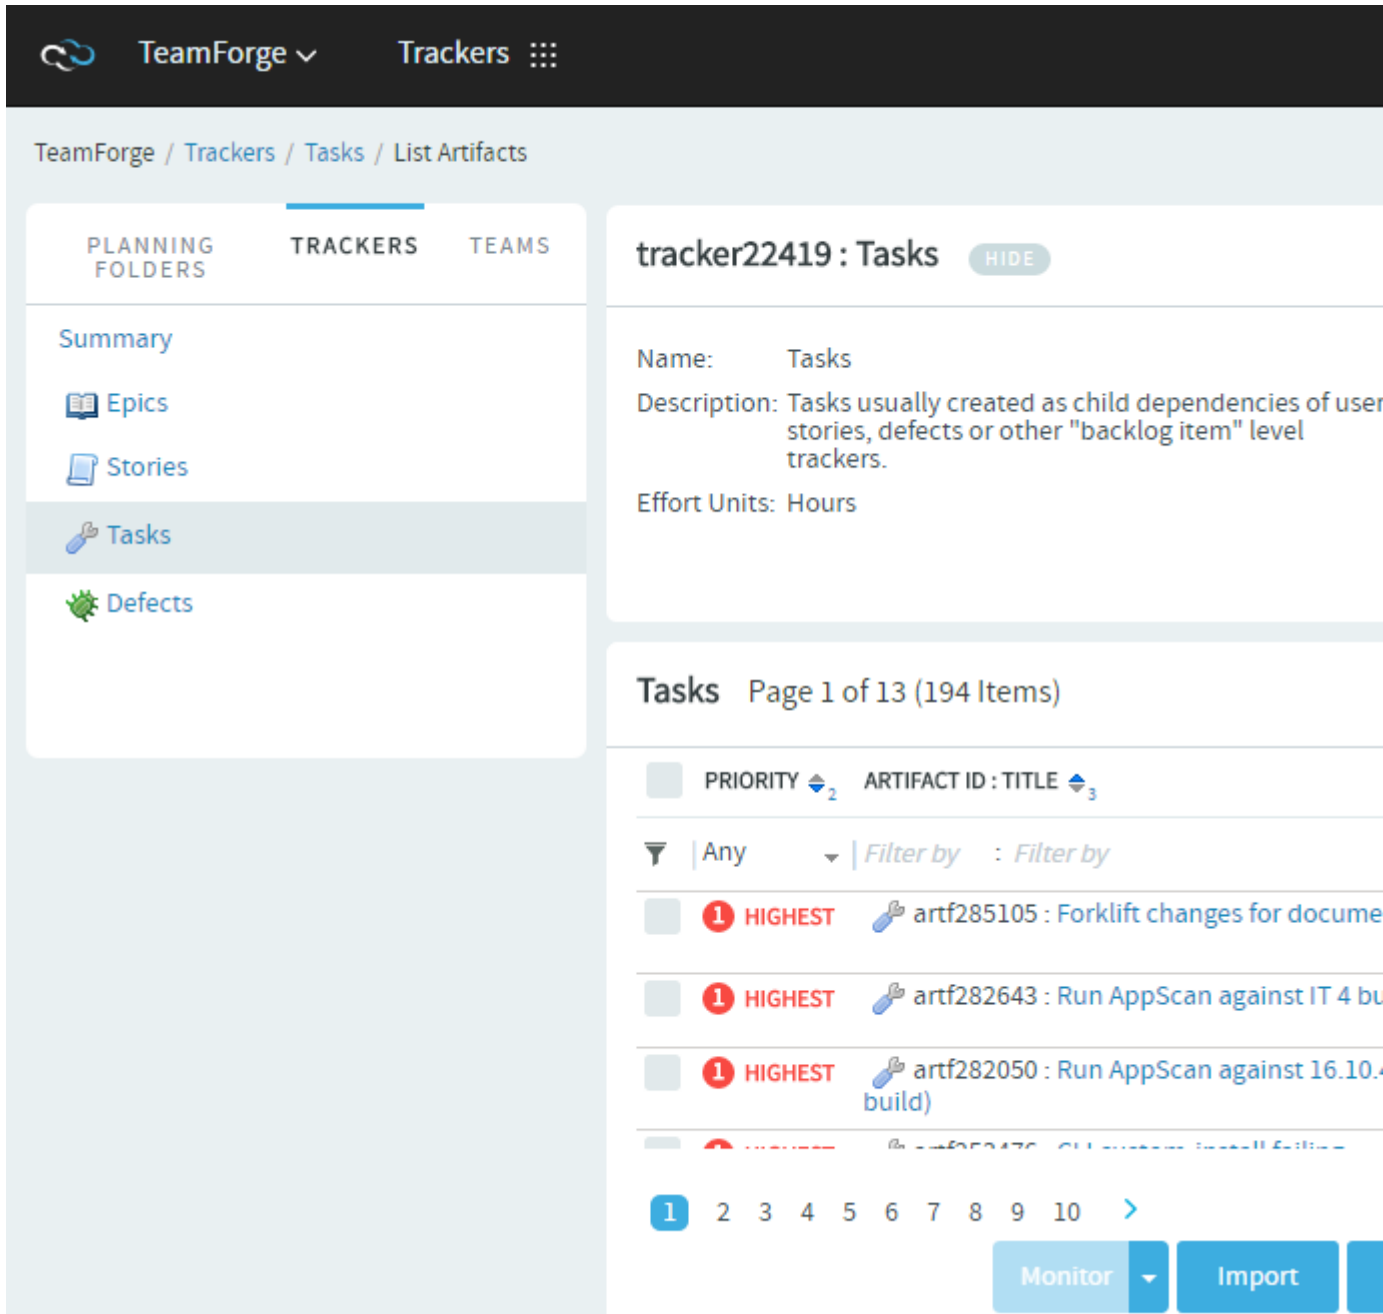
To create a tracker artifact without being logged into CollabNet TeamForge , send an email to the tracker.

You do not have to be logged into TeamForge to submit a tracker artifact using email, but you must have the tracker submit permission for the tracker to which you are submitting.

Send an email message to <tracker id>@< TeamForge server>.



 **Tip:** You can find the tracker ID on the **List Artifacts** page.



The screenshot shows the TeamForge interface. At the top, there's a navigation bar with 'TeamForge' and 'Trackers'. Below that, a breadcrumb trail reads 'TeamForge / Trackers / Tasks / List Artifacts'. The main content area is divided into two columns. The left column has a sidebar with 'PLANNING FOLDERS' (TRACKERS, TEAMS) and a list of items: Summary, Epics, Stories, Tasks (highlighted), and Defects. The right column shows the details for 'tracker22419 : Tasks', including a 'HIDE' button, Name: Tasks, Description: Tasks usually created as child dependencies of user stories, defects or other "backlog item" level trackers., and Effort Units: Hours. Below this is a table of tasks, titled 'Tasks Page 1 of 13 (194 Items)'. The table has columns for 'PRIORITY' and 'ARTIFACT ID : TITLE'. The first three rows show tasks with 'HIGHEST' priority and titles like 'artf285105 : Forklift changes for docume...', 'artf282643 : Run AppScan against IT 4 bu...', and 'artf282050 : Run AppScan against 16.10. build)'. At the bottom right, there are 'Monitor' and 'Import' buttons.

TeamForge maps your email to the tracker record like this:

**Email field**

To

Subject

Body

Attachments

**Tracker field**


Tracker email address

Artifact title

Artifact description

Attachments

**Artifacts support @mentions:** Artifact description and comments now support @mentions and users called out via @mentions are added to the monitoring list. Include usernames with "@" as prefix (for example, @mhippard) to add users to the monitoring list.

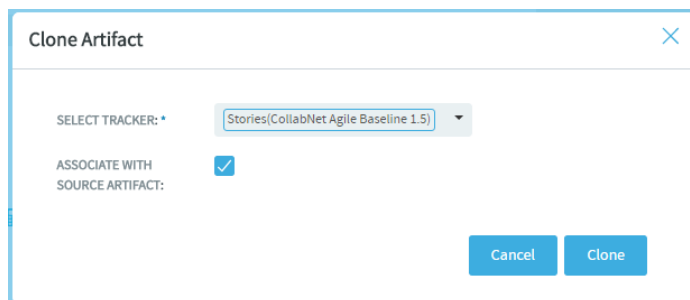
 **Note:** Users called out via @mentions must have "Artifact View" permission to be added to the monitoring list.

### Clone a tracker artifact

Save effort in duplicating a tracker artifact within the project and across projects by cloning artifacts in TeamForge.

You can now clone an artifact from one tracker to another within the project or across projects. For system defined mandatory fields, artifact data are copied from the source tracker to the target tracker as long as the fields are available on both the source and target trackers. Values for any conflicting fields are set to "None" otherwise.

1. Click **My Workspace** and select a project.
2. Select **Project Home > Trackers** and select the tracker that contains the artifact to be cloned.
3. Open the tracker you want to clone.  
The **View Artifact** page appears.
4. Click **Clone**.



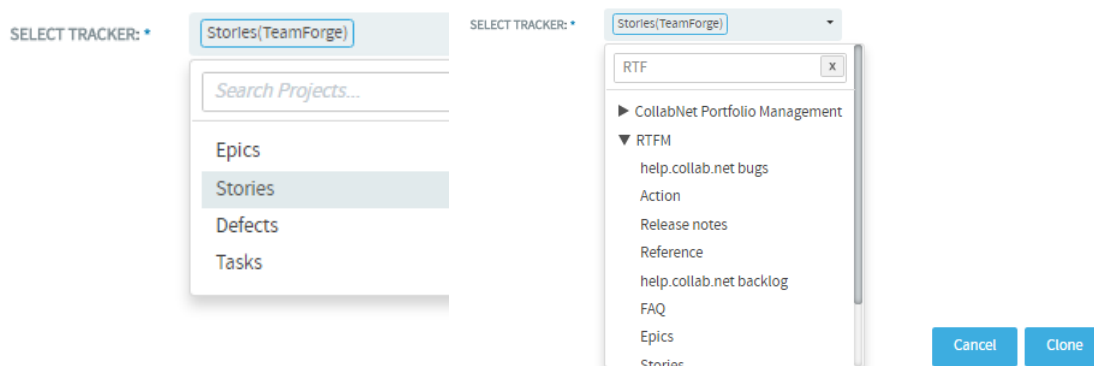
The **Clone Artifact** dialog box appears.

5. Select the **Associate With Source Artifact** check box (selected by default) to create an association with the source artifact.
6. Click **Select Tracker** drop-down list and select the target tracker.

You can either select a tracker from within the project in context or search and select a tracker from a different project. Use the **Search Projects** text box to find the project that has the target tracker and select the tracker.

**Selecting a tracker from within the same project**

**Selecting a cross project tracker where you want to clone the artifact**



7. Click **Clone**.

The **Submit Artifact** page appears. For system defined mandatory fields, artifact data are copied from the source tracker to the target tracker as long as the fields are available on both the source and target trackers. Values for any conflicting fields are set to "None" otherwise.

8. Update the fields, if required, and click **Save** or **Save and View**.

The artifact is cloned and saved in the target tracker.

### Edit a tracker artifact

Updating the information in tracker artifacts is one important way that project members can work together effectively.


When work has been done on a tracker item, or more information is needed, the project member to whom the item is assigned should update the item's status accordingly. Comments from other project members help the artifact's owner decide how to handle the work.

For example, when the work defined in the tracker item is completed, change its status from **Open** to **Fixed**.


Your tracker administrator may have set up work flow rules that constrain your ability to do certain kinds of updates. For example, an administrator may have specified that only users with the "QA Engineer" role can change an artifact's status from **Open** to **Closed**.


A project manager might also change an artifact's priority, return it to the submitter for additional clarification, or assign it to a project member for resolution or action.

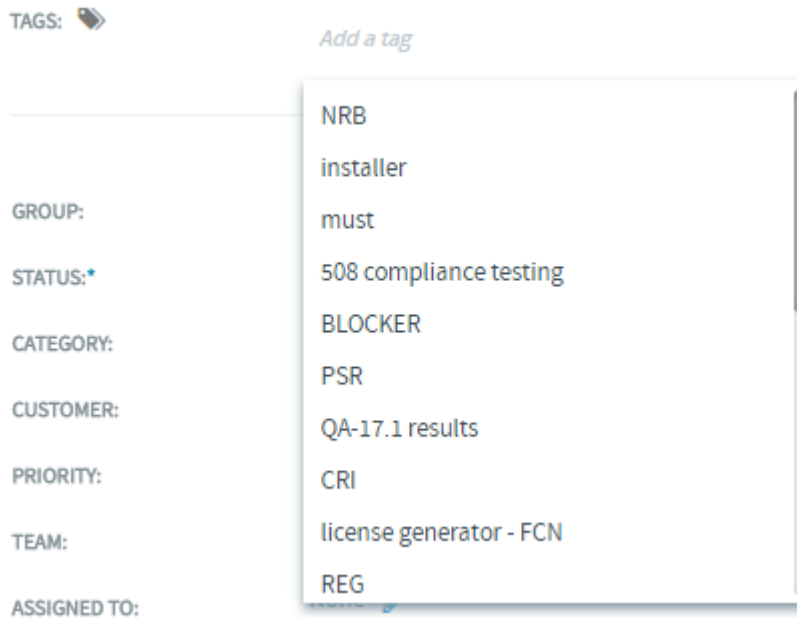
Generally, the project member to whom the tracker artifact is assigned will update the status. Any project member with the appropriate permissions can add comments.

 **Tip:** Each comment in an artifact or task has a unique ID with its own URL. To link directly to a particular comment, copy that comment's URL and paste it into an email, a project page, or another comment. For example, to point to the third comment in artifact 12345, write `artf12345#3` in your comment. (If the artifact or task you are linking to is on a different site, give the complete URL, like this: `http://mysite.com/sf/go/task1234#3`.)

**Artifacts support @mentions:** Artifact description and comments now support @mentions and users called out via @mentions are added to the monitoring list. Include usernames with "@" as prefix (for example, @mphippard) to add users to the monitoring list.

 **Note:** Users called out via @mentions must have "Artifact View" permission to be added to the monitoring list.

1. Click **TRACKERS** from the **Project Home** menu.
2. On the list of project trackers, click the tracker you want to update.
3. On the **List Artifacts** page, select the tracker artifact you want to update.
4. On the **View Artifact** page, click **Edit**.
5. On the **Edit Artifact** page, make your changes.
  - a) Select a tracker type from the **Tracker** drop-down list to change the artifact's tracker type.
  - b) Edit the artifact's title and description.
  - c) Click **Save**.
6. On the **View Artifact** page, change any other fields if required, and click **Update**.
  - a) Now project members with CREATE/EDIT permissions can also create new tags or add existing tags, if required from **View Artifact** page. Tags creation from **View Artifact** page enables you to create tags on the go and overrides the limitation of creating tags only from the **Tags** page. However, you cannot rename or delete a tag from **View Artifact** page. Click the  button next to **Tags** field to get the list of tags mapped to your project. You can add up to a maximum of 10 tags to any artifact and a message is displayed if you try to add more than 10 tags. If the entered tag name is not available already, a context menu **Create a new tag** shows up for you to create a new tag with the desired tag name.



- 👉 **Note:** Wherever the tag widget is not applicable, the associated tags are displayed as read only tags. For example in **View Artifact** page.

All updates to the artifact are recorded in the **Comments** section of the **Status/Comments** tab.

### **Edit a tracker artifact by email**

To comment on a tracker artifact when you are not logged into TeamForge, send an email to the tracker, or reply to an automatic update email about the artifact.

You can also add a comment or an attachment to a tracker artifact that you are monitoring by responding to the monitoring email notification.

You can also use email to add an attachment. But to edit any other fields, you must make the changes in TeamForge.

You do not have to be logged into TeamForge to edit a tracker artifact using email, but you must have the tracker edit permission for the tracker containing the artifact you want to edit.

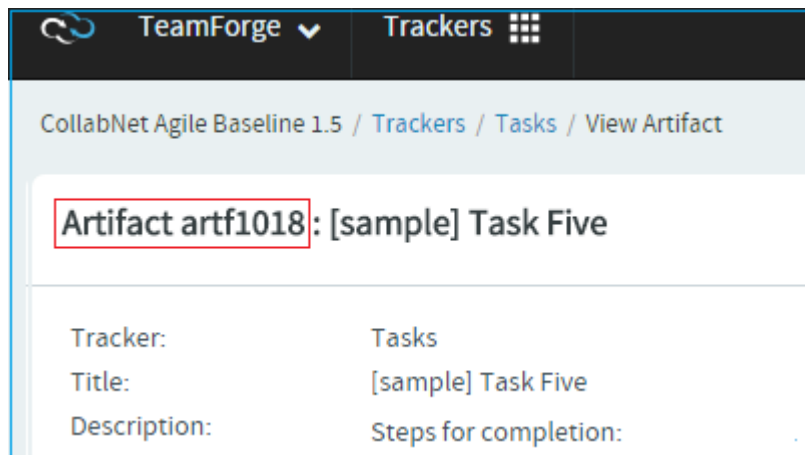
- 👉 **Tip:** Each comment in an artifact or task has a unique ID with its own URL. To link directly to a particular comment, copy that comment's URL and paste it into an email, a project page, or another comment. For example, to point to the third comment in artifact 12345, write `artf12345#3` in your comment. (If the artifact or task you are linking to is on a different site, give the complete URL, like this: `http://mysite.com/sf/go/task1234#3`.)

**Artifacts support @mentions:** Artifact description and comments now support @mentions and users called out via @mentions are added to the monitoring list. Include usernames with "@" as prefix (for example, @mhippard) to add users to the monitoring list.

- 👉 **Note:** Users called out via @mentions must have "Artifact View" permission to be added to the monitoring list.

Send an email message to `<artifact id>@<CollabNet TeamForge server>`.

- 👉 **Tip:** You can find the artifact ID on the **Artifact Details** page.




TeamForge maps your email to the tracker record like this:


Email field	Tracker field
To	Tracker email address
Subject	Artifact title
Body	Artifact description
Attachments	Attachments

### Change multiple artifacts


When you have a large number of artifacts to update (for example, all the artifacts in a tracker, or a filtered list of artifacts in a planning folder), you can edit all the artifacts at once.

 **Note:** When you update two or more artifacts at a time, each user who is monitoring any of the changed artifacts gets a single email describing all the updates.


**Artifacts support @mentions:** Artifact description and comments now support @mentions and users called out via @mentions are added to the monitoring list. Include usernames with "@" as prefix (for example, @mhippard) to add users to the monitoring list.

 **Note:** Users called out via @mentions must have "Artifact View" permission to be added to the monitoring list.


1. Click **TRACKERS** from the **Project Home** menu.
2. Go to the tracker or planning folder that contains the tracker artifacts you want to edit.

 **Tip:** In a tracker, you can use the filter to help you find the desired artifacts.

3. Select the artifacts you want to edit, and click **Mass Update**.
4. For artifacts in a tracker, choose which ones you want to update.
  - **Selected** - Updates only the artifacts that you selected.
  - **Filtered Set** - Updates all artifacts returned by your filter, or all artifacts in the tracker if you did not apply a filter.

 **Tip:** Choose **Filtered Set** when the artifacts span multiple pages and you want to select them all.

5. Make your changes and click **Update**.


 **Important:** Some fields in your tracker may have values that depend on values in other fields, or use validation rules to ensure correct content. If your mass update operation breaks any such dependencies, you must fix the errors before running the mass update.

### Edit multiple artifacts inline

When you have a list of tracker or planning folder artifacts to update, you can edit all the artifacts inline at once.

1. Click **TRACKERS** from the **Project Home** menu.
2. In the **List Trackers, Planning Folders and Teams** page, click the tracker or the planning folder that contains the artifacts you want to edit inline.
3. In the **List Artifacts** page, click **Edit Inline**.  
To help you identify editable columns, all non-editable columns are disabled. You can see a hand symbol when you hover the mouse pointer over the editable columns.
4. Click a field in a column and edit the selected artifact.  
For example:
  - Clicking a field in the Assigned To column lets you edit the person assigned to the specific artifact.
  - Clicking a field in the Planned For column displays the **Planning Folder** dialog box to let you change the planning folder for the specific artifact.

**Artifacts support @mentions:** Artifact description and comments now support @mentions and users called out via @mentions are added to the monitoring list. Include usernames with "@" as prefix (for example, @mphippard) to add users to the monitoring list.


 **Note:** Users called out via @mentions must have "Artifact View" permission to be added to the monitoring list.

5. When you are done, click **Save**.

### Move a tracker artifact


As work proceeds on a tracker artifact, its focus may change. If this happens, it may be appropriate to move the artifact into a different tracker, or to a tracker in a different project.

To move a tracker artifact between projects, you must have either the tracker administration permission or tracker 'view, submit and edit' permissions in both the source and destination projects.

 **Important:** Tracker 'delete' permission is not required to cut and paste tracker artifacts.

You can move one or more tracker artifacts in the same operation.

1. Click **TRACKERS** from the **Project Home** menu.
2. On the list of project trackers, click the title of the tracker containing the tracker artifact that you want to move.
3. Select the tracker artifact (or artifacts) that you want to move, and Then click **Cut**.  
CollabNet TeamForge removes the tracker artifacts and places them on the clipboard.
4. Go to the tracker into which you want to paste the tracker artifacts.

 **Note:** To move the tracker artifacts into another project, first find the destination project using your **Projects** menu.

5. Click **Paste**.
  - If all artifact assignees have the appropriate permissions in the destination tracker, the tracker artifacts are now moved.
  - If not, the **Paste Artifacts** window tells you how many tracker artifacts cannot be reassigned automatically.
6. Choose one of the following methods to reassign any unassigned tracker artifacts.
  - **Reassign Artifacts to:** - Assigns all unassigned tracker artifacts to the project member you select.
  - **Myself** - Reassigns all unassigned tracker artifacts to you.
7. Click **Next**.

The tracker artifacts are now moved to the destination tracker. The details of the move are recorded in the **Change Log** tab of the **Artifact Details** page.

- If the project members to whom the tracker artifacts are assigned have the appropriate permissions in the destination project, the tracker artifacts are automatically reassigned to the same project members. If not, you are offered several options for reassigning them.

- If both trackers use the same user-defined fields and field values, values in these fields are retained.
- If each tracker uses different user-defined fields or field values:
  - Values are set to the default value, if one was specified by the tracker administrator. Default values are mandatory for required fields.
  - Values are set to **None**, if no default value was specified by the tracker administrator.

### Associate tracker artifacts with other items

You can connect tracker artifacts with other CollabNet TeamForge items such as documents or topics. Creating associations between items enables you to define relationships, track dependencies, and enforce work flow rules.

*Associate a tracker artifact with a document, task, integrated application object, or forum*

When a tracker artifact is related to other TeamForge items, such as tasks, documents, integrated application objects, or discussions, you can connect the tracker artifact to the other item by creating an association.

Creating associations between items helps you to define relationships, track dependencies, and enforce work flow rules.

1. Click **TRACKERS** from the **Project Home** menu.
2. On the list of project trackers, find the tracker artifact with which you want to create an association.  
Use the filter if needed.
3. Click the artifact title.
4. On the **View Artifact** page, click the **Associations** tab.  
The list of existing associations appears.
5. Click **Add**.
6. In the **Add Association Wizard** window, select the items with which you want to associate the artifact:
  - **ENTER ITEM ID** - If you know the item's ID, you can enter it directly.



#### Note:

- To associate an object in an integrated application from within TeamForge, use the [`<prefix_objectid>`] format. Successful associations appear hyperlinked.
  - Each integrated application displays its prefix on moving the mouse over the application name in the tool bar.
  - **ADD FROM RECENTLY VIEWED** - Select one of the last ten items you looked at during this session.
  - **ADD FROM RECENTLY EDITED** - Select one of the last ten items you changed.
7. Click **Next**.
  8. You may add a comment in the **ASSOCIATION COMMENT** text box.
  9. Save your work.
    - Click **Finish and Add Another** to add additional associations.
    - Click **Finish** to return to the **Details** page.



**Note:** When an association is added to or removed from TeamForge objects such as tracker artifacts, tasks, documents, discussions, and file releases, a notification mail is sent to users monitoring these objects.

An option is provided at site level and user level to make sure whether the notification mail has to be sent or not. For more information on this, see [Configure your site's settings](#) on page 365.

10. Click the **Associations** tab to view a graphical representation of all the associated items.

Through the Association Viewer, you can choose to view associations in the form of a list or flip over to the Trace view to explore the layers of associations (including parent/child dependencies) laid out in a timeline. You can scroll across the Trace view by dragging the mouse over the association layer or use the 'Previous' and 'Next' arrows to view all the objects as events in a timeline.

While the Associations tab shows the count of the total number of associations, you can only view the most recent 500 associations when you click the Associations tab in case the artifact has more than 500 associations. You can, however, browse through the Association Viewer to view older associations.

- 👉 **Note:** To remind yourself of the details of the association later, look in the **Change Log** tab of the associated **View Artifact** page.

#### *Associate a tracker artifact with a file release*

To track the source and resolution of a bugs or a feature request, associate its tracker artifact with the file release in which it was discovered and fixed.

Tracker artifacts associated with file releases are also displayed separately, providing a simple way to track all issues that were discovered in or fixed in a specific file release.

- 👉 **Tip:** You can also add associations from the tracker artifact's **View Artifact** page.

1. On the **Project Administration** page, enable both the **REPORTED IN RELEASE** and **FIXED IN RELEASE** fields.

- **REPORTED IN RELEASE** - When submitting a new artifact, the user can choose from a drop-down list of all releases in the project to identify the release in which the issue was discovered.
- **FIXED IN RELEASE** - After the issue is fixed, the user can choose from a drop-down list of all releases in the project to identify the release in which the issue was fixed.

- 👉 **Tip:** If either of these fields says *Unknown*, the artifact you are working on may be associated with a file release that you don't have permission to view. You can leave that as it is or change it to a file release that you do have permission to view.

2. When submitting a new artifact, choose the release in which the issue was discovered from the **REPORTED IN RELEASE** drop-down list.

The drop-down list shows all releases in the project, except those that are in **pending** status.

3. After the issue is fixed, update the artifact by choosing the release in which the issue was resolved from the **FIXED IN RELEASE** drop-down list.

The drop-down list shows all releases in the project, except those that are in **pending** status.

The associated tracker artifacts are displayed on the **View Release** page, in the **REPORTED TRACKER ARTIFACTS** and **FIXED TRACKER ARTIFACTS** sections.

- 👉 **Note:** When an association is added to or removed from TeamForge objects such as tracker artifacts, tasks, documents, discussions, and file releases, a notification mail is sent to users monitoring these objects.

An option is provided at site level and user level to make sure whether the notification mail has to be sent or not. For more information on this, see [Configure your site's settings](#) on page 365.

- 👉 **Note:** To remind yourself of the details of the association later, look in the **CHANGE LOG** tab of the associated **View Artifact** page.

#### *Associate a tracker artifact with a code commit*

When checking in files to your SCM repository, you can create links to one or more tracker artifacts or other CollabNet TeamForge items.

Associations track the links between code and the bugs, feature requests, or other tracker artifacts that the code addresses. You can also associate code commits with other TeamForge items, such as tasks or documents.

A project administrator can make associations mandatory for all code commits. When this is made mandatory, the following additional rules pertaining to code commit can also be set:

- Code commits can be performed only for open artifacts.
- To perform a code commit, the committer must be the owner of the specific artifact.

- 👉 **Note:** Once you enforce the above rules, validations are strictly enforced for commits against tracker artifacts only. In case you commit against any other TeamForge object, for example a wiki or a document, mere existence of the object ID ensures successful commit and association and no validations are performed against the status of the object or who it is assigned to.

You can create tracker artifact associations from whatever interface you normally use to check code into your SCM repository. You do not have to log into TeamForge.





Use the same syntax for commits to CVS and Subversion repositories.

When making a code commit, add the associate command in the commit message like this:


[<item id>], such as the TeamForge tracker artifact ID or task ID.

- TeamForge item IDs are always letters followed by four or more numbers, such as `task1029` or `artf10011`.
- To associate a commit with multiple TeamForge items, separate the item IDs with commas.
- All associations are displayed in the **ASSOCIATIONS** tab of the **Commit Details** page.
- The **Comment** section lists the comments made with each commit.

 **Note:** To associate an object in an integrated application, use the [`<prefix_objectid>`] format. Each integrated application displays its prefix on moving the mouse over the application name.

 **Note:** When an association is added to or removed from TeamForge objects such as tracker artifacts, tasks, documents, discussions, and file releases, a notification mail is sent to users monitoring these objects.


An option is provided at site level and user level to make sure whether the notification mail has to be sent or not. For more information on this, see [Configure your site's settings](#) on page 365.

 **Tip:** To remind yourself of the details of the association later, look in the **CHANGE LOG** tab of the associated **View Artifact** page.

### Make a tracker artifact depend on another artifact

To organize your work with tracker artifacts, you can make an artifact a "child" or a "parent" of another artifact.


- A child artifact can have only one parent.
- A parent artifact can have any number of children.
- A parent artifact cannot be closed if a child is open.

 **Note:** "Open" and "Closed," in this context, refers to a type of status, not the status itself. A tracker administrator can specify a group of statuses, such as "Deferred," "Fixed," or "Rejected," as equivalent to "Closed," while "In progress" and "Under consideration" might be specified as "Open" statuses. For example, when you look at the artifact summaries at the top of any tracker list page, you are seeing a summary of the status type, not the status, of the artifacts.

1. On an artifact page, click the **DEPENDENCIES** tab.


2. Choose or create the related artifact.

- If the parent artifact already exists, click **Choose Child** or **Choose Parent**. In the pop-up, type in the artifact ID or choose from the list of your recently edited artifacts.
- If the related artifact does not exist yet, click **Create Child in** or **Create Parent in** and select the tracker that the new related artifact will belong to. Fill in the form the same way as you would for submitting an unrelated artifact.

 **Note:** If **Choose Parent** or **Create New Parent** is not visible, the artifact already has a parent artifact. An artifact can only have one parent artifact.

3. Click **Next**.

4. Write a comment that describes the relationship, if appropriate, and click **Finish**.

 **Note:** When a dependency is added to or removed from a tracker artifact, a notification mail is sent to users monitoring the artifact.

An option is provided at site level and user level to make sure whether the notification mail has to be sent or not. For more information on this, see [Configure your site's settings](#) on page 365.

The parent-child relationship between the artifacts is established.

- To cancel a parent dependency, click **Remove**.
- To cancel a child dependency, select the child artifact and click **Remove**.



**Note:** Task dependencies and tracker artifact dependencies are different things.


- For tracker artifacts, an artifact with dependencies (a "parent" artifact) can't be considered closed unless all of its dependent artifacts ("children") are closed.
- For tasks, a dependency means one task can't start until another task is completed.

### Quickly create dependencies, associations or add attachments

Use the context menu available in planning folder List View to quickly choose a "parent" or "child", remove a "parent", add associations, add attachments or clone artifacts.


In the planning folder List View, you can see a down arrow icon next to the artifact ID (as shown in the following screenshot) when you move your mouse over artifact rows.

#### Down arrow icon




A context-sensitive menu pops up on clicking the down arrow icon.

#### Context menu if an artifact has a parent



#### Context menu if an artifact has no parent



#### To quickly choose a parent or child, add associations and attachments, or remove a parent:

1. Select a planning folder on the **List Trackers, Planning Folders and Teams** page.  
The planning folder **List Artifacts** page appears.
2. Move your mouse over an artifact's ID.  
A down arrow icon appears.
3. Click the down arrow icon to see the context menu. Depending on the context, you can do one of the following tasks.


#### Choose a child

4. Click **Choose Child** from the context menu.  
The **Selecting Children...** window appears.
5. Select the **Enter Artifact ID** option and type an artifact ID. This artifact becomes the child.




**Tip:** You can also select the **Add from Recently Edited** option and choose a child from the list of recently edited artifacts.

6. Click **Next**.
7. You may add a comment in the **DEPENDENCY COMMENT** text box.
8. Click **Finish**.

 **Note:** You may also click **Finish and Add Another** to continue adding more children for the same artifact.

#### Choose a parent

9. Click **Choose Parent** from the context menu.  
The **Selecting Parent...** window appears.
10. Select the **Enter Artifact ID** option and type an artifact ID. This artifact becomes the parent.

 **Tip:** You can also select the **Add from Recently Edited** option and select a parent from the list of recently edited artifacts.

11. Click **Next**.
12. You may add a comment in the **DEPENDENCY COMMENT** text box.
13. Click **Finish**.

#### Remove a parent

14. Click **Remove Parent** from the context menu.  
A confirmation message appears: `Do you want to remove the dependency?`.
15. Click **OK**.

#### Add associations

16. Click **Add Association** from the context menu.  
The **Add Association Wizard** appears.
17. Select the **Enter Item ID** option and type an artifact ID for creating an association.

 **Tip:** You can select the **Add from Recently Edited** option and select an artifact from the list of recently edited artifacts.

18. Click **Next**.
19. You may add a comment in the **Association Comment** text box.
20. Click **Finish**.

 **Note:** You may also click **Finish and Add Another** to continue associating more artifacts.

#### Add attachments


21. Click **Add Attachments** from the context menu.  
The **Add Attachments** window appears.
22. Type a comment for the attachments in the **COMMENT TEXT** box.
23. Click **Choose File**.
24. Browse and select the file you want to attach.
25. Click the **Attach another file** link to add more attachments. Repeat this step for adding more attachments.
26. Click **Add** to attach the selected files to the artifact.

#### Clone an artifact

27. Click **Clone** from the context menu.  
The **Clone Artifact** window appears.
28. Provide a name and description for cloning the artifact.
29. Click **Clone** to clone the artifact.

#### Schedule work on an artifact



To plan how and when an artifact is to be addressed, assign it to a planning folder.

 **Note:** A planning folder can contain artifacts from many different trackers. When you assign an artifact to a planning folder, the artifact is still in the tracker where it was created.

1. Click **TRACKERS** from the **Project Home** menu.
2. Find the tracker that contains the artifact you want to assign, and select the artifact.  
You can select any number of artifacts at once.
3. Click **Plan for** and select a planning folder.

The planning folder you assigned is now shown in the **Planned For** column on the **Planning Folder Summary** page.

You can always reassign an artifact to another planning folder as conditions change. Any effort data you have provided is recalculated automatically.

-  **Note:** When a tracker is disabled, artifacts from that tracker do not contribute to the effort totals calculated for any planning folder they are in.
-  **Tip:** When you are working with a single artifact, you might prefer another way to assign it to a planning folder: open the artifact you want to assign and look for the **Planning Folder** field. Select a planning folder for the artifact and save the artifact.

### Set up trackers




A tracker is a collection of related artifacts that describe work to be done or issues to be resolved. Every project should have one or more trackers.

When you start a tracker, you decide which fields will be used, who will use them, and how they will use them.

#### *Create a tracker*

Create a tracker whenever you need to report and track bugs, feature requests, support requests, or any other type of issue where ownership, status, and activity must be managed.

Individual tracker entries are referred to as tracker artifacts. A tracker is a set of tracker items with a common purpose, such as bug reports, feature requests, or tasks.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Tracker Settings**.
3. Click **Create**.
4. On the **Create Tracker** page, provide a name and description for the tracker.
  -  **Tip:** Descriptions help users learn how best to provide the information you want from them. To maximize your chances of getting useful data, make your description as informative as you can.
5. Select an icon that suggests the type of work the tracker is handling.  
This icon will appear with any artifact in this tracker, wherever it is viewed on the site. For example, if someone brings an artifact from this tracker into a planning folder, users of the planning folder can glance at the artifact's icon to see where it comes from.
6. Select the relevant unit from the **DISPLAY EFFORT IN** field. The units displayed here are configured based on the size of the artifacts in the tracker. Eg. Select the unit as **HOURS** for a tracker of small defects, **DAYS** for a tracker of tasks, and **WEEKS** for a tracker of epics.
  -  **Note:** Configure the units at the project level and not at the planning folder level.
7. Select **INCLUDE FOREIGN CHILDREN** to include points and efforts from children artifacts across the projects in TeamForge.
  -  **Note:** In a parent artifact, enabling **CALCULATE POINTS** field sums and rolls up the points from all its children artifacts within the project. In this total, if you want to include children artifacts from other projects across TeamForge, have the **INCLUDE FOREIGN CHILDREN** option enabled.
8. Click **Create**.  
The new tracker appears at the bottom of your list of trackers.
9. If necessary, drag the tracker to a place in your tracker list that makes sense.  
The order you set here controls the order of every tracker list in your project.
10. You'll probably need some custom fields to capture information that's specific to your project.  
See [Create your own tracker fields](#) on page 54.

11. To speed up the team's work, you may want to set up some rules for automatically reassigning artifacts when their contents change.


See [Create a tracker workflow](#) on page 55.

### Clone a tracker

To save efforts in duplicating a tracker within the project and across the projects, choose cloning options available in TeamForge.

You can clone a tracker within a project along with the workflow. This can be done across the projects as well. If you are the Project Admin for the source and destination projects, you can create new roles that are not available in the destination project. However, the permissions associated with the role are not copied from the source project. A tracker admin cannot create new roles while cloning a tracker across projects.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. From the **Project Admin Menu**, click **Tracker Settings**.
3. Select **Clone External Tracker** from the drop-down list available in **Clone** button to clone a tracker from the source project.

 **Note:** To clone a tracker within the project, you can select the tracker from the **Tracker Settings** page and click **Clone Tracker**.

4. On the **Cloning Tracker** page, name the new tracker and enter a suitable description.
5. Enter **SOURCE TRACKER ID** of the tracker available in the source project and click **Next**. The **Cloning External Tracker** page appears.
6. On the **Cloning External Tracker** page, name the new tracker and enter a suitable description.
7. Click **Create**.


The cloned tracker appears at the bottom of your list of trackers.

### Enable or disable fields


If a tracker field is disabled, it does not appear on the **Artifact** page. Most fields can be disabled.

Disabled fields are accessible only to tracker administrators on the **TRACKER FIELDS** tab.

You can enable or disable any field that is user-defined or configurable. However, you can't disable all configurable fields. For example, the **Status**, **Priority**, **Category**, and **Planning Folder** fields can't be disabled.

 **Note:** If your goal is to prevent users from entering data into a field when submitting an artifact, but still display the field on the **Edit Artifact** and **Tracker Search** pages, don't disable the field. Instead, clear the **DISPLAY ON SUBMIT** option on the **Edit Tracker Field** page.


1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Tracker Settings**.
3. From the list of current trackers, click the tracker you want to configure.
4. On the **TRACKER FIELDS** tab, select the fields you want to enable or disable.
  - Click **Disable** to remove them from the **Artifact** page.
  - Click **Enable** to allow them to be configured and displayed.


 **Note:** Data in disabled fields is still search-able, but disabled fields do not appear as inputs on the **Search** pages.

### Require users to complete a tracker artifact field

If a field is set as required, users cannot submit artifacts without completing it. Most tracker fields can be required or optional.

Each tracker can have its own required and optional fields. Required fields are marked with a red asterisk \* on the **Submit Artifact** page.

 **Note:** When you make a field required, any field whose values depend on that field's values is also required. See [Help users select options in a tracker artifact](#) on page 149 for more information about dependent field values.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
  2. Click **Tracker Settings**.
  3. From the list of current trackers, click the one you want to configure.
  4. On the **Tracker Fields** tab, click the name of the field you want to set as required or optional.  
By default, only the **Title**, **Description**, and **Status** fields are set as required.
  5. On the **Edit Field** page, select or clear the **Required** option to make a field required or optional.  
Required fields automatically show up on the **Submit Artifact** page.
-  **Note:** System-defined fields and the **Status** field are always required.
6. For optional fields, select or clear the **DISPLAY ON SUBMIT** option.  
This specifies whether the field will appear on the **Submit Artifact** page.
  7. Click **Save Field**.

#### *Create your own tracker fields*

To track data that is not captured by the default set of fields, create new fields that fit your project's purposes.

You can create the following user-defined fields in each tracker:

- Up to 30 text entry fields.
- Up to 30 date fields.
- Up to 30 single-select fields.
- An unlimited number of multiple-select fields.



**Caution:** Creating a large number of user fields and multiple-select fields may affect site's performance.

#### *Configure tracker "Select" field values*

To help users provide meaningful information, supply them with useful field values to choose from in the input fields in the tracker entry form.



**Tip:** Once a tracker has been created, you may:

- Create one or more user-defined single-select or multiple-select fields.
- Add predefined values to the fields. Remove values, if required.
- Enable or disable fields.
- Change the default values for fields.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Tracker Settings**.
3. From the list of existing trackers, click the name of the tracker that you want to configure.
4. On the **TRACKER FIELDS** tab, click the name of the field whose values you want to edit.
5. On the **Edit Field** page, set up the field values you want users to see when they create a tracker artifact.
  - To define a new value, click **Add**.
  - To rename a value, edit the existing text. If you rename a value, the value is renamed in all existing artifacts.
  - To remove a value, check the box and click **Delete**. If you delete a value, the value is changed to **None** in all existing artifacts.
  - Select **DEFAULT VALUE** to set which option will be chosen if the user makes no selection. When you move a tracker artifact from one tracker to another, the default field value is the value that comes along.



**Note:**

- When you edit the values of the **Status** field, you are also asked to describe what each value's status means, as shown in the **Values** section of the **Edit Field** page. This status meaning is used in **Advanced Search** to define which values are returned when searching for open or closed artifacts.

- As always, when you create a new tracker, the default value for the '**Priority**' field is set as '4 - Low'. However, you can change the default value by editing this configurable single select field, '**Priority**'. You cannot delete or disable the **Priority** tracker field.
- When you change the tracker fields, the values in the existing artifacts remain unchanged.

6. Click **Move Up** or **Move Down** to order the list the way you want it.

7. Click **Save Field**.

All values are now available in the selection menu for the field.

#### *Assign a tracker artifact automatically*

You can configure the tracker to automatically assign newly submitted artifacts to specific project members.

You can assign artifacts to individuals based on the values in the **Category** or **Release** field.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Tracker Settings**.
3. From the list of existing trackers, click the tracker you want to configure.
4. On the **AUTO ASSIGNMENT** tab, select **Auto Assign By Category** or **Auto Assign By Release**.
5. For each category or release value, choose a project member from the **AUTOMATICALLY ASSIGN TO** drop-down menu.

This menu contains all project members with the tracker edit permission. Click the Search icon to display a list of project members to whom you can auto-assign artifacts. Choose **None** if you do not want artifacts with a specific value automatically assigned.

6. Click **Save**.

Whenever a new artifact is submitted with a value in the relevant field, the assignee receives an email notification and the artifact appears on the assignee's **My page**.

#### *Change a tracker*

A tracker's real-world uses often outgrow the name or description you gave it when you created it. When that happens, it's a good idea to update the tracker to reflect its changing role.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Tracker Settings**.
3. From the list of existing trackers, click the tracker you want to edit, and click **EDIT**.
4. On the **Edit Tracker** page, provide a new name or description for the tracker, and update the icon.
5. Update the units from **DISPLAY EFFORT IN** and click **Save**.



**Note:** These units are configured in the **Units** page at the project level, and not at the planning folder level.

6. Select **INCLUDE FOREIGN CHILDREN** to include points and efforts from children artifacts across the projects in TeamForge.



**Note:** In a parent artifact, enabling **CALCULATE POINTS** field sums and rolls up the points from all its children artifacts within the project. In this total, if you want to include children artifacts from other projects across TeamForge, have the **INCLUDE FOREIGN CHILDREN** option enabled.

7. If necessary, drag the tracker to a place in your tracker list that makes sense.

The order you set here controls the order of every tracker list in your project.

#### **Create a tracker workflow**


To channel project members' work on tracker items, set up rules for how a tracker item can move forward.

Before creating a tracker workflow, see that these criteria are met:


- You have a tracker to work with.
- The tracker has a set of statuses defined, such as "In progress" and "Ready for QA."
- Roles exist, and you can assign project members to them.

A workflow is a sequence of changes from one status to another. You can define status transitions for any combination of tracker statuses in the tracker.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Tracker Settings**.
3. From the list of existing trackers, select a tracker.
4. Click the **WORKFLOW** tab.  
The **Workflow** page lists all of your status values and the tracker workflow that you have configured.
5. On the **Workflow** page, click the status value for which you want to create a workflow.
6. On the **Edit Field Transition** page, select a status value from the **Create Transition to Status** drop-down menu.

 **Note:** By default, the **Any** workflow is shown. This means that you are defining workflow rules for the transition from the status that you selected to *any* other status.

7. Click **Add**.  
A new workflow is added. The **Any** workflow is changed to **Remaining Statuses**.
8. In the **ROLES** section, specify which users can make this change.  
For example, only users with the QA Engineer role are allowed to change artifacts from **Open** to **Cannot Reproduce**.
9. Click the **Advanced Transition** link.

 **Note:** For any new unsaved transition, an alert is shown asking you to save the transition so that you can view it in the **Transition Status To** drop-down list. Click **OK** to configure already saved transitions or click **Cancel** and then click **Save and View** in the **Edit Field Transition** page to save the unsaved transitions.

### Advanced Transition ✕

Transition Status To Under Consideration -> Completed

NAME	TYPE	REQUIRED	VISIBLE
Group	Single Select	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Status	Single Select	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Category	Single Select	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Customer	Single Select	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Priority	Single Select	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Team	Folder	<input type="checkbox"/>	<input checked="" type="checkbox"/>

Cancel
Save And Close
Save

- a) Select the transition workflow for which you want to apply Advanced Transition settings from the **Transition Status To** drop-down list.
- b) Select the **REQUIRED** check box against the fields for which the user *must* provide values.



For example, the user must assign the tracker item to someone and enter a comment.

 **Note:**

Fields whose values depend on a required parent field are automatically required. (See [Help users select options in a tracker artifact](#) on page 149 for more information on parent and child fields.)

- c) Select or unselect the **Visible** check box for showing or hiding fields respectively for the selected status transition.

Points to note:

- **Required** fields are always visible.
- Advanced Transition rules are applied when you create or edit artifacts in a tracker and only when you edit artifacts in Planning, Task, and Kanban boards.
- Hidden field values, if updated via SOAP/REST APIs, are ignored.

#### 10. Click **Save**.

The workflow is now saved. When a user submits or edits the status of a tracker artifact, he or she sees only the options that are allowed by the workflow.

#### **Import tracker artifacts**

You can now import artifacts into TeamForge using the Excel/CSV tracker import function. Data from from both Excel and CSV files can be imported.

- You must have Site Administrator permission or Tracker Submit and Edit permission to import artifacts from Excel/CSV files.
- The Excel/CSV template can be downloaded from the Import Artifact window.
- If you want to create new artifacts, the following tracker fields are absolutely required while importing data from Excel/CSV files: Artifact ID, Title, Priority, Status and Description. However, for updating an existing artifact, a valid "Artifact ID" is all that's required along with the data for other fields, which you may want to update.
- The tracker import function supports upto 500 rows of data in Excel/CSV files. However, you can configure the number of artifacts that can be mass-imported. For more information, see [Configure your site's settings](#) on page 365.

1. Click **TRACKERS** from the **Project Home** menu.
2. Select a tracker from the **TRACKERS** tab.
3. Click **Import**.

TeamForge / Trackers / Tasks / List Artifacts

tracker22419 : Tasks HIDE

Name: Tasks  
Description: Tasks usually created as child dependencies of user stories, defects or other "backlog item" level trackers.  
Effort Units: Hours

Open by priority

Tasks Page 1 of 13 (194 Items)

PRIORITY	ARTIFACT ID : TITLE	ASSIGNED TO	TEAM	SUBMIT
HIGHEST	artf285105 : Forklift changes for document flex fields	SathishKumar Rathna Sabapathy	Services Team	Anand
HIGHEST	artf282643 : Run AppScan against IT 4 build	Venkat Thandi Sivagangai	Mavericks - CRUS	rajesw
HIGHEST	artf282050 : Run AppScan against 16.10.410 build (IT 3 build)	Venkat Thandi Sivagangai	Mavericks - CRUS	rajesw

Buttons: Monitor, Import, Export, Cut, Edit, Edit Inline, Mass Up

- Click **CHOOSE FILE**. Browse and select the Excel/CSV file that contains the artifacts to be imported.

Import Artifacts

Import artifacts to Defects Tracker from Excel or CSV files. You can use the tracker template ([Excel](#), [CSV](#)) to import artifact data.

FILE\* CHOOSE FILE

SAVE IMPORT FILE AS DOCUMENT

Cancel Import

Tracker Information

Instructions

- Tip:** You can select the **SAVE IMPORT FILE AS DOCUMENT** check box if you want to have the imported Excel/CSV file stored in **DOCUMENTS** for future use.

- Click **Import**.

### Limitations

It's worth considering the following points while importing artifacts from Excel/CSV files:


- The Excel/CSV tracker import function allows importing data for any existing tracker fields even if the fields are disabled in the tracker.
- The **Import** button is not being localized on Chinese and Korean locales.
- The import instructions on the **Import Artifacts** page is not being localized on Chinese, Japanese and Korean locales.

- The Estimated Effort, Actual Effort, Remaining Effort, and Points fields can take "0" or any positive integer as values.
- Leave the "Comment text" field empty in the Excel/CSV files if you intend to create new artifacts.
- While it is expected that the "Priority" for an artifact can range from 0 and 5, the Excel/CSV import function allows any value for the "Priority" field to be imported. Make sure your Excel/CSV input file consists of "Priority" values ranging from 0 and 5.

### Export tracker artifacts

To use the contents of artifacts from a tracker or a planning folder in other applications, export them to a .csv, .xml, .xlsx or tab-delimited file.

For example, sometimes it can be useful to look at the status of a project by sharing a spreadsheet using Microsoft Excel or Google Spreadsheets.

 **Tip:** You can also export tracker artifacts from a tracker report or a set of search results.


1. Click **TRACKERS** from the **Project Home** menu.
2. Click a tracker or planning folder.

In a tracker,

- you can export a subset of the artifacts in the list view by setting up a filter. See [View the right tracker artifacts](#) on page 33.
- You can select specific artifacts and export only those checked items.


The current sort order in the list view is retained in the exported list of artifacts as well.

3. Click **Export**.
4. In the **Export Data** window, select an export format that you can import into the other application. For example, to use the data in a spreadsheet program, select .csv.
5. Move the fields you are interested in from the **AVAILABLE COLUMNS** list to the **SELECTED COLUMNS** list, then click **Export**.

 **Note:** When you export tracker artifacts directly from a tracker, you can choose from all the fields in the tracker. When you export from a planning folder, only the standard fields provided by TeamForge can be exported.

For more detail on what you can do with different kinds of fields, see [What fields can I use in a tracker?](#) on page 1050

6. Artifacts are exported and the file is downloaded automatically.

 **Tip:** Starting from TeamForge 16.10, while exporting tracker artifacts, artifacts are exported to one of the file formats such as Excel, CSV, Tab-delimited or XML (as selected by the user) and the file is downloaded automatically. To improve performance, the file download link, which you would have used in the past, is no longer available.

### Boards

You can start using the task and kanban boards, once your project administrator has configured them.

#### Use the Task Board

During a sprint, TeamForge project members can use the Task Board to view tasks, create tasks for backlog items, edit tasks and drag and drop tasks across swimlanes as they progress.

With the Task Board, you can:

- View backlog items of a selected planning folder in the **Backlog Items** swimlane, which is the leftmost swimlane of the task board.
- View tasks (belonging to backlog items) pinned to swimlanes based on the status.
- Edit backlog items and tasks by clicking the artifact ID link.
- Move tasks (cards) from one swimlane to the other as tasks progress.
- Add new tasks for backlog items.

1. Click **TRACKERS** from the **Project Home** menu.

- In the **List Trackers, Planning Folders and Teams** page, click **TRACK**.



A Task Board for the current project context is displayed.


- Select a planning folder from the drop-down list.

The Task Board is populated with backlog items and their tasks.

- If a selected planning folder has subfolders, the artifacts within the subfolders are not displayed.
- Cards have a color-coded bar to visually identify the priority. In addition, the open and closed cards' background is color-coded to uniquely identify the status.
- For each card, the artifact ID, title, priority, assigned to and effort (in terms of 'Points' for backlog items and 'Hours' for tasks) are displayed. Status information is shown on backlog cards only. Relevant tooltips appear when you hover your mouse over these data elements. You can also hover your mouse over the tracker icon in backlog and task cards to know the tracker name.

### Add new tasks

- Click the '+' icon.



The **Create Task** window appears.


- Type a title and description and click **Submit**.



**Tip:** To use the title you type as the description, select the **Use Title as Description** check box.

### To edit a backlog item or task using the task ID link:

- On the **List Trackers, Planning Folders and Teams** page, select the planning folder.
- From the listed backlog items and tasks, click the artifact ID link you want to edit.



The **View Artifact** page appears.

- Modify the backlog item or task and click **Update** to update the artifact and return to the Task Board.

### To quickly edit a task

- Click the edit artifact button in a task. The **Edit Task** window appears.



**Note:** To be able to edit tasks, you need to have the required permissions.

### Edit task button



**Edit Task window**

10. You can edit the following on the **Edit Task** window:

- **Artifact title**
- **Artifact description**
- **Primary and secondary attributes**

You have two expandable frames, **Primary Attributes** and **Secondary Attributes**, that list fields you can edit. Primary attributes include Priority, Status, Assigned To, Estimated Effort, Remaining Effort, Actual Effort, Points and Team. The expandable frames and the fields listed in them depend on your application configuration that you set in Tracker Settings. For example:

- Configurable fields which are enabled. For more information, see [Enable or disable fields](#) on page 53.
- User-defined fields (flex fields) with the field type Text Entry and Single Select and set as **Required** (Tracker Settings > Tracker Administration > Create Field). For more information, see [Create your own tracker fields](#) on page 54.
- Fields which are workflow configured. For more information, see [Create a tracker workflow](#) on page 55.

By default, the following field types, which you select while creating a user-defined field (Project Admin > Tracker Settings), are non-editable from the task board's **Edit Task** window: Multiple select (**Multiple Select**), Date picker (**Select Date**) and User picker (**Select User(s)**).


- **Comments**

In this expandable frame, in addition to adding comments, you can view the last five comments for an artifact.


11. After editing an artifact, click **Update**.

#### Move tasks across swimlanes

12. To move tasks across swimlanes, drag a task card from the source swimlane and drop it on the destination swimlane.

 **Note:** You need to have task tracker 'edit' permission to move cards across swimlanes.

#### View both 'open' and 'closed' backlog items

13.  **Note:** When you select a planning folder, the Task Board shows only 'open' backlog items by default.

Click the 'Show/hide closed artifacts' toggle button to view both 'open' and 'closed' backlog items.

#### Show/hide closed artifacts toggle button



#### Refresh the Task Board

14. Click the following 'Refresh' button to refresh the Task Board.



#### Auto assign task


If a task is not assigned to anyone, that is, if it is assigned to **None**, then it can be assigned to the user who has logged in using the **Auto Assign Task to Me** feature. This check box appears only if:

- The Task Board is configured.
- The user who logged in has the edit permission to the task tracker.
- The project is not locked.

15. Select the **Auto Assign Task to Me** check box and drag the artifact from one swimlane to another (from one status to another) and the artifact is automatically assigned to you.

#### Position artifact cards

16. If you have the requisite permission, you can drag and drop a card above or below the other cards within a swim lane to position it in the order you want the artifact addressed.

 **Note:** The positioning of cards is retained when you switch between the list and board views (planning, task and kanban).

For related information on ranking artifacts in the list view, see [Reorder the contents of a planning folder](#) on page 176.

#### Use Kanban Board

After the kanban board is configured, you can use it to view, plan and track work items for the selected planning folder or the selected team within that planning folder.

1. Click **TRACKERS** from the **Project Home** menu.
2. In the **List Trackers, Planning Folders and Teams** page, click **Kanban**.



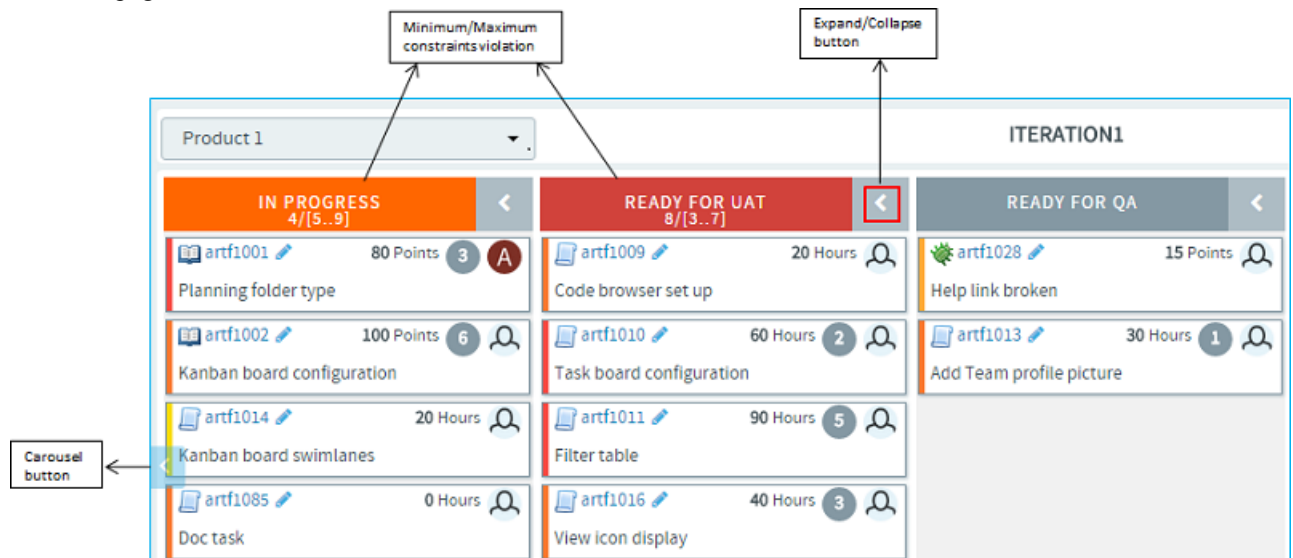
3. Select the required kanban board from **Manage Boards** (available only for project administrators). If you are a project member, select your kanban board from the **Select Board** list. The selected kanban board for the current project context is displayed.

4. Select the planning folder for which you want to view the status of the artifacts.


If your project contains teams, the **Select a Team** drop-down list appears. From this drop-down, select **All artifacts** or if you want to view artifacts specific to a team within the selected planning folder, select the relevant team.

Depending upon your selection, artifacts pertaining to the mapped trackers are displayed in the appropriate swimlane.

- You can view a maximum of 5 swimlanes at a time on your Kanban Board. If there are more, use the carousel scroll to slide across the swimlanes. You can expand or collapse each swimlane. When there are no results to display in a swimlane, the **Expand/Collapse** button does not appear.
- Based on the configuration values, if the constraints are violated, the relevant status headers are highlighted appropriately. A violation of minimum constraint is highlighted in orange indicating that resources are underutilized whereas that of maximum constraint is highlighted in red indicating resources being overloaded and bottlenecks to be fixed. For more information on how a kanban board is configured, see [Set up Kanban Board](#) on page 193.



- Each swimlane header displays the status label, the total number of artifacts within the selected planning folder, and the minimum and maximum constraints that were configured for the kanban state.




In the above scenario, for the 'In Progress' status:

'3' is the total number of artifacts in the 'In Progress' status within the selected planning folder.

'2' indicates the minimum constraint and '\*' indicates the maximum constraint '-1'. (-1 as the minimum constraint translates to 0 whereas -1 as maximum indicates there is no limit and so represented by an \*.)

- Important:** These values and constraint validation are planning folder specific and not team specific. For example, in the above scenario, the total number of artifacts is the overall count of artifacts which are in the 'In Progress' status within the selected planning folder and not within the selected team.

- Each artifact card displays its points (story points) or estimated effort if these fields are enabled in your tracker settings. If both are enabled, only the estimated effort is displayed hovering the mouse on which the artifact's points is displayed. Similarly, if all the effort fields (estimated, remaining and actuals) are enabled, you can view them by hovering the mouse on the estimated effort.
- Cards have a color-coded bar to visually identify the priority. In addition, the open and closed cards' background is color-coded to uniquely identify the status.
- A parent artifact card displays the count of its child artifacts (**Show/Hide child artifacts** button). Click this button to view or hide the child artifacts. If there are more than 5 child artifacts, when you expand the parent artifact, you will see pagination arrows ('Previous'/Next') at the bottom of the parent artifact card.



- If the tracker types of both the parent and child artifacts have been mapped with the kanban statuses, the child artifacts appear within their parent artifact card and also as an individual card. In the above scenario, Epic (parent artifact) and Story (child artifact) trackers have been mapped to the kanban status 'In Progress'. Story 'artf1014' is the child artifact of the epic 'artf1002'. So you will see 'artf1014' within the parent artifact card as well as outside of it as an individual artifact card.
- When a child artifact is closed, a 'Closed' tag appears next to the child artifact ID within the parent artifact card.


### Edit artifacts

Use Kanban Board to edit an artifact using the Edit icon on the artifact card or move artifacts from one status to another appropriately. Based on the status changes you make, the swimlane headers get updated appropriately.

#### Remember:

- You can edit an artifact only if you have the tracker edit permission. Otherwise, you can only view it.

**Artifacts support @mentions:** Artifact description and comments now support @mentions and users called out via @mentions are added to the monitoring list. Include usernames with "@" as prefix (for example, @mphppard) to add users to the monitoring list.


**Note:**  Users called out via @mentions must have "Artifact View" permission to be added to the monitoring list.

5. When you drag the artifacts from one swimlane to another, note the following validations:

- The changes you see in swimlane headers do not apply to any specific team but takes into account the total count of artifacts within the selected planning folder.




- You cannot move an artifact card to a status that is not mapped to that particular tracker type. For example, a 'Ready for QA' kanban status may not have been mapped to any of the epic tracker statuses. So when you attempt to move an epic to that status, you will get an `Invalid status configuration` error.
- You can move a parent artifact to the 'Closed' status only if all its child artifacts are closed. However, a child artifact is independent of its parent with regard to the status change, that is, when you move a parent artifact from one status to another, the status of the child may still remain unchanged. For example, an epic may have many stories, the status of some may change from 'Not Started' to 'In Progress' whereas some may not. So when the status of an epic as one single unit may change, it need not apply to all of its child artifacts.
- When your move violates a constraint (minimum or maximum), a warning is displayed; but you can still move the artifacts.
- If a kanban status is mapped to more than one tracker status, when you move an artifact, you have an option to choose a status as shown in the following screen shot.



### Create artifacts and child artifacts


Using the Create Artifact icon on the top right of your kanban board, you can quickly create an artifact. The Create Artifact icon appears only when you configure a kanban board and select a planning folder.

Similarly, using the Create Child Artifact icon available on the artifact card, you can quickly create a child artifact for any artifact card displayed on Kanban Board.

 **Note:** The Create Child Artifact icon does not appear on closed artifact cards.

You can create an artifact or a child artifact only if you have the required tracker permission.

**Artifacts support @mentions:** Artifact description and comments now support @mentions and users called out via @mentions are added to the monitoring list. Include usernames with "@" as prefix (for example, @mphippard) to add users to the monitoring list.

 **Note:** Users called out via @mentions must have "Artifact View" permission to be added to the monitoring list.

6. Click the required icon: Create Artifact or Create Child Artifact.



7. Enter the required information in the relevant window and click **Submit**.

- Only trackers configured for the kanban board appear in the Trackers drop-down list.
- While you quickly add artifacts with data for just three fields such as the tracker type, title and description, the artifacts are, however, saved with default values for other required fields, which you may choose to update later. If the default state of the selected tracker is not mapped for the newly created artifact or child artifact, the artifact card does not show up on the kanban board.

### Show / Hide closed cards

8. When you have a large volume of closed cards in a planning folder, you can restrict the number of closed cards you want to view by toggling between the 'Show all closed artifacts' icon:



Or the 'Hide artifacts older than 60 days' icon using which you can hide closed cards older than 60 days. This is the default option:



**Note:**

- The Show/Hide toggle icon appears only on configured kanban boards.
- Your selection is saved for the subsequent sessions as well.

### Position artifact cards

9. If you have the requisite permission, you can drag and drop a card above or below the other cards within a swim lane to position it in the order you want the artifact addressed.



**Note:** The positioning of cards is retained when you switch between the list and board views (planning, task and kanban).

For related information on ranking artifacts in the list view, see [Reorder the contents of a planning folder](#) on page 176.

### Work on tasks

Manage your project plan by defining units of work and assigning them to project members.

#### Create a task

Create a task when you identify a unit of work that you want to track to completion as part of your project.

1. Click **TASKS** in the project navigation bar.
2. On the **Task Summary** page, go to the task folder in which you want to create the task.
3. Click **Create**.
4. On the **Create Task** page, provide a name and description for the task.
5. Assign the task a priority between 1 (highest) and 5 (lowest).
6. Assign the task to an individual.  
Click the people icon to list the project members to whom you can assign the task.
7. Assign the task a status.
  - Not Started
  - OK
  - Warning
  - Alert
  - Complete
8. Give the task a start and end date.  
Click the calendar icon to select dates.
9. In the **ESTIMATED HOURS** field, write the time you think is needed to complete the task.
10. Specify whether the time estimate is based on a five-day or seven-day calendar.



**Note:** Choose a seven-day calendar to assign a start or end date that falls on on a weekend.

11. Click **Save**.

#### Edit a task

Tasks often change assignee, start and end dates, and dependencies before they are finished. You can update all those aspects of a task, along with its status.

1. Click **TASKS** in the project navigation bar.
2. On the **Task Summary** page, find the task folder containing the task you want to edit.
3. Click the task name.

4. On the **View Task Details** page, click **EDIT**.
5. On the **Edit Task** page, make your changes and click **Save**.

If you change the end date of a task that has dependent successor tasks, you may want to change the start and end dates of the successor tasks. For example, if a successor task normally starts on a Friday, and is moved out two days by a change in a predecessor task, its new start date is Sunday (if it uses a seven-day calendar) or Tuesday (if it uses a five-day calendar).

- To change the start and end dates of all successor tasks, click **Move Dependent Task Dates Forward/Backward x Days**. All dependent task dates are moved forward or backward by the same number of days that the original task end date was moved.
- To apply the date change to only the selected task, without changing the start and end dates of any successor tasks, click **Do Not Change Dependent Task Dates**.
- If your end date change impacts successor tasks that you do not have permission to view, these tasks are not changed.
- If your end date change impacts successor tasks that you have permission to view, but not to edit, these tasks are marked for change approval by an administrator. The start and end dates for tasks that you have permission to edit are changed.
- Successors of tasks that are marked for change approval are not changed until approved by an administrator.
- From this page, you cannot change the start and end dates of successor tasks selectively. You must change all or none of them. To edit the start and end dates of selected successor tasks only, select **Do Not Change Dependent Task Dates**, then edit each selected task manually.


### Find tasks

Use the task manager to find a specific task or set of tasks quickly. You can filter tasks by predefined fields or search by keyword, date or other values.


#### Search for tasks

Use the Search Tasks option to refine your search if task filtering returns too many results.

1. Click **TASKS** in the project navigation bar.
2. On the **Task Summary** page, click **Search Tasks**.
3. On the **Task Search Criteria** page, enter the keywords for the search.

 **Tip:** Wildcards are allowed.

4. Search by the type of content that makes sense to you.
  - Search for keywords in the **Task Name** field.
  - Search for keywords in the **Task Description** field.
  - Search the task progress fields: **Accomplished**, **Planned**, and **Issues**.
  - If you are searching for a specific task, and you know its task ID, enter the task ID.
  - To search by date, select date ranges for start date or end date.
  - To search by project member, select one or more people.

 **Note:** Click the people icon to display a list of project members with permissions to edit or create tasks.

- To search by location, select the Task Folder to search.
- To search by Status and Priority, select those.
- Click **Search**.

#### See the right tasks

Filtering and sorting allows you to find a task or group of tasks using a number of pre-defined fields.

The filter option is available on every item list view in TeamForge.


1. Click **Tasks** in the project navigation bar.
2. On the **Task Summary** page, click the title of the task folder that contains the tasks you are interested in.
3. Specify the filter criteria in one or more filter fields (at the top of each column) and click **FILTER**.

- You can find a filter field at the top of each column in most of the tables in the TeamForge application.
- The filter field could be a text box or a drop-down list with multi-select check boxes.

The screenshot shows a table titled "Stories" with "Page 1 of 40 (600 Items)". The table has columns for PRIORITY, ARTIFACT ID : TITLE, ASSIGNED TO, TEAM, and SUBMITTED BY. Each column has a filter field at the top. The filter fields are: PRIORITY (Any), ARTIFACT ID : TITLE (Filter by), ASSIGNED TO (Filter by), TEAM (Any), and SUBMITTED BY (Filter by). The table contains 11 rows of data, each with a priority level (e.g., MEDIUM, LOW, HIGH), an artifact ID and title, an assigned person, a team, and a submitted by person. At the bottom of the table, there are buttons for Monitor, Import, Export, Cut, Edit, Edit Inline, and Mass Update.

PRIORITY	ARTIFACT ID : TITLE	ASSIGNED TO	TEAM	SUBMITTED BY
3 MEDIUM	artf284437 : Create UI for viewing all diffs for entire review on one screen	C. Michael Pilato	SCM	Mark Phippard
3 MEDIUM	artf277197 : Workspace widget for artifacts	Baybora Aksoy	SCM	Mark Phippard
4 LOW	artf246564 : Reports - write protractor tests to cover e2e tests.	None	Reports (Datamart)	Muhammed Ashik
2 HIGH	artf287743 : Include 'Artifact Tags' in Project Templates	Ramkumar Rathinasamy Viveganandan	TF Core > Avengers - Tags	Aishwarya Rajan
2 HIGH	artf287742 : Tags in Edit Artifact Page	Ramkumar Rathinasamy Viveganandan	TF Core > Avengers - Tags	Aishwarya Rajan
2 HIGH	artf283822 : Synchronize SVN Indexer execution	akil roy	Devops/Releng	Johannes Passing
2 HIGH	artf278265 : Cleanup the old EventService	Nithya Pattabiraman	TF Core > Jaguars - Architecture	Sriram
3 MEDIUM	artf288250 : EventQ meta RPM	Sinhuja Harikrishnan	Devops/Releng	Sinhuja Harikrishnan
3 MEDIUM	artf287737 : Update the default/trial license to include DevOps (and perhaps new license types)	Gowthaman Ravindran	Mavericks - CRUS	Mark Phippard
3 MEDIUM	artf287193 : SVN load balancing	akil roy	Devops/Releng	Johannes Passing
3 MEDIUM	artf286210 : Monitor Snapshot before restart	h.karthik	Devops/Releng	Johannes Passing

- You can type your filter criteria in the text boxes. The search text is case-insensitive.
- You can also select the filter values from one or more drop-down lists. By default, you can only select up to 10 filter values in a drop-down list. However, you can set a value that suits your requirement for the `FILTER_DROPDOWN_MAX_SELECTION` token in the `site-options.conf` file to increase or decrease the count.
- **Filter-as-you-type:** You can find the **Enter keywords** text box in all filter drop-down lists. As you type your filter keyword, instant search results are shown in the drop-down list. For example, in the following illustration, typing "R" instantly shows all statuses having the alphabet "R". The search text is case-insensitive.




- Some search filters may not appear if your site administrator has not enabled them.
4. After filtering, if you want to clear the filters, click **FILTER** and select **Clear** from the drop-down list.
  5. Use the up-down arrow at the top of any column to sort your list by that column.
    - Your primary sort column is identified by a superscript **1** next to the up-down arrow, and your secondary and third-level sort columns, if any, are likewise marked.
    - Click the up-down arrow again to reverse the sort order.

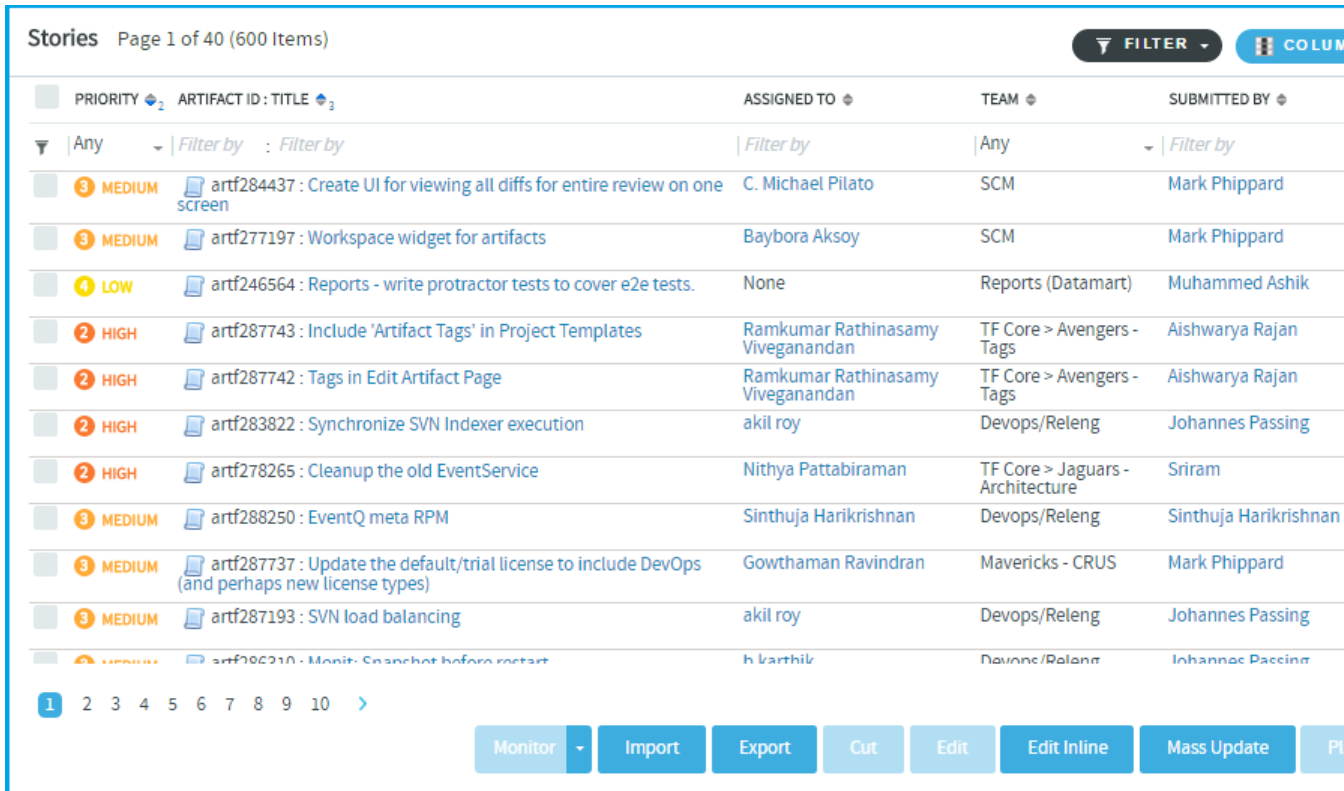
### Check your assigned tasks

It's a good practice to regularly look at all of your assigned tasks in the **My Tasks** section of your **My Page**.

Your **My Page** provides a central location where you can view and manage all tasks assigned to you from all the projects that you are a member of.

 **Note:** To see tasks assigned to you on a specific project, look at the **Tasks** section of that project.

1. Go to **My Page**.
2. If you have more than 20 tasks, click **View All** to list all of them.
3. Specify the filter criteria in one or more filter fields (at the top of each column) and click **FILTER**.
  - You can find a filter field at the top of each column in most of the tables in the TeamForge application.
  - The filter field could be a text box or a drop-down list with multi-select check boxes.

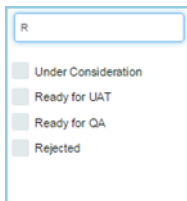


The screenshot shows a table of stories with the following columns: PRIORITY, ARTIFACT ID : TITLE, ASSIGNED TO, TEAM, and SUBMITTED BY. The table is filtered to show 10 items on page 1 of 40. The filter criteria are: PRIORITY: Any, ARTIFACT ID : TITLE: Filter by, ASSIGNED TO: Filter by, TEAM: Any, SUBMITTED BY: Filter by. The table contains the following data:

PRIORITY	ARTIFACT ID : TITLE	ASSIGNED TO	TEAM	SUBMITTED BY
3 MEDIUM	artf284437 : Create UI for viewing all diffs for entire review on one screen	C. Michael Pilato	SCM	Mark Phippard
3 MEDIUM	artf277197 : Workspace widget for artifacts	Baybora Aksoy	SCM	Mark Phippard
4 LOW	artf246564 : Reports - write protractor tests to cover e2e tests.	None	Reports (Datamart)	Muhammed Ashik
2 HIGH	artf287743 : Include 'Artifact Tags' in Project Templates	Ramkumar Rathinasamy Viveganandan	TF Core > Avengers - Tags	Aishwarya Rajan
2 HIGH	artf287742 : Tags in Edit Artifact Page	Ramkumar Rathinasamy Viveganandan	TF Core > Avengers - Tags	Aishwarya Rajan
2 HIGH	artf283822 : Synchronize SVN Indexer execution	akil roy	Devops/Releng	Johannes Passing
2 HIGH	artf278265 : Cleanup the old EventService	Nithya Pattabiraman	TF Core > Jaguars - Architecture	Sriram
3 MEDIUM	artf288250 : EventQ meta RPM	Sinhuja Harikrishnan	Devops/Releng	Sinhuja Harikrishnan
3 MEDIUM	artf287737 : Update the default/trial license to include DevOps (and perhaps new license types)	Gowthaman Ravindran	Mavericks - CRUS	Mark Phippard
3 MEDIUM	artf287193 : SVN load balancing	akil roy	Devops/Releng	Johannes Passing
3 MEDIUM	artf286210 : Monitor Snapshot before restart	h karthik	Devops/Releng	Johannes Passing

At the bottom of the table, there are buttons for Monitor, Import, Export, Cut, Edit, Edit Inline, and Mass Update. The table is paginated with 10 items per page.

- You can type your filter criteria in the text boxes. The search text is case-insensitive.
- You can also select the filter values from one or more drop-down lists. By default, you can only select up to 10 filter values in a drop-down list. However, you can set a value that suits your requirement for the `FILTER_DROPDOWN_MAX_SELECTION` token in the `site-options.conf` file to increase or decrease the count.
- **Filter-as-you-type:** You can find the **Enter keywords** text box in all filter drop-down lists. As you type your filter keyword, instant search results are shown in the drop-down list. For example, in the following illustration, typing "R" instantly shows all statuses having the alphabet "R". The search text is case-insensitive.



The screenshot shows a search input field with the text "R". Below the input field, a list of filter options is displayed, each with a checkbox:

- Under Consideration
- Ready for UAT
- Ready for QA
- Rejected

- Some search filters may not appear if your site administrator has not enabled them.
4. After filtering, if you want to clear the filters, click **FILTER** and select **Clear** from the drop-down list.
  5. Use the up-down arrow at the top of any column to sort your list by that column.
    - Your primary sort column is identified by a superscript **1** next to the up-down arrow, and your secondary and third-level sort columns, if any, are likewise marked.
    - Click the up-down arrow again to reverse the sort order.


6. Click the task name to see the task details.

### Update task status

When you have completed some work on a task assigned to you, you must update its status.

If you don't mark a task as completed, you may get an automatic reminder from the task's owner, its administrator, or other project members.


You can provide task status reports for all of your assigned tasks from a single page.

 **Note:** If you want to enter a status update for only one task, you can do so from the Task Progress section of the **Task Details** page.

1. Click **Tasks** in the project navigation bar.
2. On the **Task Summary** page, click the **MY TASK PROGRESS** tab.  
Your **My Task Progress** page displays a list of all tasks assigned to you.
3. Enter your task status updates in the **ACCOMPLISHMENTS**, **PLANNED**, and **ISSUES** fields.

- **Accomplishments** - Record any accomplishments since your last update.
- **Planned** - Describe work you plan to do on the task.
- **Issues** - Report any issues with the task.

4. If applicable, change the task status.

 **Note:** Project administrators may require approval for changes to task status. If change approval is required, you are notified when you submit your status report.

5. Save your changes.
  - Click **Save Page** to finish and save your changes.
  - If your tasks cover more than one page, click **Save Page and View Next** to proceed to the next page.


### Schedule a task before or after another task

When one task cannot be done unless another task is done first, set up a dependency to track the relationship between the two tasks.


For example, if the deadline for a task changes, the tasks that depend on that task can automatically be rescheduled to reflect the change.

1. Click **TASKS** in the project navigation bar.
2. On the **Task Summary** page, find the task folder containing the task to which you want to add dependencies.
3. Click the task name.
4. On the **View Task Details** page, click the **DEPENDENCIES** tab.
5. On the **Task Dependencies** page, click **Add** in the **Predecessors** or **Successors** section.
6. In the **Add Task Dependency Wizard** window, choose a task.

- **ENTER TASK ID** - If you know the task ID, you can enter it directly.
- **ADD FROM HISTORY** - The History is a list of tasks that you have submitted or edited.

 **Note:** To create multiple dependencies, enter multiple task IDs separated by commas, or use the **ctrl** key and select multiple tasks from the history list.

7. Click **Next**.
8. On the **Add Task Dependency Wizard** that shows the dependency created, click **Finish**.

 **Note:** When a dependency is added to or removed from a tracker artifact, a notification mail is sent to users monitoring the artifact.

An option is provided at site level and user level to make sure whether the notification mail has to be sent or not. For more information on this, see [Configure your site's settings](#) on page 365.

The **Predecessor** or the **Successor** section lists out the dependencies you have added.

## View task dependencies

Task dependencies allow you to establish relationships between tasks. Task dependencies show what must be done before a given task and what will be done after the task is complete.

1. Click **TASKS** in the project navigation bar.
2. On the **Task Summary** page, find the task folder containing the task that you want to view.
3. Click the task name.
4. On the **View Task Details** page, click the **DEPENDENCIES** tab.

For each predecessor and successor task, you can see the person to whom the task is assigned, any issues with the task, the start and end dates, and the effort expended on the task. For details about a predecessor or successor task, including that task's predecessors and successors, click the task name.


## Associate a task with other items

When a task is related to other TeamForge items such as documents, tracker artifacts, file releases, integrated application objects, or news items, you can connect the document to the other item by creating an association.


Creating associations between items helps you define relationships, track dependencies, and enforce work flow rules. Some example uses for task associations include:

- Associating tasks with a master document describing the project, such as a project plan.
- Associating a task with the deliverables it requires, such as a requirements document, completed code, or a file release.
- Associating a task with a related object in the integrated application.
- Associating a task with items on which it is dependent, such as closed tracker artifacts.

1. Click **TASKS** in the project navigation bar.
2. On the **Task Summary** page, find the task folder containing the task to which you want to add an association, and click the task name.
3. On the **View Task Details** page, click the **ASSOCIATIONS** tab.
4. Click **Add**.
5. In the **Add Association Wizard** window, select the items with which you want to associate the artifact:
  - **ENTER ITEM ID** - If you know the item's ID, you can enter it directly.

 **Note:**

- To associate an object in an integrated application from within TeamForge, use the [`<prefix_objectid>`] format. Successful associations appear hyperlinked.
  - Each integrated application displays its prefix on moving the mouse over the application name in the tool bar.
  - **ADD FROM RECENTLY VIEWED** - Select one of the last ten items you looked at during this session.
  - **ADD FROM RECENTLY EDITED** - Select one of the last ten items you changed.
6. Click **Next**.
  7. You may add a comment in the **ASSOCIATION COMMENT** text box.
  8. Save your work.
    - Click **Finish and Add Another** to add additional associations.
    - Click **Finish** to return to the **Details** page.

 **Note:** When an association is added to or removed from TeamForge objects such as tracker artifacts, tasks, documents, discussions, and file releases, a notification mail is sent to users monitoring these objects.

An option is provided at site level and user level to make sure whether the notification mail has to be sent or not. For more information on this, see [Configure your site's settings](#) on page 365.

9. Click the **Associations** tab to view a graphical representation of all the associated items.

Through the Association Viewer, you can choose to view associations in the form of a list or flip over to the Trace view to explore the layers of associations (including parent/child dependencies) laid out in a timeline. You can scroll across the Trace view by dragging the mouse over the association layer or use the 'Previous' and 'Next' arrows to view all the objects as events in a timeline.

While the Associations tab shows the count of the total number of associations, you can only view the most recent 500 associations when you click the Associations tab in case the artifact has more than 500 associations. You can, however, browse through the Association Viewer to view older associations.

## Manage tasks

Project members with the Task Admin permission can perform all task administrative functions such as copying and moving tasks. They can also create and administer task folders.

### *Copy a task*

Copying a task is useful if you want to use an existing task as a template for creating a new task with many of the same attributes.

When making a copy of a task, reference to the source task is maintained in the task name and history and in the **Associations** section of the **Task Details** page.

1. Click **TASKS** in the project navigation bar.
2. On the **Task Summary** page, find the task folder containing the task that you want to copy.
3. Select the task that you want to copy, and click **Copy**.  
CollabNet TeamForge creates a copy of the the task and places it on the clipboard.
4. Go to the folder where you want to place the copy and click **Paste**.

You can place a copy in the same task folder or a different task folder.

The task copy is now placed in the selected folder, with the name `Copy of <Task Name>`.


### *Move a task*

When a task's focus or status changes, it may be useful to move it to a more appropriate task folder.


You can move a task to another folder in the same project or in another project.

To move a task between projects, you must have task administration permission in both the source and destination projects.

If the project members to whom the tasks are assigned have the appropriate permissions in the destination project, the tasks are automatically reassigned to the same project members. If not, you must reassign them.

 **Important:** If you move tasks across projects, the dependencies, resources, and change approval settings might be changed.

1. Click **TASKS** in the project navigation bar.
2. On the **Task Summary** page, go to the task folder containing the task you want to move.
3. Select the tasks you want to move and click **Cut**.
4. Go to the folder to which you want to move the tasks, and click **Paste**.

 **Note:** To move the tasks into another project, first find the destination project using your **Projects** menu.

- If the destination folder is in the same project, the tasks are now moved.
  - If the destination folder is in a different project, the **Paste Tasks** window appears.
5. If the project members to whom the tasks are assigned don't have the appropriate permissions in the destination project, reassign the tasks to another project member or to yourself.
  6. Click **Next**.

The tasks are now in the destination folder. The history of the move is recorded in the task history.

## Organize tasks

Use task folders to organize tasks according to work group or project phase.

### *Create a task folder*

Set up a task folder to contain tasks that are related to each other.



Task folders, like file directories, let you organize your tasks in a logical, hierarchical manner with labels that are meaningful to you and your organization.

You can create folders in the root folder or as subfolders in any other folder. There is no limit to the number of folders you can create.

1. Click **TASKS** in the project navigation bar.
2. On the **Task Summary** page, go to the location where you want to place the task folder.  
You can create a task folder in the root folder or in any other task folder.
3. Click **New**.
4. In the **New Task Folder** window, provide a name for the task folder.
5. Click **Save**.

The task folder is created.

#### *Rename a task folder*

When the focus or function of a group of tasks changes, you may wish to rename the folder they are in.


1. Click **TASKS** in the project navigation bar.
2. On the **Task Summary** page, find the task folder you want to rename.
3. Click the folder title and choose **Rename** from the **Edit** menu.
4. In the **Rename Folder** window, provide a new name for the folder.
5. Click **Save**.

The folder is now renamed.

#### *Reorder task folders*

To make it easier for project members to find things, you can arrange task folders alphabetically or otherwise.


By default, task folders are ordered according to when they were created. New task folders are added to the end of the task tree.

 **Note:** After you reorder folders, you cannot restore them automatically to their original order. You must reorder them manually.

1. Click **TASKS** in the project navigation bar.
2. On the **Task Summary** page, find the folder containing the subfolders that you want to reorder.  
You can reorder the task folders in the root directory, or the subfolders in any folder.
3. Choose **Reorder Subfolders** from the **Edit** menu.
4. In the **Reorder Subfolders** window, order the subfolders.
  - Click **Alphabetize** to sort in alphabetical order.
  - Click a folder and click **Move Up** or **Move Down** until the folder is where you want it.
5. Click **Save**.

#### *Delete a task folder*

If you no longer need a task folder, it's a good practice to delete it.

 **Important:** Deleting a task folder deletes all of the tasks within it. Delete a task folder only if you are sure that you no longer need any of the tasks within it. Move any tasks that you want to save before deleting the task folder.

1. Click **TASKS** in the project navigation bar.
2. On the **Task Summary** page, click the task folder that you want to delete.
3. Choose **Delete** option from the **Edit** menu.

The task folder is removed.

## Communicate with project members

You can communicate with project members via discussion forums, mailing lists, wiki pages, and project news announcements.

### Participate in discussions and mailing lists

CollabNet TeamForge discussions provide workspaces where project members can discuss project-related topics online or by email.


#### *Create a forum topic*

Create a new forum topic to begin discussion of a new subject.


A topic starts a message thread to which other users can reply. A forum can have any number of topics.

A forum topic is similar to an email, in that you can use it to communicate with other people subscribed to the forum, as if it were an email list. If the forum owner has enabled the forum to work as a mailing list, then you can post to the forum by email as well.

1. Click **DISCUSSIONS** from the **Project Home** menu.
2. On the **Forum Summary** page, click the title of the forum in which you want to create a topic.
3. On the **Topic Summary** page, click **Create**.
4. On the **Create a Topic** page, describe the topic in the **Subject** field.
5. Write the message in the **Message** field.  
After the topic is created, other users can reply to this message.
6. If you want the message sent by email to people who are not members of the forum, add their email addresses to the **Other Recipients** field. If there is more than one, put commas, semicolons or spaces between them.

 **Note:** In a moderated discussion forum, addresses in the `Other Recipients` field get your message only after the moderator approves the message.

7. To add an attachment to the topic, click **CHOOSE FILE** and select the file.
8. Click **Save**.


 **Note:** If this discussion forum is moderated, the topic is held until a moderator approves or rejects it. (Except if it is from a trusted user, these messages don't require moderation.)

#### *Search for posts*


Search for a post by using keywords, specifying the forum it belongs to, selecting the sender of the post or entering relative date or date range.

You can search for posts either across all forums or within specific forums.

1. Click **DISCUSSIONS** from the **Project Home** menu.
2. On the **Forum Summary** page, click **Search Posts**.
3. On the **Post Search Criteria** page, enter the desired search criteria.

 **Tip:** You can use wildcards.

- To search by subject, body or attachments, enter the text in the **Search Text** field and select **Subject**, **Body** and/or **Attachments** options.
- To search by forum name, select the forum in the **Forum** field.

 **Note:** When none of the forum options are selected, the system searches for matching posts across all forums.

- To search by post-sender, select the user name in the **Posted By** field.

 **Note:** You can select **User running search** option to display your posts.

- To search by time span, specify relative dates such as "Within the last 7 days".
- To search by date range, enter the start and end dates for the search. Click the Calendar icon to select dates from a calendar.


4. Click **Search**.

All posts matching your search criteria are displayed in the **Search Results** page.

#### *Reply to a forum message*


You can post a message in any topic in any forum you have access to. You can also post a message in reply to another message.

If the forum is moderated, you must have posting permission. Contact the forum moderator.

 **Tip:** If you are getting your forum messages delivered as email, you can reply to a post by email too.

1. Click **DISCUSSIONS** from the **Project Home** menu.
2. On the **Forum Summary** page, find the topic in which you want to post a message.
3. In the section containing the message to which you want to respond, do one of the following:
  - Click **Quote** to quote the original message in your response message.
  - Click **Reply** to omit the original message from your response message.
4. Write the message.
5. To add an attachment to the message, click **CHOOSE FILE** and select the file.
6. Click **Save**.

The forum message is posted. Other project members can reply to it using the same process.

 **Note:** A message to a moderated discussion forum is held until a moderator acts on it. (Except if it is from a trusted user, these messages don't require moderation.)

#### *Subscribe to a discussion or mailing list*

When you monitor a discussion forum, you are notified of contributions to the forum by email. Monitoring a forum is the same thing as subscribing to a mailing list.

A bell icon in the **Monitoring** column indicates that you are subscribed to a forum.

It looks like this:



- If you prefer to do it by email:
  - a) To subscribe in message by message type subscription, send an email to <Email address name in Mailing list>-<project name>-subscribe@<domain name>
  - b) To subscribe in digest type, send an email to <Email address name in Mailing list>-<project name>-digest-subscribe@<domain name>
  - c) You can also change the subscription type from message by message to digest and vice versa.
- If you prefer to do it through the web interface:
  - a) Click **Discussions** from the **Project Home** menu.
  - b) Click **Monitor**.
    - Set the notification frequency using "Monitoring Preference" as it suits you.

#### *Subscribe others to a discussion or mailing list*

You can add other users to a discussion.

If you are a forum administrator, you can also add users to the discussion as a group.

1. Click **DISCUSSIONS** from the **Project Home** menu.
2. Select the forums that you want to add users to.
3. In the **Monitor** list, click **Users Monitoring Selected**.
4. On the **USERS** tab, Click **Add**.
5. On the **Find a User** window, move the users you want into the **USERS TO ADD** column.



**Tip:** To add users as a group (assuming you are a forum administrator), do the same operation on the **USER GROUPS** tab.

### *Post to a forum by email*

If the forum also works as a mailing list, you can create a forum topic by sending an email message to the forum. You can also reply to a post by replying to the email.

If the forum is moderated, you must have posting permission. Contact the forum moderator.

The forum's email address, if it has one, appears on the **Topic Summary** page.



**Tip:** If a forum is not set up to work as a mailing list, you can still get posts by email when you monitor the forum. See [Monitor many items](#) on page 130.

1. Click **DISCUSSIONS** from the **Project Home** menu.
2. On the **Forum Summary** page, click the forum in which you want to create a topic.
3. On the **Topic Summary** page, find the email address in the **Mailing List** field. Send your email message to that address.

CollabNet TeamForge maps your email to the forum topic like this:

Email Field	Forum Topic field
To	Forum email address
Subject	Title of forum topic
Body	Text of forum topic
Attachments	Attachments

Emails from the forum can be read in RTF (Rich Text Format) or HTML format. The format of the message is delivered as an attachment to the post. Embedded attachments, such as text or images, are also delivered.

A message to a moderated discussion forum is held until a moderator acts on it. (Except if it is from a trusted user. These messages don't require moderation.)

In a moderated discussion forum, if you add other addresses in the `cc:` field of your email, those addresses get your email only after the moderator approves the message.

### *Associate forum messages with other items*

When a forum message concerns some other CollabNet TeamForge item, such as a document, a tracker artifact, a file release, a code commit, or a task, link the message to the item under discussion by creating an association.

Creating associations between items enables you to define relationships, track dependencies, and enforce work flow rules.


1. Click **DISCUSSIONS** from the **Project Home** menu.
2. On the **Forum Summary** page, find the forum message with which you want to create an association.  
Any existing associations are displayed in the **Associations** section.
3. Click **Associate**.
4. In the **Add Association Wizard** window, select the items with which you want to associate the artifact:
  - **ENTER ITEM ID** - If you know the item's ID, you can enter it directly.



#### **Note:**

- To associate an object in an integrated application from within TeamForge, use the [`<prefix_objectid>`] format. Successful associations appear hyperlinked.
- Each integrated application displays its prefix on moving the mouse over the application name in the tool bar.
- **ADD FROM RECENTLY VIEWED** - Select one of the last ten items you looked at during this session.
- **ADD FROM RECENTLY EDITED** - Select one of the last ten items you changed.

5. Click **Next**.
6. You may add a comment in the **ASSOCIATION COMMENT** text box.
7. Save your work.
  - Click **Finish and Add Another** to add additional associations.
  - Click **Finish** to return to the **Details** page.

 **Note:** When an association is added to or removed from TeamForge objects such as tracker artifacts, tasks, documents, discussions, and file releases, a notification mail is sent to users monitoring these objects.

An option is provided at site level and user level to make sure whether the notification mail has to be sent or not. For more information on this, see [Configure your site's settings](#) on page 365.

### *Administer forums and mailing lists*

Forum administrators create forums and do what is needed to keep them on track, such as editing or moderating forum posts.


You can choose to make a discussion forum either public or private.

You can be a forum or mailing list administrator without being a project administrator. Ask your project administrator or site administrator to grant you forum administration permissions.

Forum administrators can enable or disable forum moderation and add or remove moderators and trusted users. Any project member with forum post permissions can be a moderator.

Forum administrators can also make forums work like mailing lists.

Guest users can monitor a forum if email monitoring is set to "Allow all site users and guests". Additionally, guest users can email-post or subscribe to a forum if the mailing list is enabled and "Email Posting" is set to "Allow all site users and guests" from the discussion settings.

 **Tip:** Who can post by email to a discussion forum is controlled by the "Email Posting" options set for the forum. It does not depend on the permissions set for users of the Web forum.


### Create a discussion forum


Encourage project members to work together by creating discussion forums to which project members with the appropriate permissions can post messages.

Discussion forums can also function as mailing lists.

As a forum administrator, you can choose to make a discussion forum either public or private.

1. Click **DISCUSSIONS** from the **Project Home** menu.
2. On the **Forum Summary** page, click **Create**.
3. On the **Create Discussion Forum** page, enter a title and description for the forum.
4. To make it a private discussion forum, set the **TYPE** to Private.  
Private discussion forums restrict anyone without specific access permissions from posting to the forum. For example, you may want to restrict a preliminary planning discussion to your project's core team before sharing it more widely.
5. If you want the forum to work as a mailing list, select **Enable Mailing List**.
  - a) Provide a name for the mailing list in the **EMAIL ADDRESS** field.
 

 **Note:** The mailing list name must be unique within a project.
  - b) Set who can [post to the forum via emails](#). Choose either User with Roles and Permissions (default) or Allow only forum admins from the **Email Posting** drop-down list.
  - c) Set who can [subscribe to monitoring via emails](#). Choose either User with Roles and Permissions (default) or Allow only forum admins from the **Email Monitoring** drop-down list.

 **Note:** Selecting Allow only forum admins for **Email Monitoring** will not restrict users with 'Discussion-View' permission from getting monitoring emails in case they choose to *monitor the forum via the web UI*.

d) Choose how replies to posts are handled by setting the **REPLY BEHAVIOR**.

 **Tip:**

- Many users are accustomed to having their replies go automatically to the whole list. Others are used to having replies go just to the original sender. You should check with your users to see what makes more sense for a particular mailing list.
- When **REPLY BEHAVIOR** is set to "To the list," email replies are sent to the list as a whole, not to the individual post. This may be a change from what some users are used to.

e) Specify a prefix for the subject lines of messages from this list.

This can help users sort their incoming messages, if they are subscribed to multiple lists.

f) You may limit the size of emails (including attachments, if any). Enter the size (in MB) in the **MESSAGESIZE** field.


g) Under **FOOTER TEXT**, provide any information you want to show up at the bottom of each email that subscribers receive.

For example, you may want to offer useful web locations or email addresses.

6. To make this a moderated forum, select **ENABLE MODERATION**.

a) Click the Search icon to add moderators.

A moderated discussion must have at least one moderator.

 **Note:** If your project includes members of a parent project, you can select those members too.

b) Click the Search icon to add trusted users.

Posts by trusted users do not need moderator approval.

7. On the **Create Discussion Forum** page, click **Save**.


If you set your forum to work as a mailing list, all project members monitoring the forum will receive notifications whenever a new topic or message is posted.

Rename or edit a discussion forum

To help keep a forum or mailing list focused, try updating its title.

As a forum administrator, you can enable/disable mailing list or moderation features or just update their settings for the discussion forum.

1. Click **DISCUSSIONS** from the **Project Home** menu.
2. On the **Forum Summary** page, select the forum you want to change, and click **Edit**.
3. On the **Edit Discussion Forum** page, make your changes.
4. Click **Save**.

 **Note:** All project members monitoring the forum receive notifications of the update.


Add or modify moderators

As a forum administrator, you can add or remove forum moderators.

 **Note:**


- If a forum is moderated, it must have at least one moderator.
- When you designate a forum moderator, you also become a moderator yourself.

1. Click **DISCUSSIONS** from the **Project Home** menu.
2. On the **Forum Summary** page, select the forum for which you want to add or modify the moderators.
3. On the **Topic Summary** page, click **Edit**.
4. On the **Edit Discussion Forum** page, add or modify the forum moderators.

 **Note:** The existing moderators are listed.

- Click the **Search** icon to add or remove forum moderators.

You can select one or more moderators.

 **Note:** You can select the inherited project members also from the list.

- To add moderators, select the required users, click **Add** and click **OK**.
- To remove moderators, select the required users, click **Remove** and click **OK**.

- On the **Edit Discussion Forum** page, click **Save**.

#### Moderate a discussion forum by email

If your discussion forum is also a mailing list, you can approve or reject the post by email.

A message to a moderated discussion forum is held until a moderator acts on it. (Except if it is from a trusted user. These messages don't require moderation.)

As a moderator, you get an email when a message is awaiting moderation. The email contains the URL where you can approve or reject the message.

See the notification email for your options.


Option	Description
<b>To accept the message:</b>	Send an email to <postId>-accept@<domain>, or just click <b>Reply</b> .
<b>To accept the message and add the sender to the "Trusted Users" list:</b>	Send an email to <postId>-allow@<domain>, or click <b>Reply-to-All</b> .
<b>To reject the message:</b>	Send an email to <postId>-reject@<domain>.
<b>To add your comments to the message:</b>	Include your comments in the <b>Start Comment</b> and <b>End Comment</b> blocks in your response email. Your comments will appear in the approved post.


#### Approve a forum post

If a post is appropriate, you can add it to the forum by approving it.


A message to a moderated discussion forum is held until a moderator acts on it.

As a moderator, you get an email when a message is awaiting moderation. The email contains the URL where you can approve or reject the message.

 **Note:** If your discussion forum is also a mailing list, you can approve or reject the post by email. See the options in the notification email.


 **Note:** If the sender consistently contributes useful and appropriate input, you can save the moderation time by designating that user as a "trusted" user. Posts from trusted users don't have to be moderated.

- On **My Page**, click the **ITEMS AWAITING MY APPROVAL** tab.  
The number of posts awaiting approval are displayed against the project names.
- To view forum details, click the hyperlinked **Number of Posts** for your project.
- Click the **Number of posts awaiting approval** link on the **Forum Summary** page.

 **Note:** The **Forum Summary** page displays all the forum names and the corresponding number of posts awaiting approval, if any.

- Select either of the three approval techniques in the **Posts Awaiting Approval** tab.

As the moderator of the post, you will be able to view the topic title in the **All Topics** tab; and post title in the **Posts Awaiting Approval** tab on the **Topic Summary** page.

 **Note:** In the **All Topics** tab, the hourglass icons differentiate the topics that contain posts awaiting approval. To view posts nested within a topic, you can use the hyperlinked topic title.


Option	Description
<b>To approve posts and senders individually:</b>	Select the post and click <b>Approve</b> or <b>Approve and Trust</b> .
<b>To approve multiple posts and senders at once:</b>	Select all the posts to be approved and click <b>Approve</b> or <b>Approve and Trust</b> below the post details.
<b>To view the post details and approve:</b>	Click the post title and click <b>Approve</b> or <b>Approve and Trust</b> on the <b>Review Post Awaiting Approval</b> page after reading the details.


### Reject a forum message

If a proposed message is not appropriate or does not contribute to the goals of your discussion forum, you can reject it.


A message to a moderated discussion forum is held until a moderator acts on it. (Except if it is from a trusted user. These messages don't require moderation.)

As a moderator, you get an email when a message is awaiting moderation. The email contains the URL where you can approve or reject the message.

 **Note:** If your discussion forum is also a mailing list, you can approve or reject the post by email. See the options in the notification email.

 **Note:** You can reject posts with or without comments or reasons for rejection.

1. On **My Page**, go to **ITEMS AWAITING MY APPROVAL** tab.
2. On the **Forum Summary** page, click **Number of posts awaiting approval**. Topics that contain posts awaiting approval have an hourglass icon.
3. On the **Posts Awaiting Approval** tab, choose your method of rejection.


Option	Description
<b>To reject the posts individually:</b>	Select the each post, then click <b>Reject</b> at the end of the post.
<b>To bulk-reject the posts:</b>	1. Select all the posts you want to reject, then click <b>Reject</b> below the post details table.
<b>To view the post details and reject:</b>	Click the hyperlinked post title, then click <b>Reject</b> on the <b>Review Post Awaiting Approval</b> page.
<b>To explain your rejection with a comment:</b>	Click <b>Reject With Comment</b> instead of <b>Reject</b> .  <b>Note:</b> The rejection comment is posted to the message sender.

Rejected messages are deleted from the posts awaiting approval list and the message senders are notified by email.

### Delete a forum message

When a message in a forum is off topic or potentially harmful, you may want to delete it.

Before deleting a forum message, consider leaving it in place so that future users can consult it if they need to. Consider this even if the message does not seem very useful right now.

 **Important:** You cannot delete the topic starter's original message without deleting the entire topic.


1. Click **DISCUSSIONS** from the **Project Home** menu.
2. On the **Forum Summary** page, find the forum message that you want to delete.
3. Click **Delete** and confirm that you want to delete the forum message.

The forum message is deleted.



## Delete a forum topic

If you no longer want a forum topic in your project, you can delete it.

 **Important:** Deleting a forum topic deletes all of the forum messages in the topic. Delete a forum topic only if you are sure that you no longer need any of the forum messages in it.


1. Click **DISCUSSIONS** from the **Project Home** menu.
2. On the **Forum Summary** page, click the title of the forum containing the topic that you want to delete.
3. On the **Topic Summary** page, select the topic you want to delete.
4. Click **Delete** and confirm that you want to delete the topic.


The forum topic is deleted.

## Stop moderating a forum


If a moderated discussion forum does not require posts to be moderated anymore, moderation can be disabled.

To change a moderated discussion forum to unmoderated discussion forum in CollabNet TeamForge, select the forum and turn off its moderation feature.

 **Note:** Users with forum admin permissions only can enable/disable moderation.


 **Note:** On disabling moderation on a moderated forum, the posts awaiting approval are automatically approved.

1. Click **DISCUSSIONS** from the **Project Home** menu.
2. On the **Forum Summary** page, select the moderated forum that does not require moderation anymore.
3. On the **Topic Summary** page, click **Edit**.
4. On the **Edit Discussion Forum** page, to disable moderation, de-select the **ENABLE MODERATION** check box.

 **Note:** The following message is displayed: "Any post awaiting moderation will be approved automatically."

5. Click **Save** to turn off the moderation.

To keep moderation enabled, select **Cancel**.

 **Note:** The moderated discussion forum changes to an unmoderated discussion forum and any posts sent to the forum will be displayed without any restrictions.

## Share project news

Regular project news announcements help members stay in touch with events that can affect their work on the project.

### *Post a news item*

To help project members stay up to date, post regular news items on the project home page.

News items are posted and displayed on the project home page and on the CollabNet TeamForge home page.

1. Go to the home page of the project in which you want to post the news item.
  - From within the project, click **Project Home** in the project navigation bar.
  - From anywhere in TeamForge, choose the project from the **Projects** menu in the TeamForge navigation bar.
2. In the **Project News** section, click **CREATE NEWS POST**.
3. On the **Create News Post** page, provide a title for the news item.
4. Write the news item in the **BODY**.

The news item can be up to 4000 characters long, including spaces.

5. Click **Create**.

The news item shows up on the project home page and the TeamForge home page immediately.

### *Delete a news item*

It's a good idea to promptly delete a news item that is out of date or otherwise incorrect.

Deleting a news item from a project also deletes it from the TeamForge home page.

1. Find the project home page.

- From within the project, click **Project Home** in the project navigation bar.
  - From anywhere in TeamForge navigation bar.
2. In the **Project News** section, click **Delete** next to the news item that you want to delete, and confirm that you want to delete the news item.

The news item is deleted.

#### *Edit a news item*

To keep the project news in sync with the developments, modify the news items as and when required.

1. Find the project home page where you want to update the latest news.
  - From within the project, click **Project Home** in the project navigation bar.
  - From anywhere in TeamForge , choose the project from the **Projects** menu in the TeamForge navigation bar.
2. In the **Project News** section, click **Edit** next to the news item that you want to modify.
3. On the **Edit News Post** page, modify the title and/or the content of the news item.
4. Click **Save**.

The news post appears on the project home page, along with your name and the time at which you modify the post.

#### **Contribute to the project wiki**

The TeamForge Wiki allows you to create an unlimited number of Wiki pages in each TeamForge project. Wiki pages are a tool for managing project information as unstructured, linkable content.

#### *Start a Wiki*

To start communicating with other project members via Wiki, create a new page in your project's Wiki.

Every TeamForge project starts with a blank Wiki. You do not need to create a Wiki before you can begin adding content. After a Wiki is started, any user with the appropriate permissions can add or edit content; however, you cannot delete all of the content to start over with a new, blank Wiki.

1. Click **WIKI** from the **Project Home** page.
2. On the Wiki home page, click **Edit**.
3. On the **Edit Wiki Page**, write your Wiki text.

Wiki content is a combination of plain text, markup for font elements such as bold or italics, headers, bulleted and numbered lists, and links.


4. Customize your Wiki entry with any of these optional steps:
  - a) Use the buttons at the top of the text area to add Wiki markup to your text in the **WYSIWYG Editor** mode. You can also enter Wiki syntax directly into the text area in the **Plain Editor** mode. For an explanation of Wiki syntax, click [Syntax Reference](#).
  - b) To insert a link to another TeamForge item, just type the item id. You do not need additional Wiki syntax to create the link.
  - c) To change the size of the display window, drag the arrow available at the bottom-right of the window.
  - d) To attach an external file to a wiki page, click **CHOOSE FILE**, then browse for the desired file.
  - e) To add a version comment, write it in the **Version Comment** field.
5. Click **Preview Changes** to see how your Wiki content will look. You can make further edits from the Previewing Home page before saving your changes.
6. Click **Update** to save your changes.

Your Wiki is now started.


#### *Add Wiki content*

When the information you want to share with other project members does not fit neatly into a tracker comment or a document review, use a Wiki page for a more free-form communication flow.

After a Wiki is created, any user with the Wiki create, edit, or view permission can add or edit Wiki content. You can also add associations or multiple attachments, or create additional Wiki pages.

 **Note:** You can now add multiple attachments in a Wiki page without updating the Wiki each time to include another attachment. If attachments are not required they can be deleted before updating the Wiki page.

1. Click **Wiki** from the **Project Home** menu.  
Any existing Wiki content appears.
2. For more information about this Wiki, click **View Details**.
3. On the Wiki home page, click **Edit**.
4. On the **Edit Wiki** page, make your changes or additions to the Wiki content.
  - If you prefer to use buttons to do your text formatting, the way you would with a word processor, click [NAME OF BUTTON].
  - If your tastes run more to typing in your wiki formatting, see [Wiki syntax](#) on page 377 for the choices available.

 **Tip:** You can use a variety of preconfigured queries to generate up-to-date content for your Wiki page.

For more preconfigured Wiki content, see [Wiki syntax](#) on page 377.

5. In the **Version Comment** box, note the reason for your change.  
It is optional, but advisable, to get in the habit of recording a version comment. If your project manager has made it mandatory, then you must record a version comment before you can save your changes.
6. Click **Preview Changes** to see how your Wiki content will look.  
You can make further edits on the **Previewing** page before saving your changes.
7. Click **Update** to save your changes.

### *Create a new Wiki page*


A TeamForge site can have any number of Wiki pages. All Wiki pages are linked, and their relationships are traced on the **Back Links** tab of each Wiki page.

1. Click **WIKI** from the **Project Home** page.
2. In the large text field, insert the title of your new page between square brackets, like this:


```
Please post comments about [test results] here.
```

Then click **Update**.

The text between the square brackets becomes a link on your Wiki page.

 **Note:** You can also type the new link in CamelCase (each word starts with an uppercase letter, no spaces) and skip the square brackets.

3. Click the link to open the new Wiki page.  
The new page is created. The title of the new page is the same text as the link, rendered in CamelCase.
4. On the **Create Wiki** page, write the content you need, then click **Save**.

 **Tip:** You can use a variety of preconfigured queries to generate up-to-date content for your Wiki page.

For more preconfigured Wiki content, see [Wiki syntax](#) on page 377.

The new Wiki page is created. The back link to the page from which it was created appears on the **Back Links** tab of the **Show/hide Details** section.

### *Search a wiki*

Use the wiki search page to find content in a project wiki.

1. Click **WIKI** from the **Project Home** page.
2. On the **Wiki home** page, click **Search Wiki Pages**.
3. Try one of the predefined searches.  
These can save you time by running some of the most widely used content searches with a single click:
  - List the wiki pages that have changed in the last 15 days.
  - List the wiki pages that no other wiki page links to.

- List all the pages in this project's wiki.
4. If you need to narrow your search beyond the predefined searches, enter some search terms under **Wiki Pages Search Criteria**.

Wildcards are allowed.

- To search active wiki content, enter the search keywords in the in the **SEARCH TEXT** field.
- To search active and inactive wiki page versions, select **Search All Versions**. (By default, searches are performed on active wiki page versions only.)
- To search wiki attachments, select **Include Attachments**.
- If you know approximately when the wiki content was created or last edited, enter the start and end dates for the search and click **Search**. Click the calendar icon to select dates from a calendar.
- To search by author, click the Search icon in the **CREATED OR EDITED BY** field that displays a list of project members.

5. Click **Search**.

A list of wiki pages matching your search criteria appears.

*View a wiki page as HTML*

One way to resolve messy or incorrect formatting on a wiki page is by converting the page to clean HTML.

1. Click **WIKI** from the **Project Home** page.
2. Find your wiki page by navigating or searching.
3. Click the **View HTML** button.

The wiki uses HTML Tidy to display the cleanest HTML it can manage. View the page source to see the results.

*View a wiki page as a PDF*

When you want to send a wiki page to someone outside the project, it can be handy to convert it to a PDF document.

1. Click **WIKI** from the **Project Home** page.
2. Find your wiki page by navigating or searching.
3. Click the **View PDF** button.

Depending on how your browser is set up, the resulting PDF document appears or the browser offers to download it for you.


## Manage source code


Browse a project's source code repositories for detailed information about code commits, changed files, and associations with other CollabNet TeamForge items.

### Get the code

Browse TeamForge to find the code you want to work on, then check out the code.

You can view the contents of each file in a repository, plus additional information about each file such as revision history, comments, date and time of submission, and branch and tag information. You can also view differences (diffs) between any two files. You can also search for code in a repository using TeamForge Code Search.

 **Restriction:** You can see only those paths in the repository that the repository administrator has given you access to.

 **Note:** If you're getting code from a Subversion replica repository, the TeamForge account used while setting up the replica determines what's available to be checked out. This account could have been provided total access to the master repository, or restricted access using path-based permissions.

1. Click **SOURCE CODE** from the **Project Home** menu.
2. On the list of project repositories, click the name of the repository in which you want to view code. For each file, the revision number, time since check-in, author, and last log entry appear in the **Repository Browser**.
  - To view a file, or to view the diffs between two files, click the file name.
  - To view a specific version of the file, click **Download**.


- To view the differences between two files, do either of these:
    - Click **[select for diffs]** next to each of the two files that you want to compare.
    - Enter the file revision numbers in the Diffs between boxes at the bottom of the page.
- If you need to diff files, choose a display from the **Type of Diff** menu, then click **Get Diffs**. The differences between the two files are displayed.
  - Use your source control client to check out the code to your local machine.

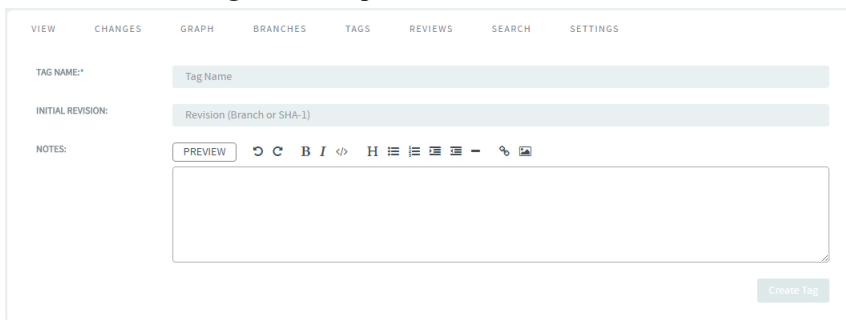
### Internal code browser

For Subversion and Git repositories, you have the option to use the TeamForge code browser which is turned on by default while integrating the source code server. For more information, see [Integrate a source code server](#).

- On the list of project repositories, click the name of a Subversion or a Git repository in which you want to view code.

On the top right of code browser, you can select the branch/tag (for Git) or specify the revision (for SVN) you want to browse.

- View:** This tab allows you to do the following:
  - Browse through the folder hierarchy of the repository and view the content of specific files. For any folder or file you are viewing within a branch (Git) or revision (SVN), you can obtain the commit information pertaining to its last update.
  - While viewing a single specific commit or a file, you can see the paths that were modified in that commit, the associations including JIRA such as builds, code reviews and so on (from EventQ events) for the specific commit and the differences between files in that commit.
    -  **Important:** To view the associations, you must have installed EventQ and must have RBAC (role-based access control) permissions to use the "EventQ READ" or "Reporting API" of TeamForge EventQ. If either of these requirements is not met, this section will not show up at all.
  - While viewing a folder, if there is a file named `readme`, `readme.txt` or `readme.md` that file will automatically be rendered beneath the list of files in the folder. If the file contains markdown formatting, it will be rendered as rich text.
- Changes:** This tab lets you view all of the commits that touched a specific path you are browsing within a branch or revision. Click a commit to view its details.
- Graph:** This tab provides a graphical representation of the changes made including branching and merging of repositories.
- Branches** (for Git): This allows you to see all of the branches in the repository in their relation to the default (master) one. Using **Compare Branch** you can see the commits in the branch that do not exist in the default branch.
- Tags:** This tab lets you create Git tags and tag specific points in history as being important. Typically you can use this functionality to mark release points (v1.0, and so on) with an option to add Release Notes for the tagged revision. Once you create a tag, you can use it to download source code as a zip/tar file and view the tag information in **Changes** and **Graph** tabs.



The screenshot displays the 'TAGS' tab in the TeamForge code browser. At the top, a navigation bar includes tabs for VIEW, CHANGES, GRAPH, BRANCHES, TAGS, REVIEWS, SEARCH, and SETTINGS. The main content area contains three input fields: 'TAG NAME:' with a text box containing 'Tag Name', 'INITIAL REVISION:' with a text box containing 'Revision (Branch or SHA-1)', and 'NOTES:' with a rich text editor containing 'PREVIEW' and various formatting icons. A 'Create Tag' button is positioned at the bottom right of the form.

The image displays three screenshots of the TeamForge interface, illustrating the 'TAGS' feature. The first screenshot shows the 'TAGS' tab with a list of tags, including 'build/17.1.522', which is highlighted with a red box. The second screenshot shows the 'COMMITS' tab for the date 01/27/2017, with a commit message 'b92f308: Add validation to moveDocument3() SOAP API' and a tag 'build/17.1.522' highlighted with a red box. The third screenshot shows the 'GRAPH' tab with a commit history, where the tag 'build/17.1.522' is highlighted with a red box.

- **Reviews:** This tab lists all the Open, Merged and Abandoned reviews, both Pull Requests and Gerrit single-commit reviews. Pull requests allow developers to collaborate with each other on a code change before merging it into another branch on a GIT repository. You can access this tab only when the repository owner has enabled this feature. For more information, see [Pull request step-by-step](#) on page 100.
- **Support for both Pull Requests and single-commit Gerrit Reviews:** Supports all types of code review policies, which include Pull Requests and single commit Gerrit Reviews.

VIEW CHANGES GRAPH BRANCHES TAGS **REVIEWS** SEARCH SETTINGS

Open Merged Abandoned Includes both Gerrit and Pull Request reviews

<b>p</b>	#5061 - [artf255076] updated title :)	🔗 master	M Miikka Ander
<b>p</b>	#5050 - [artf260665] Add test cases for server controller	🔗 master	Stephen Else
<b>b</b>	#5057 - [artf271495] testing testing	🔗 master	M Miikka Ander
<b>p</b>	#5055 - Use rest api to delete associations	🔗 master	B Baybora Aksc
<b>p</b>	#5042 - [artf262520] : Deprecate *_REPOSITORY_BASE and move default rep...	🔗 master	V vetrivel vaithi
<b>p</b>	#5025 - [artf275684] : Basic changes to list comments [artf255243] : Paginatio...	🔗 master	R Rakesh Kuma
<b>p</b>	#5012 - [artf274066] : backend changes for Wiki component	🔗 master	A Aishwarya Ra
<b>p</b>	#5043 - [WIP] [artf260017] Upgrade to Wildfly 10.0	🔗 master (wildfly10-upgrade)	J Johannes Pa

Gerrit code review

Pull Request review

**Auto refresh when a Pull Request changes:** When a pull request changes, the page is automatically refreshed to reflect the changes.

- **Comments to support @mentions:** Inline comments are parsed for @mentions and users called out via @mentions are added as reviewers.
- **Open your Gerrit Reviews in Gerrit's user interface:** A new button has been added to let you open your Gerrit Reviews in Gerrit's user interface.

VIEW CHANGES GRAPH BR


**K** #4456 - [artf263918]- Package Compon

OPEN 🔗 refs/changes/56/4456/1 → 🔗 m


**↗** Messages 3 Files 5

View this code review in Gerrit

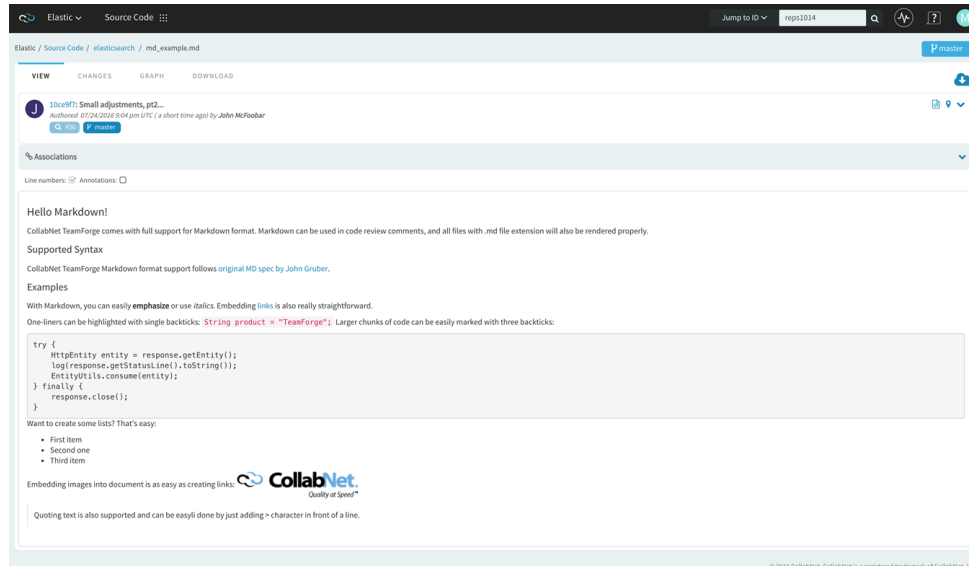
- **Code commenting:** During code reviews, you can now add line comments in context while looking at the files in diff view. You can double-click to block a line/text and add a comment.



You can also reply to line comments.



- **Ability to diff the change against the Base or a previous Patch Set:** As part of the Gerrit review workflow, you now have the ability to diff the change against the Base or a previous Patch Set.
- **Markdown support:** Markdown support for all .MD files: Render Markdown files when viewed through Code Browser



- **Mass delete/resurrect options:** Mass delete/resurrect options in History Protect tab:



Elastic / Source Code / elasticsearch

VIEW CHANGES GRAPH **BRANCHES** REVIEWS SETTINGS

Default branch: master

93cc3e4: testing pr  
Authored 06/13/2016 8:22 pm UTC (1 month ago) by John McFoobar

Active Merged All Deleted/Rewritten

	TYPE	REFERENCE NAME	USER	REWRITE DATE	ORIGINAL COMMIT ID	NEW COMMIT ID
<input checked="" type="checkbox"/>		build_update	john	03/07/2016 4:45 pm UTC	c310cc3	
<input type="checkbox"/>		conflict_branch	mike	03/07/2016 8:40 pm UTC	265edc5	
<input checked="" type="checkbox"/>		license_update	john	03/07/2016 7:11 pm UTC	2e42c1c	
<input type="checkbox"/>		master	admin	03/07/2016 2:13 pm UTC	5194590	050bfc3

Resurrect Delete Permanently

© 2016 CollabNet. CollabNet is a registered trademark of CollabNet, Inc.

Elastic / Source Code / elasticsearch

VIEW CHANGES GRAPH **BRANCHES** REVIEWS

Default branch: master

93cc3e4: testing pr  
Authored 06/13/2016 8:22 pm UTC (1 month ago) by John McFoobar

Active Merged All Deleted/Rewritten

**Resurrect A Rewritten Reference**

Do you want to resurrect 2 selected references to their original name?

Cancel Resurrect

	TYPE	REFERENCE NAME	USER	REWRITE DATE	ORIGINAL COMMIT ID	NEW COMMIT ID
<input checked="" type="checkbox"/>		build_update	john	03/07/2016 4:45 pm UTC	c310cc3	
<input type="checkbox"/>		conflict_branch	mike	03/07/2016 8:40 pm UTC	265edc5	
<input checked="" type="checkbox"/>		license_update	john	03/07/2016 7:11 pm UTC	2e42c1c	
<input type="checkbox"/>		master	admin	03/07/2016 2:13 pm UTC	5194590	050bfc3

Resurrect Delete Permanently

© 2016 CollabNet. CollabNet is a registered trademark of CollabNet, Inc.

- **Inline editing of files:** Quick changes to files, if required only to few files, can be done using the inline edit feature from within the code browser without having to clone an entire repository. Browse the repository, locate and open the file in the **View** tab, click **Edit** to open the file in the **File Editor**, make your changes, **Create code review** and **Publish** your changes for review.

The screenshot shows the TeamForge interface. At the top, there are navigation tabs: VIEW, CHANGES, GRAPH, BRANCHES, TAGS, REVIEWS, SEARCH, SETTINGS, and DOWNLOAD. Below the tabs, a commit is displayed with the message "4266727: [artf285315] Updated documentation to reflect HAProxy changes". The commit is authored by Johannes Passing. Below the commit, there are buttons for "#8236" and "build/17.1.514". Underneath, there is an "Associations" section. At the bottom of the commit view, there is a "Show diff against:" dropdown menu with "No commit selected" and an "Edit" button. A red box highlights the "Edit" button. Below the commit view, there is a "File Editor" window for the file "README.md". The editor shows the following content:

```

11 - ### Cloning the repository
12
13 Go to [the repository viewer of the TF Git repo]
    (https://forge.collab.net/ctf/code/git/projects.teamforge/scm.TeamForge_C
  
```

At the bottom of the file editor, there are three buttons: "Preview", "Cancel", and "Create code review". The "Create code review" button is highlighted with a red box. A red line connects the "Edit" button from the commit view to the "Create code review" button in the file editor.

- **Submit whole topic:** You can now bundle related changes (code reviews) by topic and submit the whole topic for review instead of just submitting changes one-by-one. Just open a review, click the **Set Topic** link and enter the topic name.

The screenshot shows the TeamForge interface for a code review. At the top, there are navigation tabs: VIEW, CHANGES, GRAPH, BRANCHES, TAGS, REVIEWS, and SEARCH. Below the tabs, a code review is displayed with the message "#7969 - [artf284935] Enabled FindBugs for 'core' module". The review is in an "OPEN" state. Below the message, there is a link for "Set topic" with "(No topic)" next to it. A red box highlights the "Set topic" link. Below the link, there is an input field for the topic name. A red box highlights the input field. A red line connects the "Set topic" link to the input field. Below the input field, there is a "VIEW" button and a "REFS" section showing "refs/changes/69/7969/10" and "master". Below the "VIEW" button, there is a "SEARCH" button. Below the "SEARCH" button, there is a "REFS" section showing "refs/changes/69/7969/10" and "master". Below the "REFS" section, there is a "CHANGES" section showing "Enabled plugin".

- **Search:** This tab lets you search for code via TeamForge Code Search powered by *Elasticsearch*. You can search all files in a repository or narrow your scope to specific file types such as C, C++, C# and so on. Type your search keyword, select a file extension (optional) and click **Search**. For more information, see [Search code](#) on page 107.


The screenshot shows the TeamForge interface for the search tab. At the top, there are navigation tabs: VIEW, CHANGES, GRAPH, BRANCHES, TAGS, REVIEWS, SEARCH, and SETTINGS. Below the tabs, there is a search input field with the placeholder text "Search this repository". To the right of the input field, there is a dropdown menu for "Extensions..." and a "Search" button. A red box highlights the "Search" button. Below the search input field, there is a "Last indexed revision" section showing "9cbfcbf 04/18/2016 3:27 am PDT (9 months ago)".

- **Settings:** This tab lets you configure the repository settings.
- **Download:** This tab allows you to download a copy of the required file.

## View code commits

To stay up to date with code development on a project, browse the code commits made to each repository integrated into your TeamForge site.

For each code commit, you can view the list of files that were checked in, the version history of each file, and any associations with other TeamForge items, such as tracker artifacts or tasks.

 **Restriction:** You can see only those paths in the repository that the repository administrator has given you access to.

1. Click **SOURCE CODE** from the **Project Home** menu.
2. From the list of project repositories, select the repository you want to look at, then click **View Commits**.
3. If you have internal code browser disabled (see [Integrate a source code server](#) on page 347):

The **Commits** section of the **Repository Details** page lists all the code commits in the repository. By default, it shows the commits made over the preceding seven days.

- a) Specify the filter criteria in **Commit Name**, **Committed By** or **Committed On** date ranges and click **FILTER**.
- b) After filtering, if you want to clear the filters, click **FILTER** and select **Clear** from the drop-down list.
- c) To view the details of a commit, click its title.

The **Files** section of the **Commit Details** page lists all files that were checked in with the code commit, including the version number of each file and the last operation that was performed, such as modified, deleted, moved, copied, or added.

- To view the file information, click the file name.
- To view the latest version of the file, click the file version.

- d) To look at other items related to this commit, click the **ASSOCIATIONS** tab.

4. If you have the internal code browser enabled (see [Integrate a source code server](#) on page 347):

The **Changes** tab lists all the commits in the repository sorted by date.


- Use the **Expand all** toggle button to expand or hide commit log messages for all the commits or selectively show or hide the commit log for the commit you are interested in.
- You can also **browse repository from a specific commit** you are interested in.

## Associate code commits with other items

Create associations between code commits and other CollabNet TeamForge items, such as tracker artifacts or documents, to help define relationships, track dependencies, and enforce work flow rules.

For example:

- Associate a code commit with the bugs, feature requests, or other tracker artifacts that the code addresses.
- Associate a code commit with the task requiring its completion.
- Associate a code commit with an object in an integrated application.
- Associate a code commit with a requirements document.

 **Note:** When you commit files to your source code repository, a source code commit notification mail is sent to users who are monitoring that source code repository.

An option is provided at site level and user level to make sure whether the notification mail has to be sent or not. For more information on this, see [Configure your site's settings](#) on page 365.


### *Associate code with other items while committing*

When you commit files to your source code repository, use the commit comment to quickly link your commit with one or more tracker artifacts or other TeamForge items.

Associations track the links between code and the bugs, feature requests, or other tracker artifacts that the code addresses. You can also associate code commits with other TeamForge items, such as tasks or documents.

A project administrator can make associations mandatory for all code commits. When this is made mandatory, the following additional rules pertaining to code commit can also be set:

- Code commits can be performed only for open artifacts.
- To perform a code commit, the committer must be the owner of the specific artifact.

 **Note:** Once you enforce the above rules, validations are strictly enforced for commits against tracker artifacts only. In case you commit against any other TeamForge object, for example a wiki or a document, mere existence of the object ID ensures successful commit and association and no validations are performed against the status of the object or who it is assigned to.


You can create tracker artifact associations from whatever interface you normally use to check code into your SCM repository. You do not have to log into TeamForge.


Use the same syntax for commits to CVS and Subversion repositories.

When making a code commit, add the associate command in the commit message like this:


[<item id>], such as the TeamForge tracker artifact ID or task ID.

- TeamForge item IDs are always letters followed by four or more numbers, such as `task1029` or `artf10011`.
- To associate a commit with multiple TeamForge items, separate the item IDs with commas.
- All associations are displayed in the **ASSOCIATIONS** tab of the **Commit Details** page.
- The **Comment** section lists the comments made with each commit.

 **Note:** To associate an object in an integrated application, use the [`<prefix_objectid>`] format. Each integrated application displays its prefix on moving the mouse over the application name.

 **Note:** When you commit files to your source code repository, a source code commit notification mail is sent to users who are monitoring that source code repository.

An option is provided at site level and user level to make sure whether the notification mail has to be sent or not. For more information on this, see [Configure your site's settings](#) on page 365.

 **Tip:** To remind yourself of the details of the association later, look in the **CHANGE LOG** tab of the associated **View Artifact** page.

#### *Create associations with code that is already committed*

At any time after a code commit is made, you can associate the code commit with other CollabNet TeamForge items, such as tasks, integrated application objects or documents.

1. Click **SOURCE CODE** from the **Project Home** menu.
2. On The list of project repositories, select the repository containing the code commit with which you want to create an association.
3. Click **View Commits**.
4. On the **Repository Details** page, click the name of the commit with which you want to create an association.
5. On the **Commit Details** page, click the **ASSOCIATIONS** tab.
6. On the list of existing associations, click **Add**.
7. In the **Add Association Wizard** window, select the items with which you want to associate the artifact:
  - **ENTER ITEM ID** - If you know the item's ID, you can enter it directly.

 **Note:**


- To associate an object in an integrated application from within TeamForge, use the [`<prefix_objectid>`] format. Successful associations appear hyperlinked.
- Each integrated application displays its prefix on moving the mouse over the application name in the tool bar.
- **ADD FROM RECENTLY VIEWED** - Select one of the last ten items you looked at during this session.
- **ADD FROM RECENTLY EDITED** - Select one of the last ten items you changed.

8. Click **Next**.

9. You may add a comment in the **ASSOCIATION COMMENT** text box.

10. Save your work.

- Click **Finish and Add Another** to add additional associations.
- Click **Finish** to return to the **Details** page.

 **Note:** When you commit files to your source code repository, a source code commit notification mail is sent to users who are monitoring that source code repository.

An option is provided at site level and user level to make sure whether the notification mail has to be sent or not. For more information on this, see [Configure your site's settings](#) on page 365.

**11.** Click the **Associations** tab to view a graphical representation of all the associated items.

Through the Association Viewer, you can choose to view associations in the form of a list or flip over to the Trace view to explore the layers of associations (including parent/child dependencies) laid out in a timeline. You can scroll across the Trace view by dragging the mouse over the association layer or use the 'Previous' and 'Next' arrows to view all the objects as events in a timeline.


While the Associations tab shows the count of the total number of associations, you can only view the most recent 500 associations when you click the Associations tab in case the artifact has more than 500 associations. You can, however, browse through the Association Viewer to view older associations.

### Create a source code repository


Each project can have one or more source code repositories.

Before you can create a source code repository, a site administrator must first add one or more SCM servers to the TeamForge environment.

1. Click **SOURCE CODE** from the **Project Home** menu.
2. In the list of the project repositories, click **Create Repository**.
3. Choose the server on which you want to create the repository.

 **Note:** The drop-down list shows all of the SCM servers that the TeamForge administrators have added to the TeamForge environment.

4. On the **Create Repository** page, enter the directory name for the repository.

 **Note:** For CVS repositories, the directory name is the name of the directory relative to the CVS server's repository root directory. A UNIX group by the same name is also created to enforce permissions.

5. Enter a name and description for the repository.

If you plan to use an SCM server that requires approval for new repositories, use the **DESCRIPTION** field to provide your reason for asking to create this repository.

6. If you are creating a Git repository:
  - a) Select a **Repository Category**.

This is where you define the code review policy for your repository.


- **Default - no review** (default): No review is required for pushing changes to the repository.
- **Pull request:** Users can create and push to feature branches but require pull requests to certain “protected branches”. See [Review code](#) on page 97 for more information.
- **Optional review:** Every change submitted to the repository can be pushed for code review or directly pushed to the repository bypassing review. This depends on the TeamForge user having the appropriate permissions — source code "Delete/View" or "Commit/View" permission for the former and source code "Admin" permission for the latter.
- **Mandatory review:** Every change pushed to the repository must pass through a review process before it can be pushed (merged) to the repository.
- **Custom:** Users with TeamForge source code "Admin" permission to directly fine tune permissions (access rights) in Gerrit's web interface. Such changes are not overridden by TeamForge.
- **User defined:** User defined code review policy.

For more information, see the following blog posts that talk about access right mappings between TeamForge and its Gerrit integration.

- [What Gitflow is, and how to best use it with TeamForge/Gerrit user defined repository policies.](#)

- [How to implement your own user-defined repo category with branch based permissions.](#)
- [Alternative approaches to implement branch based permissions \(without user-defined repo categories\) and when to use which approach.](#)

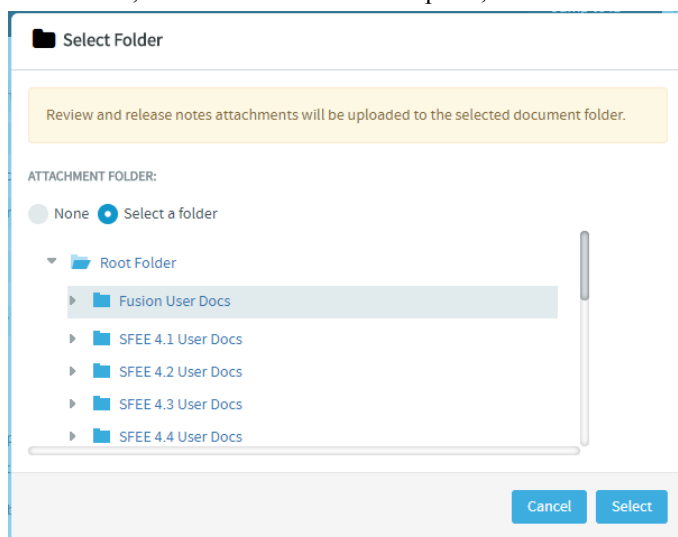
b) Enable history protection, if required. Select the **Protect History** check box.

 **Note:** The `GERRIT_FORCE_HISTORY_PROTECTION` site-options token is no longer supported. Remove this token from the `site-options.conf` file while upgrading to TeamForge17.4 or later. History Protection is enabled by default in TeamForge17.4 and later. However, Site Administrators can disable History Protection at the site level if need be, after which Project Administrators can choose to have History Protection enabled or disabled for individual repositories.

c) If you are enabling Git Large File Storage (LFS), select a value for **LFS Enabled** and **MAX LFS Object Size** fields. For more information, see [Set up LFS](#).


d) Select a folder to upload code review and release notes file attachments. All the files attached to the repository are stored in the selected folder.

Click **None**, select the Select a folder option, browse and select a folder from the folder tree and click **Select**.




7. If you want to require that each code commit be associated with an artifact (or a task or some other work item), select **Required on Commit**. Selecting this check box enables the following options:

- **Artifact must be in open state:** Select this check box to ensure that users cannot perform a commit for closed artifacts. Therefore, if a user attempts to commit changes for a closed artifact, an appropriate error message is displayed.

 **Note:** When the commit log includes multiple artifact ids, then all artifacts must be in the OPEN state for the user to perform a commit.


- **Committer must Own Artifact:** Select this check box to ensure that only an artifact owner can perform the relevant commit. If the committer is not the artifact owner, then an appropriate error message is displayed preventing them from committing the changes to the specific artifact.

 **Note:**

- In the case of a Git repository, the pusher who pushes the commits must be the owner of the artifact.
- Once you enforce the above rules, validations are strictly enforced for commits against tracker artifacts only. In case you commit against any other TeamForge object, for example a wiki or a document, mere existence of the object ID ensures successful commit and association; no validations are performed against the status of the object or who it is assigned to.

8. For security reasons, you may want to restrict email notifications to the essential information. If so, select **HIDE DETAILS IN MONITORING MESSAGES**.


9. By default, the **Available in Search Results** check box is selected. Clear the check box if you want to exclude the repository from searches.

 **Tip:** You have the ability to control whether source code is included in search results. For example, you may want to exclude source code from searches that have sensitive information. You may also want to omit some repositories, such as the branding and publishing repositories, from search results. The **Available in Search Results** setting in **Create Repository** or **Edit Repository** page allows you to configure this. By default, this option is selected for all repositories, except for the **branding** repository in the **look** project and for **Publishing** repositories.

#### 10. Click **Save**.

Your request for a new repository is submitted. You will receive an email notification when your repository is created or if your request for a new repository is denied.

- If the SCM server that you chose does not require approval for new repositories, the repository is created.
- If the SCM server that you chose requires approval for new repositories, a TeamForge administrator must approve your repository before it is created.

 **Note:** By default, repository requests on unmanaged CVS servers require approval, because the SCM administrator must manually create the repository.

#### Delete a source code repository


When you delete a repository, a request is submitted to the administrator for approval.

You need to have the required permission to delete SCM repositories.

1. Click **SOURCE CODE** from the **Project Home** menu.
2. In the list of the repositories, select the repository you want to delete and click **Delete**.  
The following confirmation message appears: All SCM data in this repository will be lost. Are you sure you want to delete this repository?
3. Click **OK** to delete.

Your request for deleting a repository is submitted. You will receive an email notification when your repository is deleted or if your request for deleting a repository is denied.

- If the SCM server that you chose does not require approval for deleting repositories, the repository is deleted right away.
- If the SCM server that you chose requires approval for deleting repositories, a CollabNet TeamForge administrator must approve your request to delete a repository before it is deleted.

 **Note:** By default, repository delete requests on unmanaged CVS servers require approval, because the SCM administrator must manually delete the repository.

#### Replicate a repository

When a Subversion Edge replica has been successfully registered with a TeamForge SCM integration server, it is available to project administrators in projects using that server to house repositories. To replicate a Subversion repository, you need to add it to a replica server.

Before you can replicate a Subversion repository, an administrator must first add one or more replica servers. This involves converting a *Subversion Edge server*, and then *approving the replica* in TeamForge.

1. Click **SOURCE CODE** from the **Project Home** menu.
2. In the list of project repositories, select the one you want to replicate and click **Edit**.

The **Edit Repository** page lists the available replica servers. Here's an example:

Replica servers hosting repository: Desktop				
Status	Name	Hostname/IP address	Description	Owner
No results found.				

Available Replica Servers				
Status	Name	Hostname/IP address	Description	Owner
<input checked="" type="checkbox"/>	Tokyo	cu126.cloud.sp.collab.net	Replica in Tokyo office	Subversion Replication

[Replicate repository](#)

If you don't see any available replica server listed here, it may be because none were created for this Subversion server, or there are pending replicas which haven't yet been approved by a TeamForge administrator.

3. In the Available Replica Servers section, select a replica server and click **Replicate Repository**.

You will see it in the list of replica servers hosting the repository.

- To see details such as the hostname and the user managing the replica repository, and recent replication command history, click the **Status** icon for the repository.
- For the Subversion checkout command, see the project's **Source Code** page.

**Note:** The TeamForge account used while setting up the replica determines what's available to be checked out from a replica repository. This account - the **Managed By** user in the **Repository Replica Details** page - could have been provided total access to the master repository, or restricted access using path-based permissions.

### Check command history

Recent command history for a replica server or a specific repository allows you to check for errors and see whether there are pending commands.

- Maybe a repository revision is not showing up. You can check for errors to know it's not just because the repository is still synching.
- You can also check if there are commands waiting in the queue -- you'd be able to see whether the repository is truly in synch.
- To check the command history for a replica repository, follow these steps:
  - a) Click **SOURCE CODE** from the **Project Home** menu.
  - b) In the list of repositories, click the icon for the one you're interested in. You'll see repository details and command history. Here's an example:

Repository Replica Details				
Replica Server:	Tokyo	Server:	Subversion	
Hostname:	cu126.cloud.sp.collab.net	Repository Name:	Benchmark	
Managed By:	Subversion Replication	Description:	SVN benchmarks	
Last Contact:	Wed Apr 06 12:23:32 PDT 2011			
<a href="#">Browse Repository Replica</a>				

Replication Command History Page 1 of 3 (45 Items)				
Status	Command	Revision	Created Date	Last Contact
Any			From 04/05/2011  (MM/dd/yyyy)	
			To 04/06/2011  (MM/dd/yyyy)	
	Sync repository	70	Wed Apr 06 12:21:50 PDT 2011	Wed Apr 06 12:21:54 PDT 2011
	Sync repository	69	Wed Apr 06 12:21:46 PDT 2011	Wed Apr 06 12:21:54 PDT 2011
	Sync repository	68	Wed Apr 06 12:21:41 PDT 2011	Wed Apr 06 12:21:43 PDT 2011

**Tip:** You can also check command history by clicking on the **Status** icon in the **Edit repository** page.



- You need to be a TeamForge administrator to check the command history across a replica server.
  - a) On the site administration navigation bar, click **INTEGRATIONS**.
  - b) On the **SCM INTEGRATIONS** page, click the name of the replica server you're interested in. The **Edit System** page displays the command history.

### Check out code

You can use the checkout command to check out the code from Subversion or GIT repository.

#### Check out Subversion code anonymously

When you want to experiment with the code, you can do an anonymous checkout from the Subversion repository. The checkout command uses a unique system-created user called "guest" and works without authentication.

To make anonymous checkout possible, the project administrator must set public and repository view permission to "All Users." The checkout command differs based on whether a user is logged in or not.

- When you are not logged in, use this command to check out: `svn checkout --username guest <domain>/svn/repo URL name`
- When you are logged in, use this command to check out: `svn checkout --username <logged in_username><domain>/svn/repo URL name`


#### Check out GIT code

Generally, GIT repositories can be accessed by more than one protocol. To support this, a **Protocol** drop-down list is available on the source code repositories page. This feature applies only to GIT repositories and so selecting a protocol determines the check out command for a GIT repository. This also gives the user an option to override the default protocol which is set while configuring the GIT integration server.

#### Access CVS with a stored SSH key

If you don't want to enter your password each time you access a CVS repository, you can create a set of SSH authorization keys and use the public-key method for automatic authentication.

Secure shell (SSH) can use public-key cryptography to confirm your identity, authenticate you to the remote host, and encrypt data transmission. On many TeamForge sites, access to CVS repositories is authenticated using SSH.

 **Note:** These instructions will help you only if all the following are true:

- You are using CVS.
- The `CVSROOT` (or "`cv s -d`") info given you by your admin (or the CTF "Repositories" list) does not include `:pserver:`
- Your admin has set up SSH access for you.

Check with your site administrator if you are not sure about any of these conditions.

1. Generate an SSH key pair according to the instructions in the documentation for your SSH client. (Each SSH client provides its own mechanism for key pair generation.)
2. Log into TeamForge .
3. On your **My Workspace** page, select **MY SETTINGS** from the **My Page** menu.
4. On your **User Details** page, click **AUTHORIZATION KEYS** tab.
5. On the **Authorization Keys** page, copy the public key from the `.pub` file in your SSH installation directory and paste it into the **Authorized Keys** field.
6. Click **Update**.

Your SSH public key is now saved. When you log into a CVS repository on your TeamForge site, SSH automatically checks your private key against this public key and authenticates you. You do not have to enter a password.

 **Note:** You can store multiple SSH public keys in TeamForge, but only one for each client machine.

### Review code

TeamForge provides a unified code review experience as it supports both Pull Request and Gerrit single-commit reviews.

Git repositories are hosted and served in TeamForge via Gerrit. Gerrit is most widely known for providing powerful code review features. While Gerrit includes a powerful code review feature, the way it works and the workflow is different from the Pull Request style that was introduced by GitHub. A Google search of “*Gerrit Pull Request*” will yield a bounty of passionate viewpoints on these differences. Hashing through the pros and cons would just add another result to that search. Instead, just know that with TeamForge you are free to use either style of code review methodology, even on the same Git repository.

VIEW CHANGES GRAPH BRANCHES TAGS REVIEWS SEARCH SETTINGS

Open Merged Abandoned Includes both Gerrit and Pull Request reviews

Review ID	Title	Target Branch	Reviewer
#5061	[artf255076] updated title :)	P master	Miikka Ander
#5050	[artf260665] Add test cases for server controller	P master	Stephen Else
#5057	[artf271495] testing testing	P master	Miikka Ander
#5055	Use rest api to delete associations	P master	Baybora Aksc
#5042	[artf262520] : Deprecate *_REPOSITORY_BASE and move default rep...	P master	vetrivel vaithi
#5025	[artf275684] : Basic changes to list comments [artf255243] : Paginatio...	P master	Rakesh Kuma
#5012	[artf274066] : backend changes for Wiki component	P master	Aishwarya Ra
#5049	[WIP] [artf260017] Upgrade to Wildfly 10.0	P master (wildfly10-upgrade)	Johannes Pa

Gerrit code review

Pull Request review

While TeamForge supports both pull request and single-commit Gerrit reviews, this topic focuses more on the Pull Request type reviews.

### Pull request configuration

In order to use pull requests in your Git repository there is some configuration that must be done first. This is to set up the proper permissions in your repository so that your policies are being followed.

Open the repository in the code browser, click the **Settings** tab then click **Policies** tab. This is where you configure the pull request-based code review policy. For more information, see [Configure Pull Request for repositories step-by-step](#).

**Repository Category and Protected Branches:** A new category named “Pull requests” has been added. What this category does is set up the repository permissions so that users can create and push to feature branches but require pull requests to certain “protected branches”. Once you change the repository category to “Pull request”, the **Protected Branches** field shows up. This will be the list of branches that you will be merging your pull requests into. Typically this would be the “master” branch but you may also have various “release” branches that you would like to protect. Users will not be able to directly push changes to these protected branches. Instead, the user will create a feature branch with their changes in it, and then create a pull request when they are ready for their changes to be reviewed and merged to the protected branch.

**Note:** “master” branch is added as a protected branch by default.

**Review Rules:** These review rules govern the requirements for a given change to be eligible to be merged. There are four new policies available:

- **No Approval Required:** This is similar to the policy on sites like Github. This basically means anyone with the proper TeamForge permissions can accept and merge any pull request. This means you probably favor “social” policies and trust that your reviewers will do the right thing. Pull requests become a tool to aid with code review

and it is still possible for users to use the voting tools in the review to communicate their feedback but the votes on a review do not prevent the review from being merged.

- **Code Review Required:** With this policy, the voting tools begin to matter. The change cannot be merged until it has a net positive vote total, not counting the owner of the review. In other words, if two users give a thumbs up and one a thumbs down then it will be eligible to be merged -- assuming the owner of the review is not one of the two thumbs up. The owner can vote, but their votes do not count towards the total.
- **CI Required:** The assumption here is that the relevant votes are being cast by a “bot” or “process” such as a Jenkins CI job. There is no UI in the pull request provided to cast these votes, it will be done via API or by using the Gerrit UI. The pull request shows a check mark when a positive Verified vote has been cast and an X when a negative vote has been cast. With this policy the change cannot be merged unless there is at least one positive verified vote and no negative votes. Users can still provide thumbs up and down votes but they do not control whether or not the change is eligible to be merged. Of course the person that decides to merge the change can still factor in the code review votes and comments.
- **Code Review and CI Required:** This is obviously just a combination of the two previous policies. So a positive Verified vote is necessary, with no negative verify votes, and a total positive Code Review vote is required.
- **Default:** The final policy is to just use the Gerrit code review default policy. This requires a +2 code review vote and a +1 verified vote and there cannot be a -2 code review vote as that acts as a veto. Users that are not familiar with Gerrit tend to find these voting rules confusing. For example, two +1 votes does not equal a +2 vote and your permissions determine what votes you are able to cast. We do not recommend you use this policy if you are using pull requests, but it is an available option and might be desired if you are already using Gerrit reviews and do not want to change the voting rules.

## Pull request workflow walk-through

The primary difference between the pull request workflow and the normal Gerrit change-based model is the use of branches. In the pull request model the assumption is that work will begin on a feature branch and you create a pull request when you are ready to start receiving feedback on the branch. This could mean the work is ready to be merged, but it could also mean that you just want to get feedback from the CI system or initial feedback from code review. Once all feedback and review is complete and the pull request is eligible to be merged, then the request can be merged and the feature branch deleted.

### Create feature branch (locally)

If working in a small team you might want to create the feature branch on the server or create one locally and push to the server right away. For now, we will assume that there is just a single developer. Typically, it is best to just begin the process by creating a feature branch locally. It is a good idea to fetch all changes from the server before beginning this process:

```
$ git checkout master && git pull origin master
$ git checkout -b feature_branch
```

Give your feature branch a meaningful name.

### Commit to feature branch and push to server

The next steps are of course to just do your work and commit changes locally. Before you have pushed the changes to the server it is OK to do things like squashing your commits or rebase your branch on master, but once you have pushed your branch to the server you should no longer do this. The first time you push to the server you will need to set the upstream branch to the name you want to give your feature branch on the server.

```
$ git push --set-upstream origin feature_branch
```

### Create pull request

When you are ready to merge the change, or at least to start getting feedback, you should create a pull request, add reviewers and have reviewers share feedback on your changes. See [Create Pull Request step-by-step](#) for more information.

Pull requests are implemented as merge commits between your feature branch and the target branch. The pull request subject and description will combine to form the commit message for the merge commit. You can also provide a Markdown summary of the change that will be captured as the first comment on the pull request. If no summary is provided, then the commit message will serve as the first comment. If you have automated CI configured, then it will typically run as soon as the pull request is created and whenever it is updated. So you could also create a pull request early in your process so that you can benefit from the feedback of your CI system. If you are posting a pull request that is not ready to be merged it is a good convention to follow to cast a Thumbs Down vote in the pull request to signify this to potential reviewers.

- **Reviewers:** Reviewers are anyone that you potentially want to provide feedback on the change. Adding a user does not require the user to review the change, it just notifies the user of its existence and see it in lists and filters of reviews they are assigned to. Likewise, users that have not been added to the review are still free to cast votes on the review.
- **Voting:** Tools are provided to cast votes on the review. If you are using one of the four code review policies provided with TeamForge you will see simple Thumbs Up/Down buttons to cast your vote.


### Commit and Push More Changes

As you continue to work on your feature branch you can just commit changes to your local branch and push them to the server. If you are working in a team on the same branch, then you will need to fetch and rebase changes made by other team members to the feature branch.

```
$ git add/commit etc.
$ git push
```

### Update Feature Branch and Pull Request

If new changes are pushed to the feature branch, the pull request must be updated to include the new changes. This involves updating the merge commit to recognize the new HEAD of the feature branch. To update the pull request you simply need to open it in the web UI. It will then update itself as needed.

 **Note:** When a pull request is updated, all existing votes will be reset and need to be cast again based on the new review.

### Resolve Conflicts

When working in a feature branch, it is not uncommon for conflicts to arise between your feature branch and the target. When this happens, you will not be able to merge the pull request and you will see a warning of the conflict in the web UI. To resolve the conflicts you must fetch and merge the changes from the target branch into your feature branch and resolve and commit the conflict resolution. Then push the result back to your feature branch on the server and update the pull request.

```
$ git fetch $ git merge origin/master
$ git add/commit etc. $ git push
```

Of course you can also rebase and force push to update your feature branch as long as you understand the ramifications of this when collaborating with a team that is sharing the same branch.

### Merge Pull Request


Once a pull request is eligible to be merged, meaning there are no merge conflicts with the target and all voting requirements have been satisfied, the Merge button will be enabled in the web UI. Anyone who can see this button has the permissions to merge the pull request and just needs to click the button to merge it. See [Merge pull requests step-by-step](#) for more information. This ends the life cycle for this pull request and if the user has the necessary permissions they will also be given the ability to delete the feature branch from the server.

It is possible to continue to use the feature branch and create new pull requests to merge subsequent changes, but this is not recommended. It is generally a good idea to delete feature branches once they have been merged.

#### *Pull request step-by-step*

Pull requests allow developers to collaborate with each other on a code change before merging it into another branch on a GIT repository.

Pull request is a fully integrated solution of the code browser component. It supports all the basic functionality such as creating, viewing, updating, abandoning, rebasing and merging of pull requests. Using a pull request, you notify others about a feature or fix change that needs attention.

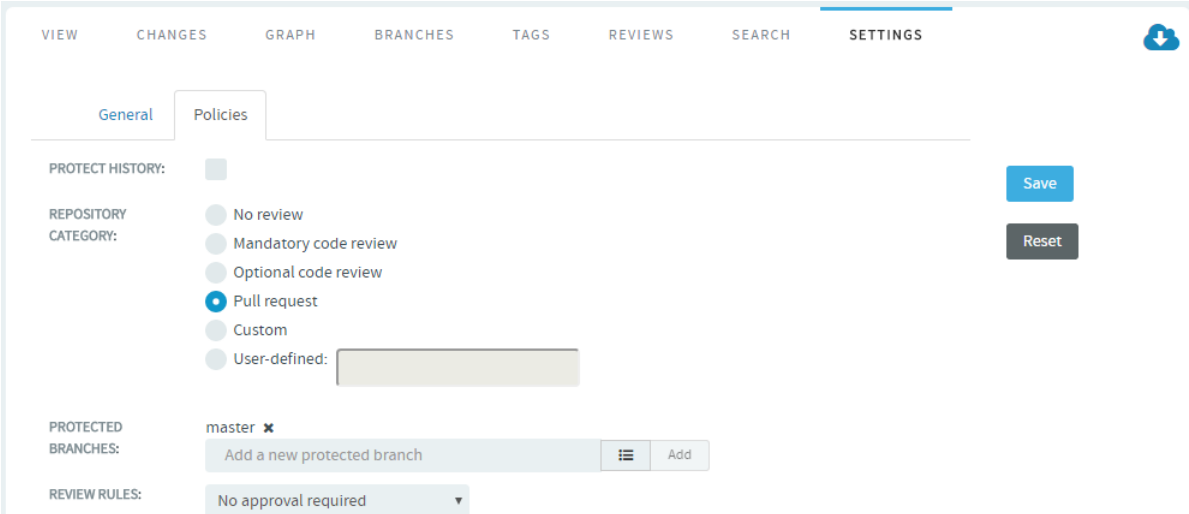
 **Important:** You can access the pull request feature only when the repository owner enables the feature and sets the code review policy on the **Settings > Policies** tab.

### Configure "Pull Request" for repositories

In order to use pull requests in your Git repository, you need to set up proper permissions in your repository so that your policies are being followed. Configure the repository from the **Settings** tab.

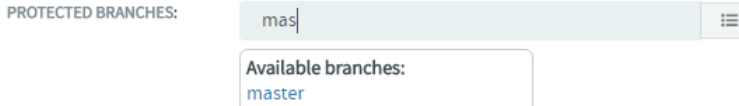
1. Click **SOURCE CODE** from the **Project Home** menu.
2. Browse and open the Git repository in the code browser.
3. Select **Settings > Policies**.
4. Select Pull request for **Repository Category**.

The **Protected Branches** and **Review Rules** fields show up. For more information about Repository Category, Protected Branches and Review Rules, see [Review code](#) on page 97.

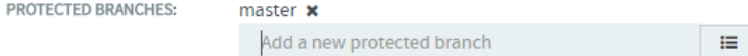


The screenshot shows the 'Settings > Policies' interface. The 'Policies' tab is selected. The 'REPOSITORY CATEGORY' section has radio buttons for 'No review', 'Mandatory code review', 'Optional code review', 'Pull request' (selected), 'Custom', and 'User-defined'. The 'PROTECTED BRANCHES' section shows 'master' as a protected branch with an 'Add' button. The 'REVIEW RULES' section has a dropdown menu set to 'No approval required'. 'Save' and 'Reset' buttons are on the right.

Add one or more protected branches. Type the branch name, select the branch and click **Add**.

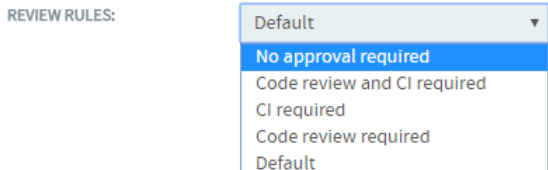


The screenshot shows the 'PROTECTED BRANCHES' field with a text input containing 'mas' and a dropdown menu showing 'master' as the selected branch.



The screenshot shows the 'PROTECTED BRANCHES' field with a text input containing 'Add a new protected branch' and a dropdown menu.

Select one of the review rules from the drop-down list.



The screenshot shows the 'REVIEW RULES' dropdown menu with the following options: Default, No approval required (selected), Code review and CI required, CI required, Code review required, and Default.

5. Click **Save**.

### Create a pull request

When you are ready to merge the change, or at least to start getting feedback, you should create a pull request. You can do this easily from the **Branches** tab by clicking on the **Create** button for your feature branch.

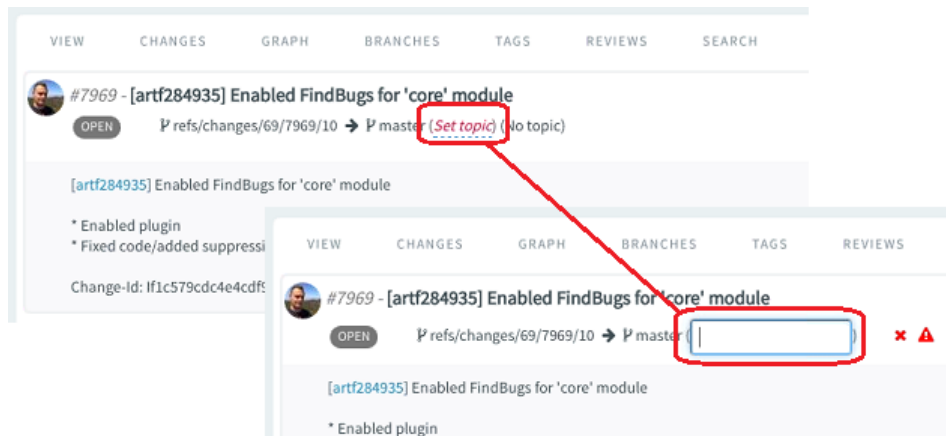
6. Go to the **Branches** tab on a Git repository page.
7. Click **Create**.
8. Select the source branch which is wanted to be merged.
9. Select the target branch to which you want the changes to be merged.
10. Give an appropriate subject line and description that will be used as a commit message for a merged pull request.

Optionally you can provide a summary of the pull request. This supports markdown formatting.

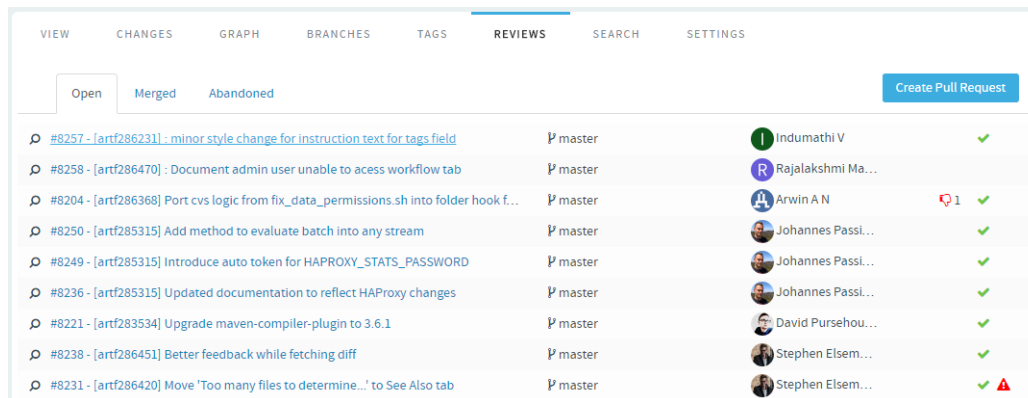
The **Commits** tab at the bottom of the page displays the list of commits made in the selected source branch. The **Files** tab shows the difference between the source and target branch.

11. Click **Create Pull Request**.

12. **Submit whole topic:** You can now bundle related changes (code reviews) by topic and submit the whole topic for review instead of just submitting changes one-by-one. Just open a review, click the **Set Topic** link and enter the topic name.




13. As a reviewer, click the pull request that you want to review from the list of open pull requests on the **Reviews** tab.




### Review a pull request

In addition to the requested reviewers, anyone with access to the repository and who wishes to comment on the pull request can review and post their comments on the pull request details page.

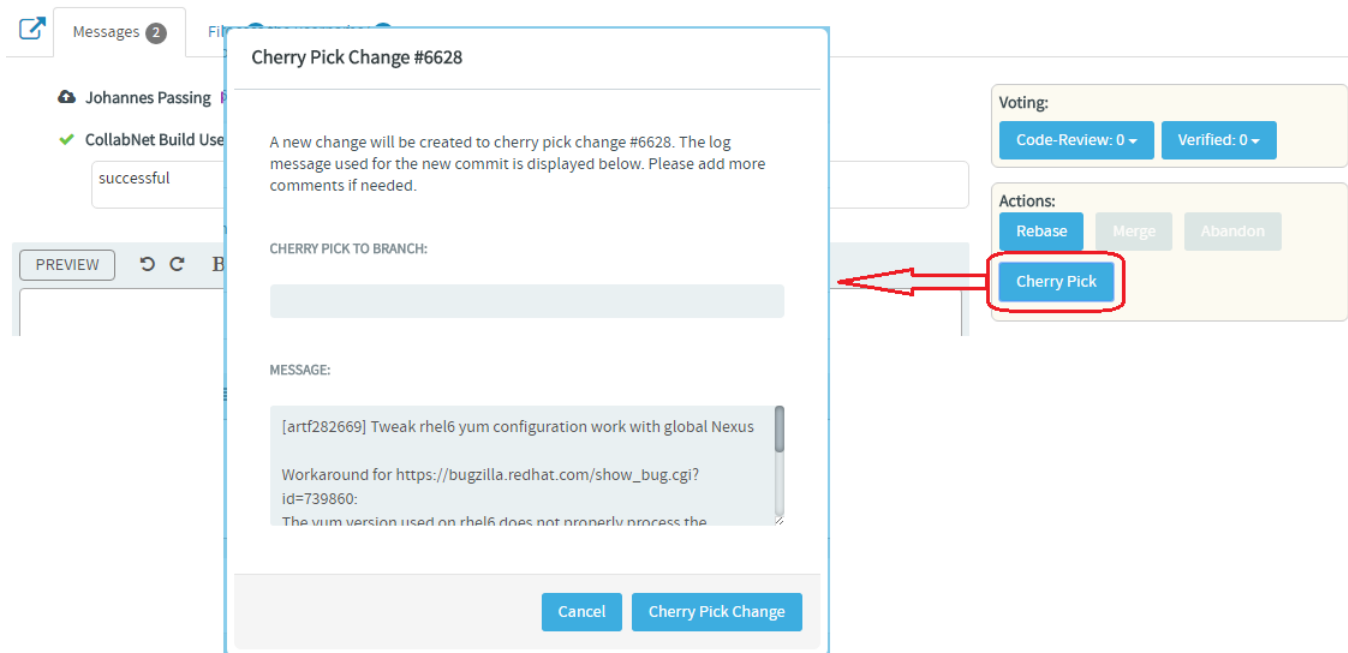
14. On the pull request details page, you can switch between three views: **Messages**, **Commits** and **Files**. Click the **Commits** tab to view the list of commits. Click the **Files** tab to review the code changes made in each file. You have the option to view the difference between the source and target branches.



15. Once you have reviewed the changes, on the **Messages** tab, enter your comments and give an appropriate voting as well.

 **Note:** The message section supports markdown formatting with a preview option.

16. **Cherry Pick: Apply the changes introduced by existing commits:** You can also cherry pick and apply changes introduced by existing commits to another branch. For example, you can now use this Cherry Pick function in TeamForge's native code browser to apply a commit in master to a release branch.



### Inline editing of files

17. Quick changes to files, if required only to few files, can be done using the inline edit feature from within the code browser without having to clone an entire repository. Browse the repository, locate and open the file in the **View**

tab, click **Edit** to open the file in the **File Editor**, make your changes, **Create code review** and **Publish** your changes for review.

18. You can also add new files to a review and delete files from a review by clicking the **Edit Files** icon and then the "+" and "-" icons respectively.




Type the name of the file to see results matching the file name, select a file and click **Add File**.

Type the name of the file you want to delete to see results matching the file name, select the file and click **Delete File**.

19. Click the **Complete File Edits** icon.

### Merge (close) a pull request


Once the pull request is reviewed, it is ready to be merged, that is the **Merge** button on the pull request details page is enabled only if the pull request satisfies all the repository specific qualification criteria. For example, it is possible to merge pull requests even without any voting if “no voting” has been defined as gating criteria. Also, it might require both voting AND acceptance by Continuous Integration; basically it totally depends on the gating criteria of the repository in question.

 **Note:** These criteria are set by the repository owner in **Settings > Policies** tab for the **Pull Request** Repository Category.

20. Click the **Merge** button to merge the source branch into the target branch.

- If the source branch is not updated with the latest changes, a merge conflict is detected prompting you to rebase your request.
- Click the **Rebase** button. Once rebased, the pull request has to be revalidated after which you can merge the pull request into the target branch.

21. The newly-merged pull request is added to the list of merged pull requests.

 **Note:** Once merged, on the **Graph** view, you can see that the merged branch has been added to the graph. A link to the pull request is also provided.

### View pull requests

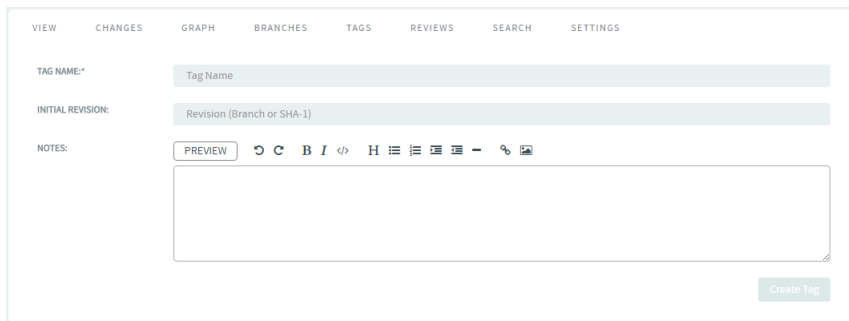
22. Click the **Reviews** tab on a GIT repository page. The pull request details page displays the list of **Open**, **Merged** and **Abandoned** pull requests under appropriate tabs.

- For each pull request, the author, title, number of thumbs up/down, and the time elapsed since the pull request was created, are shown.
- On the **Abandoned** pull request details page, you have the option to restore an abandoned pull request.

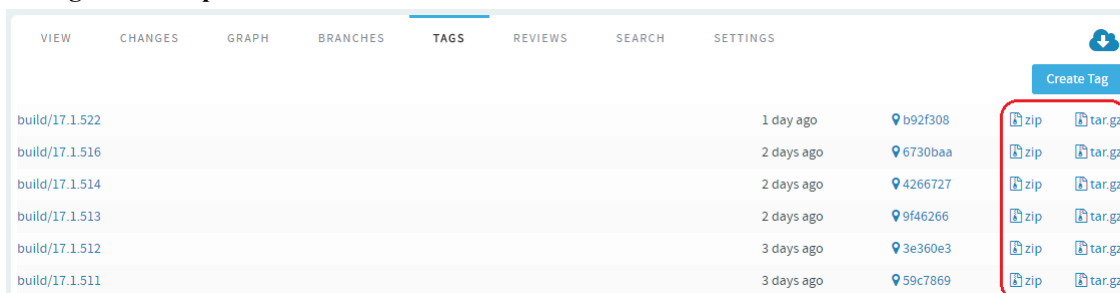
### Tagging

23. To create Git tags and tag specific points in history as being important (to mark release points, for example, v1.0, and so on), select the **TAGS** tab and click **Create Tag**.

24. Type a tag name and revision number and add a Release Notes for the tag. Click **Create Tag**.



Once you create a tag, you can use it to download source code as a zip/tar file and view the tag information in **Changes** and **Graph** tabs.



Tag Name	Time	Commit Hash	Download Links
build/17.1.522	1 day ago	b92f308	zip tar.gz
build/17.1.516	2 days ago	6730baa	zip tar.gz
build/17.1.514	2 days ago	4266727	zip tar.gz
build/17.1.513	2 days ago	9f46266	zip tar.gz
build/17.1.512	3 days ago	3e360e3	zip tar.gz
build/17.1.511	3 days ago	59c7869	zip tar.gz

The top screenshot shows the 'Commits on 01/27/2017' view. The commit list includes:

- b92f308: Add validation to moveDocument3() SOAP API** (by Subhashini Marimuthu) with tags #8329, master, and build/17.1.522.
- 0322dae: [artf286666] Planning Folder REST API is not working** (by Annapoorni Dakshinamurthy) with tag #8328.

The bottom screenshot shows the 'GRAPH' view of the commit history. The commit list includes:

- b92f308 2 days ago** by Subhashini Marimuthu: Add validation to moveDocument3() SOAP API (tags #8329, master, build/17.1.522).
- 0322dae 2 days ago** by Annapoorni Dakshinamurthy: [artf286666] Planning Folder REST API is not working (tag #8328).
- a98573a 2 days ago** by Rajesh Vadivelu: Merge "[artf286653] : Column Configurations are Not saved,"
- 7b81c82 2 days ago** by Rajalakshmi Madasamy: [artf286585] : Permission denied to view documents when logged in user has permissions
- 53448d8 2 days ago** by Rajesh Vadivelu: [artf286653] : Column Configurations are Not saved, (tag #8323)

## Search code

Integration with Black Duck Code Sight (BDCS) is no longer supported in TeamForge 17.1 and later. TeamForge 17.1 (or later) is equipped with its own code search function powered by Elasticsearch. You can do away with BDCS integration while upgrading to TeamForge 17.1 (or later) and set up TeamForge Code Search, which is now one of the integral services of TeamForge. This section discusses the features of TeamForge Code Search and what it takes to set up Code Search on your site.

## TeamForge Code Search powered by Elasticsearch

TeamForge's Code Search feature uses [Elasticsearch](#) as its engine for the indexing and retrieval of documents. Elasticsearch, used in TeamForge, has no customizations other than what it provides via its engine and API. For more information about Elasticsearch, refer to its [documentation](#).

## TeamForge Code Search (Elasticsearch): How it works

TeamForge uses Elasticsearch engine 5.1.1. Elasticsearch provides the engine for the indexing and retrieval of documents. TeamForge is equipped with its own indexers that extract source files from the SCM repositories and send them to the Elasticsearch engine for indexing. These indexers run as an hourly cron job. When the job starts, it calls an internal TeamForge API that returns the list of repositories to be indexed on the specified SCM integration server. The indexer then processes the repositories, one at a time, starting with the repository that was most recently committed to. For each repository, it determines if a full indexing is needed or just an incremental indexing is sufficient to catch up since the last indexing operation. It then goes down the list of files and processes them for indexing. Only files that need to be indexed are extracted. For example, large binary files are not extracted as they would not be indexed anyway. As the indexers crawl a repository, it would skip any file that is greater than 1MB. To index a file it has to be UTF8 text. Files are scanned for file encoding and converted to UTF8 if required. Binary files are skipped by the indexers as would be files that cannot be converted to UTF8 format. All files that are skipped are logged along with the reason.

The engine in Elasticsearch is configured to convert the source to lower case and tokenize it. A Camel case filter is used so that common programming elements are broken into individual tokens based on camel case and other common separators used in function names etc. A set of common programming terms (if, then, else, return, exit etc.) are not indexed by default.

How the data is indexed does effect the search results as partial matches are not generally returned. For example, consider a function name in the source code such as: TeamForgeHelper. This would be indexed so that a search for

any of these terms would find a match: Team, Forge, Helper, TeamForgeHelper. Nothing else would match, including "TeamForge" or "Tea".

Search results are returned to users with TeamForge RBAC applied to it at a repository level. In other words, a user must have permission to view the entire repository to be able to search it. This works with SVN Path-based permissions (PBP), but only if PBPs are being used to restrict write access. If PBP's are used to deny the user view access to certain paths in the repository, then TeamForge Code Search denies the user the ability to search the repository.

### Considerations while upgrading to TeamForge 17.1 (or later)

Consider the following points while upgrading from TeamForge 16.10 (or earlier) to TeamForge 17.1 (or later) versions.

- Code Search is now an integral part of TeamForge, which is installed by default during TeamForge installation.
- You can install TeamForge Code Search either on the TeamForge Application Server or on a separate SCM integration Server. It is recommended to install TeamForge Code Search on a separate server if your site's indexing needs are considerably high (large number of repositories, for example).
- TeamForge Code Search works with GIT and SVN repositories. TeamForge Code Search has no support for CVS repositories.
- Installation of Elasticsearch is determined by adding the "codesearch" identifier to the SERVICES token of either the TeamForge Application Server or the SCM Integration Server (if Code Search runs on a separate server). Refer to the [site-options.conf](#) configuration section of TeamForge [installation/upgrade](#) instructions for more information.
- Elasticsearch needs 2GB of JVM heap size by default on a TeamForge site. You must have adequate RAM to accommodate the JVM heap requirements of Elasticsearch in addition to the JVM heap requirements of other components such as Jboss, integrated applications, and so on.
- If required, you can increase the JVM heap size for Elasticsearch. Set the [ELASTICSEARCH\\_JAVA\\_OPTS](#) token with the Elasticsearch JVM heap size required for your site.
- Elasticsearch stores every document, which for code search is each source file that is indexed, and it has its index itself as well. As a rule of thumb, the index for a repository would be roughly the same size as the working copy for that repository but in practice it will likely be smaller. Consider that all binary files and all files greater than 1 MB are not indexed at all. So, obviously any repository that has a large working copy due to these types of files will have a much smaller index. By default, Elasticsearch compresses the original files using LZ4 compression. It also offers a "best\_compression" codec that compresses the originals using DEFLATE algorithm. In TeamForge, indexes have been updated to use the maximum compression.
- The Elasticsearch data and logs are stored in `/opt/collabnet/teamforge` folder. Log rotation, startup and so on are all managed the same way as other TeamForge services.
- As Black Duck Code Sight (BDCS) is not supported:
  - remove all BDCS tokens from your [site-options.conf](#) file while upgrading to TeamForge 17.1 (or later). TeamForge create runtime fails otherwise. See [Site options change log](#) on page 730 for a list of obsolete `site-options.conf` tokens.
  - there is no migration support for existing BDCS indexes. All the repositories should be indexed afresh post upgrade to TeamForge 17.1 (or later). The list of repositories to index, remains the same though. After upgrading to 17.1 (or later), whatever repositories were being indexed by BDCS are indexed by the new Code Search without any user intervention. However, if you had customized indexing for one or more repositories using the BDCS administration UI, such information will not have been migrated and would need to be done again, but that can now be done via the TeamForge UI.

#### How to search for Source Code

Searching repositories for code snippets can be done via the Jump to ID search or via the Advanced Search.

Just select **Source Code** from the **Jump to ID** drop-down list, type the search keyword and press enter to search the repositories of the project in context. You can also select **Advanced Search** from the **Jump to ID** drop-down list and do an advanced search for the required code.

1. Click **Advanced Search** from the **Jump to ID** drop-down list.


- Type the search keyword.
- Select **Source Code** from the list of components.
- Set the scope of search: Select All Languages (default) or one of the programming languages such as C, C++, C# and so on from the All Languages drop-down list.
- Select either All Projects or a selected few projects (select one or more projects) from the **IN PROJECTS** list.

- Click **Search**.

The search results are displayed in **Files** and **Commits** tabs.

- In case you want to search for code in a specific repository, select **Project Home** > **Source Code**, select a repository you want to search, click **Browse Repository**, and select the **Search** tab.

You can search all files in a repository or narrow your scope to specific file types such as C, C++, C# and so on.



8. Type your search keyword, select a file extension (optional) and click **Search**. The search results are displayed in **Files** and **Commits** tabs.

### Manage binaries

TeamForge 8.0 and later support Nexus integration. TeamForge 16.7 and later support Artifactory integration.

Before you can create binary repositories and work with them, a TeamForge site administrator must set up one or more Nexus or Artifactory servers and integrate the servers with TeamForge.

 **Important:** TeamForge-Nexus and TeamForge-Artifactory integrations are not supported in SUSE Linux platform. See [TeamForge installation requirements](#) on page 441.

### An overview of binary artifacts and repository managers

Here is the design overview of integrating TeamForge, a full ALM suite with binary repository managers.

An important aspect of the end-to-end development lifecycle is the creation and storage of software packages that are often binary artifacts. In the Java world, these are usually reusable jars that are used by other projects. Binary artifact repository managers are software systems that manage, version, and store binary artifacts. Examples of such repository managers are *JFrog Artifactory*, and *Sonatype Nexus*.

### What is a binary artifact repository?


A binary artifact repository stores binary artifacts along with the metadata in a defined directory structure, conceptually similar to a source code repository. The metadata describes the binary software artifact and includes information such as dependencies, versioning, and build promotions. Maven is the widely used tool for dependency management, especially for Java projects. Maven represents dependencies in an XML file called Project Object Model (POM). Other tools can use similar approaches to store documentation archives, source archives, Flash libraries and applications, and Ruby libraries.

### How does a binary artifact repository manager help?

Some of the advantages of using a binary artifact repository manager are:

- **Dependency management:** Nexus and Artifactory can act as Maven repositories. Maven is a widely used Java dependency management and build tool.
- **Efficient builds:** With the help of a binary artifact repository manager, you can save the download time from public repositories as the artifacts once downloaded are cached locally.
- **Predictability and release stability:** Once published onto a release repository, the binary artifact and metadata do not change. It ensures predictable and repeatable builds.
- **Control and audit:** If you want to standardize libraries that are used in your software, the binary artifact repository helps track the versions of your software components. Also it enables you to audit the licenses of your third-party components used in your software.
- **Promotes collaboration:** The binary artifact repository enables you to share components with other teams.

### How to integrate Nexus with TeamForge?

 **Important:** TeamForge-Nexus integration is not supported in SUSE Linux platform. See [TeamForge installation requirements](#) on page 441.

TeamForge supports integration with Nexus in both the ALM and SCM modes. Nexus integration has been tested by CollabNet for Nexus and Nexus professional versions 2.9, 2.10, and 2.11.

To integrate Nexus with TeamForge:

1. Download and install the Nexus OSS if you do not have a Nexus instance running.

2. Download and install the CollabNet Nexus integration plugin.
3. Change your build system and use the CollabNet supplied Maven deploy plugin for end-to-end traceability.
4. Set up the TeamForge EventQ activity source to provide your teams with end-to-end visibility from requirements to source code all the way to deployed binary artifacts.

### Installing the TeamForge Nexus integration plugin

You need to have the following information handy before you start off with the installation:


- Installation path of the running Nexus instance.
- TeamForge host's URL.
- TeamForge site administrator credentials.
- A suitable name for your Nexus instance; the Binaries App in TeamForge refers to this name.

You must have a Nexus instance running for the integration. If you are upgrading from an earlier version of the plugin, ensure that the old plugin is completely removed from the directory and the new plugin is unzipped on the same directory before you restart the Nexus instance.

- To set up a Nexus server, see [Install Nexus](#) on page 669.
- To install the TeamForge-Nexus integration plugin, see [Install the TeamForge-Nexus integration plugin](#) on page 670.
- To upgrade the TeamForge-Nexus integration plugin, see [Upgrade the TeamForge-Nexus integration plugin](#) on page 671.

**Accessing Nexus through TeamForge:** You have to introduce a TeamForge project context in Nexus and allow authentication to use TeamForge credentials for logging into Nexus directly. Accessing Nexus through the TeamForge project toolbar provides you with Single Sign-on (SSO). It logs you into Nexus automatically with the project context. You can allow RBAC (Role Based Access Control) using TeamForge roles.

### How to integrate Artifactory with TeamForge?

 **Important:** TeamForge-Artifactory integration is not supported in SUSE Linux platform. See [TeamForge installation requirements](#) on page 441.

TeamForge supports integration with Artifactory in both the ALM and SCM modes. Artifactory integration has been tested by CollabNet for Artifactory Pro 4.7.

To integrate Artifactory with TeamForge:

1. You must have an Artifactory Pro instance running for the integration. Download and install the Artifactory Pro 4.7 if you do not have an Artifactory Pro instance running.
2. Download and install the CollabNet Artifactory integration plugin.
3. Change your build system and use the CollabNet supplied Maven deploy plugin for end-to-end traceability.
4. Set up the TeamForge EventQ activity source to provide your teams with end-to-end visibility from requirements to source code all the way to deployed binary artifacts.

### Installing the TeamForge Artifactory integration plugin

You need to have the following information handy before you start off with the installation:

- TeamForge Host URL.
- TeamForge site administrator credentials.
- A suitable name for your Artifactory instance; the Binaries App in TeamForge refers to this name.
- Artifactory Host URL.
- To set up an Artifactory server, see [Install Artifactory Pro](#) on page 674.
- To install the TeamForge-Artifactory integration plugin, see [Install the TeamForge-Artifactory integration plugin](#) on page 675.

**Accessing Artifactory through TeamForge:** You have to introduce a TeamForge project context in Artifactory and allow authentication to use TeamForge credentials for logging into Artifactory directly. Accessing Artifactory through

the TeamForge project toolbar provides you with Single Sign-on (SSO). It logs you into Artifactory automatically with the project context. You can allow RBAC (Role Based Access Control) using TeamForge roles.

## Authentication policies

### Nexus

Your site administrator can enable the integration with the following two authentication mechanisms:

- TeamForge and native Nexus login (default)
- TeamForge only

In both the cases, you can use your TeamForge credentials to log on to Nexus. If your Site Administrator has used the default setup, you can use your pre-existing Nexus credentials.

### Roles and permissions

Following are the two administrative privileges in Nexus:

- Nexus Admin (Site Admin in TeamForge will be a Nexus Admin)
- Project admin (permissions to create, update, and delete binary artifact repositories.)

For all the other users, privileges are based on the TeamForge RBAC setup.

### Artifactory

Your site administrator can enable the integration with the following two authentication mechanisms:

- TeamForge and native Artifactory login (default)
- TeamForge only

In both the cases, you can use your TeamForge credentials to log on to Artifactory. If your Site Administrator has used the default setup, you can use your pre-existing Artifactory credentials.

### Roles and permissions

- Artifactory Admin (Site Admin in TeamForge will be an Artifactory Admin)
- A user with any one of the following permissions in TeamForge becomes an Artifactory Admin: **VIEW/CREATE REPOSITORIES**, **VIEW/UPDATE REPOSITORIES** and **VIEW/DELETE REPOSITORIES**.
- All other permissions in TeamForge are mapped as is in Artifactory.

For all the other users, privileges are based on the TeamForge RBAC setup.

## Known limitations with TeamForge-Nexus/Artifactory integrations

Here is a list of known limitations of TeamForge-Nexus integration.

### TeamForge-Nexus integration

- While it is possible for users (all TeamForge users but site administrators) with 'create repository' permission to create binary repositories via the TeamForge **Binaries** application, such users cannot



create repositories directly on the Nexus server (using the Nexus UI) as there is no TeamForge project mapping available in Nexus.

- A TeamForge user cannot be configured as an anonymous user in Nexus as TeamForge users are not available in the Nexus database.
- TeamForge broadcast messages and license notifications are not visible in Sonatype Nexus pages in TeamForge. This is due to a limitation with the TeamForge-Nexus integration plugin.

## TeamForge-Artifactory integration

Artifactory repository names must be unique. You cannot have two or more repositories with the same name even if the repositories are of different types such as Remote or Hosted repositories. However, if you want to create two (or more) repositories with like-sounding names, you can devise and follow a repo naming convention to uniquely identify repositories of different types. For example, a "central" repo can be named as "central-local" and "central-remote" to uniquely refer to the "Hosted" and "Remote" repository types respectively.


### Create a binary artifact repository

With TeamForge—Nexus integration enabled, each project can have one or more binary repositories.


#### Remember:

Before you can create a repository for binary artifacts, a site administrator must set up TeamForge-Nexus integration and add one or more binary servers to your TeamForge site.

1. Click **BINARIES** from the **Project Home** menu.
2. Click **Create Repository**.  
The **Create Repository** page appears.
3. **Server**: Select a Nexus or Artifactory binary repository server from the drop-down list.
4. If you are creating a Nexus repository, type or select the values for the following fields in the **Create Repository** page.

 **Tip**: Hover your mouse over the Help icon (question mark icon) for relevant tooltip.

- **Repository Name** and **Repository ID**: Type a name and ID for the binary repository.
  - **Repository Type**: Select one of the following types from the drop-down list: Hosted Repository or Proxy Repository or **Virtual Repository**.
  - **Provider**: Select the content provider of the repository from the drop-down list.
  - **Repository Policy**: Repositories can store artifacts of a release or snapshot or both. Select a policy from the drop-down list.
  - **Override Local Storage Location**: Leave this blank to use the default local storage location. To change the local storage location, type the path (URL) in the following format.
    - **Windows**: `file://{drive-letter}:/`
    - **Other operating systems**: `file://`
- a) Modify the default access settings such as **Deployment Policy**, **Allow File Browsing**, **Include in Search** and **Publish URL**, if required.
  - b) Modify the **Not Found Cache TTL** duration (in minutes), if required.
  - c) Click **Save**.
5. If you are creating an Artifactory repository, type or select the values for the following fields in the **Create Repository** page.

 **Tip:** Hover your mouse over the Help icon (question mark icon) for relevant tooltip.

- **Repository Type:** Select one of the following types from the drop-down list: Local Repository or Remote Repository or **Virtual Repository**.
  - **Package Type:** Type the package type. For example, "maven". Package type cannot be changed once the repository is created.
  - **Repository Key:** Type the repository key. The repository key must be unique within an Artifactory instance. Can contain no spaces or special characters. Cannot begin with a number.
- a) Fill in the fields (modify the default values, if required) listed under the "**General**", "**Include/Exclude Patterns**", and "**Maven Settings**" sections. Hover your mouse over the Help icon (question mark icon) for relevant tooltip.
  - b) Click **Save**.

The binary repository is created.

After creating a binary repository, you can now configure traceability for the repository so that your Nexus or Artifactory runtime events are published to EventQ. See [Configure traceability for binary artifact repositories](#) on page 114.

### Link binary artifact repository to project

You can link existing binary artifact repositories, if any, to your project.

1. Click **BINARIES** from the **Project Home** menu.
2. Click **Link Existing Repository**.
3. Select a repository from the list of binary repositories and click **Link to Project**.

### Configure traceability for binary artifact repositories

Once a binary repository is created, configuring the repository to publish runtime data to EventQ is the next step that lets you track and visualize binary artifacts in real time using TeamForge EventQ.

Configure traceability for a binary repository.

1. Click **BINARIES** from the **Project Home** menu.
2. Select the binary repository for which you want to configure runtime.
3. Click **Configure Traceability**.
4. Click **Create**.
5. Change the **Associated Source** from the drop-down list and click **Update** to change the associated source.
6. Clicking **De-Activate** and **Re-Activate** disables and enables traceability for the repository respectively.
7. Clicking **Force Delete** deletes the traceability configuration including all historical associations. Exercise caution before deleting the traceability configuration.
8. Click **Close** to exit the **Traceability Configuration** window.

### Configure Maven build

CollabNet provides you with a Maven deploy plugin that completely inherits the default (vanilla) Maven deploy plugin. Additionally, it posts binary artifact creation event information on to TeamForge EventQ enabling end-to-end traceability.

To configure your Maven build:


1. Replace the standard deploy with CollabNet's plugin in your POM.xml.

```
<pluginRepositories>
  <pluginRepository>
    <id>collabnet</id>
    <name>Collabnet Public Repo</name>
    <url>http://mvn.collab.net/nexus/content/groups/public/</url>
  </pluginRepository>
</pluginRepositories>
```

```

<plugins>
  ....
  <plugin>
    <groupId>org.apache.maven.plugins</groupId>
    <artifactId>maven-deploy-plugin</artifactId>
    <version>2.8.2</version>
    <configuration>
      <skip>true</skip>
    </configuration>
  </plugin>
  <plugin>
    <groupId>net.collab.maven.deploy</groupId>
    <artifactId>collabnet-deploy-maven-plugin</artifactId>
    <version>8.2</version>
    <extensions>>true</extensions>
    <executions>
      <execution>
        <id>default-deploy</id>
        <phase>deploy</phase>
        <goals>
          <goal>deploy</goal>
        </goals>
      </execution>
    </executions>
    <configuration>
      <skipLinkToBinaries>true</skipLinkToBinaries>
      <amqpServer>amqp://host.with.amqp</amqpServer>
      <sourceAssociationKey>as-per-orchestrate-setup</
sourceAssociationKey>
      <associatedBuildNumber>${env.BUILD_NUMBER}</associatedBuildNumber>
    </configuration>
  </plugin>
</plugins>

```

 **Important:** Make sure the `<skip>` tag is set to true to prevent more than one Nexus/Artifactory notification for a single Nexus/Artifactory artifact deployment. If `<skip>` is not set to true, notifications are sent by both the `maven-deploy-plugin` and the `collabnet-deploy-maven-plugin` for a single binary artifact.

2. Setup the Nexus/Artifactory and EventQ credentials in `settings.xml`. You may find this file in the Maven home directory.

For example, in the following illustration, your distribution management section has a repository id of the local-nexus and the `amqpServer`, `amqp://host.with.amqp` is setup as mentioned in Step 1:

```

<settings>
  <servers>
    <server>
      <id>local-nexus</id>
      <username>your_ctf_username</username>
      <password>xxxxxxxx</password>
    </server>
    <server>
      <id>amqp://host.with.amqp</id>
      <username>myproj_nexus</username>
      <password>xxxxx</password>
    </server>
  </servers>
</settings>

```

If you already have a set of credentials against an AMQP URL in `settings.xml`, all jobs including the one that posts events against different sources can use the same credentials.

**Table 1: All available configuration items:**

Configuration parameter	Description	Mandatory	Default value	Example
amqpServer	An amqp url for the form <code>amqp://host.domain</code> where EventQ is listening to.	True	None	<code>amqp://example.com</code>
sourceAssociationKey	The source association key from orchestrate custom event setup.	True	None	uuid format
orchestrateSchemaId	The schema id posted to EventQ. Do not set it in v1.4.1 and later.	True	<code>binary_artifact</code>	<code>binary_artifact</code>
orchestrateSchemaVersion	The version of the schema, usually 1. Omit to default to 1. Do not set it in v1.4.1 and later.	False	1	1
amqpUserName	The AMQP user name. It is recommended not to set it in the project POM but using in settings.xml. Take it from the custom event setup in EventQ.	False	None	<code>proj_custom</code>
amqpPassword	The AMQP password. It is recommended not to set it in the project POM but using in settings.xml. Take it from the EventQ settings.	False	None	<code>af78aaec0c5610d21...</code>
amqpQueueName	Do not specify. Defaults to <code>orchestrate.custom</code> .	False	<code>orchestrate.custom</code>	<code>orchestrate.custom</code>
associatedBuildNumber	Specify to the env variable depending on your build system process. Set to <code>\${env.BUILD_NUMBER}</code> for jenkins and teamcity.	False	None	<code>\${env.BUILD_NUMBER}</code>
skipOrchestrateNotification	Set to true to disable notification.	False	False	False
dieOnOrchestrateError	Defaults to true. Set to false to treat orchestrate notification errors as non-fatal.	False	True	True
component	Used to identify a specific binary artifact as a component in a larger application.	False	none	An ALM platform has several components such as an application server, an indexer, an SCM integration server and so on. These components have their own build process. This property is used to uniquely identify such components in EventQ.


Configuration parameter	Description	Mandatory	Default value	Example
componentOf	Associated with the 'component' parameter to store the details of the component.	False	none	SCM as a component of Teamforge.

### Delete a binary artifact repository

You need to have the required permission to delete binary repositories. As a project admin, you can use the delete option only to unlink the repository from a project. If you want to delete a repository from the Nexus, login as a Nexus admin.


1. Click **BINARIES** from the **Project Home** menu.
2. In the list of the repositories, select the repository you want to delete and click **Delete**.  
The following confirmation message appears: This repository will no longer be accessible. Your site administrator can relink it later. Are you sure you want to unlink this repository?
3. Click **OK** to delete.

The repository is deleted.

 **Note:** This deletion disassociates the repository from your project; only the Site Admin can reinstate the repository.

### Provision hardware

Use TeamForge Lab Management to access virtual or physical machines where you can build and test your code.

 **Tip:** For more information about the new features in TeamForge Lab Management 2.2, see [this introductory video](#). (You'll need a free open.collab.net membership to watch it.)

### Get an account on this system

If you are new to this system and require access, contact a TeamForge (or CEE) administrator in your organization.

The administrator will add a role to your TeamForge (or CEE) account which will enable you to access resources on the site.

### What Lab Management systems are assigned to me?

The easiest way to see what Lab Management systems are assigned to you is to check the **My Pages** page.


Each system assigned to you is listed on the **My Pages** page, organized by project.

For more details on a particular system, click the name of the system. On the host status page, you can get more information about the host's configuration and performance, release the system, and more.

### Find an available host to use

To work in the Lab Management environment, you need to either allocate a system to yourself or log into someone else's system in your project.


1. Click **My TeamForge Lab Management** or the **Projects** tab and select your project.

 **Important:** It is usually considered impolite, and may even be a violation of your company's policies, to check out a system from one project that you are a member of in order to do work related to another project.


2. Click **Project Hosts**.

3. In the list of hosts assigned to the project, do what you need to do:

- If the host is allocated, click the user's name to ask their permission to work on their host.

 **Tip:** You may also ask the user if they still need the host, and would they mind releasing the system so that you or someone else, could check out the host.

- If the host is in the **Free** state, allocate it to yourself.

-  **Note:** It is polite to ask the owner of a machine if you can use it to log in and run jobs on. TeamForge Lab Management assumes that all TeamForge Lab Management users in a project should be able to log in to all systems in a project, to facilitate developer and QA productivity. But, along with this freedom comes the responsibility to be a "good citizen" and not step on other people's resources.

### Allocate yourself a Lab Management machine


When you have been given membership and an appropriate role in a TeamForge Lab Management-enabled project, allocate a host for your own use.

Your personal host can be rebuilt with any profile whenever you want.

You can install any software you want on it (subject to the CollabNet Terms of Service and your own organization's security policy, all depending on how and where your TeamForge Lab Management instance is deployed).

You can run any jobs you want on the server (subject to the same conditions as software installation).

1. Find an available host to allocate.
2. On your project's TeamForge Lab Management home page, click the host you've identified.
3. On the host status page, click **Allocate** to reserve the host for yourself.
4. Specify how long you wish to allocate the host.
  - Choose a specific time/date to deallocate the host using the calendar widget.
  - Choose a length of time to allocate the host.
  - Allocate the host indefinitely (until you or a project admin frees the host).


-  **Note:** Your project administrator may have set a limit on the amount of time a machine can be allocated. If this is the case, then you will receive an error message if you try to allocate a host for a longer period.

You can always change the allocation time for your host later, if you decide that you need the host for a shorter or longer amount of time (subject to the maximum limits set by your project administrator).

### Compare profiles or profile versions


To choose a profile for your TeamForge Lab Management system, compare different profiles and profile version in the Profile Library.

1. On your project's home page, select **Profile Library**.  
The table shows all profiles in your project.
2. Choose a profile from each of the sets of profiles and click **Compare selected profiles**.

-  **Note:** Profiles must first be added to your project as "allowed profiles". It is not possible to compare profiles before you have added them to your project.


3. On the **Profile Compare** screen, select the versions of each profile to compare.

The default version is the latest version of the profile, but any version of a profile can be compared against any version of any other profile.

-  **Note:** Any two profiles can be compared, even if they are from two different operating systems. For example, TeamForge Lab Management will let you compare a Solaris profile and a Windows XP profile (but do not expect the results to be very meaningful).

### Create a profile from a host

Let's say you've built a generic Lab Management host and customized it for your team. With Lab Management 2.4, you can create a profile based on the host image and store it in the Project Build Library (PBL) for your project.

-  **Note:**

- You can only create profiles from ESX guests.
- These OS profiles are supported: RHEL 5, CentOS 5, SLES/11, Windows XP and And Windows Server 2003. Windows Server 2008 and Solaris profiles are not supported.

When creating a profile that contains user-specific software or services, here are a couple of things you should be aware of:

- If the software or service is run by a user who is a locally-created user (say, an 'oracle' user who was manually created on the node), the profile should be easily portable to other projects.
- If the software or service is run by a user who is a Lab Management user, the profile will only work in projects where that user has access. In projects where that user has no access, it's not likely that the profile will work well - it may not work at all.

1. Bring up your host's **Admin** tab.
2. In the **Create Profile from This Node** section, enter a name for the profile.  
Only alphanumeric characters, hyphens, underscores, and periods are allowed.

**Create Profile from This Node**

Create a new profile based on this guest. This guest will be inaccessible while the new profile is being created.

<b>Profile Name:</b>	<input type="text"/>	Profile names may only contain alphanumeric characters, underscores, and periods (i.e., "a-z0-9_.-").
<b>Profile Description:</b>	<input style="height: 40px;" type="text"/>	
<b>Storage Location:</b>	Project freedom public PBL storage ▼	

Create Profile from cu035.h.sp.collab.net


3. Provide a description for the profile.
4. Select the PBL location for the project.  
The profile will be stored here.
5. Click **Create Profile from <hostname>**.

Your host will be switched to a clone, and inaccessible for few minutes. When it gets back to the allocated state, the profile will be created. The profile will be available in your project and in any project that allows public profiles.

### Rebuild a system

You can rebuild a host with a specific profile that you choose from within your project.

Usually you will want the latest (highest) revision of the profile you are working with. However, profiles are versioned, like source code. This means that you can always build with an older version of an operating system or package if you need to.

 **Note:** Profiles are assigned to projects by your project administrator.

1. On the **Project Host properties** screen, select the host to rebuild.
2. Click the **Configuration** tab.
3. Select a profile from the **Profile** drop-down list and click **Rebuild**.

You can check the status of the host's progress by refreshing your browser.


### Log into a TeamForge Lab Management system

You can log into any machine that is assigned to a project you are a member of and that you have the requisite permissions for.

For security reasons, all TeamForge Lab Management systems are on a private network, inaccessible from anywhere outside of CollabNet. To access these systems, you must tunnel through the TeamForge Lab Management Manager.

If you need access to a system that belongs to a project you are not a member of, you have two options:

- Ask to join the project, or to get permission within that project to access TeamForge Lab Management hosts.
- Ask the owner of the system, or a project/domain administrator, to create an account for you on the machine.

-  **Note:** This may violate your organization's security policy, so check with your site administrator before creating any accounts. If the account is temporary, you will be responsible for removing it.

### Log into a UNIX system

To log in to UNIX systems, such as Linux and Solaris machines, you must use port forwarding over SSH.

Port forwarding over SSH is a secure way to tunnel traffic over the Internet.

1. Set up your SSH tunnel.
  - a) From a Linux or Solaris system, or from a Windows system with Cygwin, the command to set up an SSH tunnel to "cu086.\$vars.mgr\_domain" through the TeamForge Lab Management Manager host "\$external\_host" would look like this:

```
ssh -L22000:cu086.$vars.mgr_domain:22 $external_host -l username
```

- b) From a Windows system without Cygwin installed, set up the tunnel using PuTTY.

All connections to "localhost" port number 22000 are now tunneled to "cu086.\$vars.mgr\_domain" port number 22 (SSH) through the TeamForge Lab Management manager host "\$external\_host"


2. Log into "cu086.\$vars.mgr\_domain" by issuing this command:

```
ssh localhost -p 22000 -l username
```

### Log into a Windows system

To log into a Windows system in TeamForge Lab Management, you must use SSH with the Remote Desktop Protocol (RDP), which renders a full Windows graphical environment on your desktop.


The TeamForge Lab Management environment includes Windows XP Professional and Windows 2003 Server machines.

-  **Tip:** You can automate this whole process with a tool called Wissh. CollabNet does not provide support for this tool, but it is something which we recommend.

1. Set up your SSH tunnel to the TeamForge Lab Management machine.
  - a) If you have Cygwin installed, use a command like this. For example, "cu087.\$vars.mgr\_domain", through the TeamForge Lab Management Manager host "\$external\_host"

```
ssh -L3390:cu087:3389 $external_host
```

- b) If you do not have Cygwin, set up the tunnel using Putty.
2. Start Remote Desktop by clicking **Start > Programs > Accessories > Communications > Remote Desktop Connection**.

-  **Note:** If you are using Linux to access a Windows system, you can use the `rdesktop` program, which provides the same functionality.

3. In the Remote Desktop Connection dialog box, enter `localhost:3390` and click **OK**.  
This initiates a remote desktop connection to the TeamForge Lab Management host.
4. At the login prompt, enter your TeamForge Lab Management username and password.  
If you are an authorized user on the machine, you are granted access and the correct privileges.
5. If your remote Desktop Connection to `localhost` does not work, try this workaround:
  - a) Copy the two files `mstsc.exe` and `msctax.dll` to another folder.  
These files should be in the directory `C:\WINDOWS\system32`, or in `C:\Program Files\Remote Desktop` if you installed the Remote Desktop Client from the Windows 2003 CD.
  - b) Right-click the `mstsc.exe` file and select **Properties**.
  - c) On the **Compatibility** tab, select **Run this program in compatibility mode for Windows 98**.
  - d) Start up the new copy of `mstsc.exe` from the new location.  
You should now be able to connect to `localhost` at IP address 127.0.0.1.




### Set up a tunnel with PuTTY

Summary to come...

1. Launch the PuTTY application.
2. For **Host name**, enter the TeamForge Lab Management Manager host "\$external\_host".
3. For **Connection type**, enter "SSH".
4. In the left pane, open **Connection list**.
5. Open **SSH list** and click **Tunnels**.
6. In the right pane, enter "3390" for **Source port** and "cu087.\$vars.mgr\_domain:3389" for **Destination**, and click **Add**.
7. Select the forwarded port from the list and click **Open**.
8. At the login prompt, enter your TeamForge Lab Management username and password.

The RDP tunnel is now set up.

 **Note:** This procedure is equivalent to running this SSH command:

```
ssh -L3390:cu087.$vars.mgr_domain:3389 $external_host
```


### Set up a VNC server on TeamForge Lab Management

To allow remote access to an X Windows display on Linux from a remote computer, set up a VNC console.

VNC is a convenient and effective way to access an X Windows display from a remote computer. Some applications require X, and since users don't have access to consoles on TeamForge Lab Management nodes, VNC might be the only way those applications can be accessed.

1. Set a VNC password using the `vncpasswd` command.

```
[grue@cu012 grue]$ vncpasswd
Password:
Verify:
```

 **Tip:** This doesn't need to be the same as your TeamForge Lab Management password, but keep in mind that this password will give the bearer access to an X session as your username, so keep it something reasonably hard to guess.


2. Edit the file `/etc/sysconfig/vncservers`. Add an entry for your session number, defaulting to 1, followed by your username.

Use this as an example file:

```
VNCSERVERS="1:grue"
```


3. Start the VNC service.

```
[grue@cu012 grue]$ sudo /etc/init.d/vncserver start
Starting VNC server: 1:grue      [ OK ]
```

 **Note:** To stop the VNC server on your machine, run the command like this: `[grue@cu012 grue]$ sudo /etc/init.d/vncserver stop`

4. If you want VNC to start up each time this server is rebooted, use the `chkconfig` command to turn it on.

```
[grue@cu012 grue]$ sudo /sbin/chkconfig vncserver on
```


 **Note:** To prevent the VNC server from starting up on your machine when it is rebooted, run the command this way: `[grue@cu012 grue]$ sudo /sbin/chkconfig vncserver off`

### Connect to a VNC server from your desktop

To access an X Windows display from a remote computer, install and connect a VNC client.

VNC saves your desktop session and allows you to connect to it later. For example, you can start an X session at work, then log off your computer, go home, start up your computer again, and resume your same X session. In this sense it is very similar to the `screen(1)` utility, which allows users to create, disconnect and reconnect to text-based terminal sessions.

1. Install a VNC client on your computer.


 **Tip:** Either of these recommended packages is available for all major desktop operating systems.

- TightVNC
- RealVNC

2. Set up an SSH tunnel to port 5901 on the proper TeamForge Lab Management node.

For example, let's say you are the user `grue` and you have set up a VNC server on the machine `cu012`. `$vars.mgr_domain`. The port forwarding command would look like this:

```
ssh -L5901:cu012.$vars.mgr_domain:5901 $external_host
```


 **Note:** Multiple `-L` options can be combined on the same command line.

3. Start your VNC client software, and give the following information:

- Hostname: `localhost`
- Port: `5901`

Your VNC desktop session appears.

When you are done with the session, you can close the VNC client application on your desktop.

 **Important:** Keep in mind that closing the VNC client does not close applications that are currently running. You can connect to this session again at any time later and your complete session is saved.

For security purposes, we recommend that you do not keep any windows open which contain root shells or other software running as the `root` user. This will mitigate the risk of someone being able to access your VNC session without you knowing it.

### Transfer files in and out of the system

You can use WebDAV, SFTP or Windows RDP to transfer files in and out of a TeamForge Lab Management system.

#### Transfer files with WebDAV

The most useful way to get files in and out of the TeamForge Lab Management environment is via WebDAV.

Connect to WebDAV with Windows Explorer on Windows XP

With a software update, Windows can be used to connect to a DAV share on a TeamForge Lab Management system.

Microsoft recommends you install this [Software Update for Web Folders](#) before using Web Folders. You may need this update for Web Folders to work properly.

1. Launch Windows Explorer.
2. On the **Tools** menu, click **Folder Options** and select the **View** tab. Clear **Hide extensions for known file types**. This will make file names that begin with a dot (".") to appear properly (for example, your `.bashrc` file).
3. Launch Internet Explorer.
4. On the **File** menu, click **Open**.
5. Enter the URL for your home directory, which will always end in `/home/USERNAME`.

For example:

```
https://$external_host/home/user
```

6. Select **Open as Web Folder**.
7. Enter your user name and password when prompted.

8. If the connection is successful, you will see a new entry added to *My Network Places*. You can re-connect to this folder in the future by going to *My Network Places* and selecting this folder.

As of this writing Microsoft does not provide a client that can mount a Web Folder as a drive letter using an encrypted, authenticated connection, but a third-party tool called Webdrive can do this.

#### Connect to WebDAV with Mozilla Firefox on Windows XP

Firefox isn't a WebDAV client, but it will interface with the Windows native WebDAV client to open up WebDAV shares.

1. Install the [openwebfolder](#) Firefox plug in and restart Firefox.
2. Log into TeamForge Lab Management and go to the **Summary** page of any host you have access to.
3. Under **Access Options**, right-click the link pointing to your home directory and select **Open As Webfolder**.
4. Enter your user name and password.


Your WebDAV home directory is now mounted.

#### Connecting to WebDAV with Mac OS X

To connect to a TeamForge Lab Management system from a Mac, open up your TeamForge Lab Management home directory via DAV access.


This procedure has been tested with Mac OS X 10.4.8 (Tiger) and should work for all Mac OS X versions.

1. In the **Finder**, click **Go > Connect to Server**.

 **Note:** You can also press **cmd-K** on the keyboard.

2. For **Server Address**, enter the URL corresponding to your home directory.  
For example, `https://$external_host/home/user` (where "user" is your user name).
3. Click **Connect**.
4. If prompted, click **Continue**.
5. In the **Verify Certificate** dialog, enter your TeamForge Lab Management user name and password.
6. Select **Remember this password in my keychain** if you want to store your user name and password locally.
7. Click **OK**

You should now see your folder.

 **Note:** Files and directories beginning with a "." are not shown in the graphical folder view, but are accessible via the command line in OS X.

#### Connect to WebDAV with Nautilus on Linux

The Nautilus file browser -- the default file browser on distributions like Red Hat, Fedora and Ubuntu -- has built-in support for DAV folders.

1. From the main desktop, go to **Places > Connect to Server**.
2. For **Service type**, select **Secure WebDAV (HTTPS)**.
3. For **Server**, enter the name of the TeamForge Lab Management server, e.g., `$external_host`
4. For **Folder**, enter `/home/user`, where "user" is your user name.
5. You can leave the **User Name**, **Port**, and **Name to Use For Connection** fields blank.
6. Click **Connect**.

This creates an icon on the desktop, which brings up your home directory when accessed.

#### *Transfer files with SFTP*

Secure File Transfer Protocol (SFTP) can be a useful way to transfer files between your desktop and TeamForge Lab Management.

SFTP is an extension to the Secure Shell (SSH) protocol, which emulates File Transfer Protocol (FTP). File transfers made with SFTP show up immediately in your TeamForge Lab Management environment. Any TeamForge Lab Management client machines you are logged into immediately reflect changes made to your home directory from your desktop.

1. On Linux, UNIX, or Windows with Cygwin, run the command-line SFTP program with a command like this:

```
sftp username@$external_host
```



**Note:** There are also graphical clients available for Linux.

2. On Windows, we recommend using the Filezilla program to manage your SFTP transfers.

Filezilla is available as a plugin for the Firefox web browser with an easy-to-use Explorer-like interface.

### *Connect to TeamForge Lab Management with Windows Remote Desktop*

If you are connecting to a Windows host via Remote Desktop, you can share drives from your local system to the remotely connected system.

1. On the **Start** menu, click **All Programs > Remote Desktop Connection**.
2. Click **Options** and click the **Local Resources tab**.
3. Click **Disk Drives > Connect**

### **Find a build**

You can search for file names (not file content) in the Project Build Library.

Search for files in Project Build Library using these fields:

- Project
- User
- File Name: Shell-like regular expressions are permitted. For example, for searching all RPMS you can type `*.rpm`.
- Description: Substring-like pattern matching is permitted. For example, to search all files added by the `pbl_fsck` script, type `pbl_fsck`.



**Note:** Exclude directories from your search results with the **Exclude** option.

### **Connect to network services outside TeamForge Lab Management**

TeamForge Lab Management systems are firewalled from the outside world, for our users' safety. However, you can use port forwarding to connect to machines outside of TeamForge Lab Management.

For example, you may need to connect to a web server holding the results of your nightly builds, or access an external web server where libraries are installed.

#### *Reach a TeamForge Lab Management server from the outside*

Use port forwarding to allow users outside the TeamForge Lab Management environment to access specific ports on specific hosts inside the TeamForge Lab Management network.

Port forwarding must be set up by a TeamForge Lab Management Domain Administrator: neither you nor your project administrator can do this. This is necessary because port forwarding can open up security holes.



**Caution:** Exposing an insecure network service outside of the TeamForge Lab Management environment could put your whole project and the entire TeamForge Lab Management environment at risk, if that service is compromised.

Contact your Domain Administrator and give them the following information.

- Your name.
- Your project.
- The host you want to expose a port from.
- The port number you would like to expose.
- The service you would like to run on this port.
- Security precautions you have taken to insure that this service is safe to expose outside of TeamForge Lab Management.



**Note:** Your Domain Administrators may have more stringent information gathering or signoff processes which you will be subject to.

### *Reach outside destinations from within TeamForge Lab Management*

TeamForge Lab Management has no built-in features to allow outbound network access to be granted by an administrator. However, if you need access to external destinations from inside TeamForge Lab Management, you have a couple of options.

Try one of these approaches:

- Contact your Domain Administrator and ask CollabNet to open up a hole in the firewall for this purpose.
- Get the files you need into your home directory from outside CollabNet, and access them that way.

### **Allocate a host from a cloud**

To allocate a host from a cloud into a project, specify its type, size and profile information.

In Lab Management, you can allocate from virtual machines on physical hosts and EC2 hosts. You cannot allocate a physical machine from a cloud

1. In your TeamForge Lab Management project, click **Hosts**.
2. Click **Allocate** for the cloud from which you want to allocate your host.


The clouds you can choose from depend on the ones available in the Lab Management domain, and which of these the administrator has made available for your project.

3. In the **Clouds** page, select a host type and size.

Different host types have different attributes such as CPU, disk size and RAM. They also have different costs and available profiles.

4. Select the system profile you want on your host.

A profile in Lab Management is a collection of software packages that comprise a runnable operating system and accompanying set of applications.

 **Note:** Profiles have different costs. For example, a profile might include some licensed software that adds to the cost.

5. Select the number of hosts you want to allocate.

The default is one.


6. To customize your host, click **Advanced options**.

a) Enter a description for the host.

b) Select whether you want the host allocated up to a certain date, or for a specific number of days.

By default, the host is allocated indefinitely until you free it.

7. Click **Allocate Host**.

 **Tip:** You can see information about the host in your **My Pages**.

### **Build and test your software**

When your project manager makes CollabNet Build & Test available, you can set up the CollabNet plugin for Hudson (or Jenkins) in your development environment.

For more information, please see:

- [Hudson wiki for the CollabNet plugin](#)
- [Jenkins wiki for the CollabNet plugin](#)

### **Release products**

Publish the output of your project to selected audiences as packages and releases.

#### **Download a release**

Downloading a release brings all the release's files to your local machine in a single .zip file.

1. Click **FILE RELEASES** from the **Project Home** menu.
2. On the **File Release Summary** page, click the title of the package containing the release you want to download.
3. On the **List Releases** page, click the title of the release.

#### 4. On the **View Release** page, click **Download Selected**.


TeamForge prepares the compressed file in .zip format.

You are prompted to open or save the file.

#### **Create a package**

A package is a folder into which one or more related releases are published.

For example, you might create a package to represent a product deliverable or major component. You can then create releases within the package for product builds or other groups of files.

 **Note:** A package must exist before you create the releases and individual files that will go into the package.

1. Click **FILE RELEASES** from the **Project Home** menu.
2. On the **File Release Summary** page, click **Create**.
3. On the **Create Package** page, provide a title and description for the package.
4. Click **Save**.

The package is created.

When you have created a package, you can publish a release into the package.

#### **Create a release**

A release is a group of one or more files that are published as a unit.

Before creating a release, you must have a package into which your release will go.

1. Click **FILE RELEASES** from the **Project Home** menu.
2. On the **File Release Summary** page, click the title of the package in which you want to create the release.
3. On the **List Releases** page, click **Add**.
4. On the **Create Release** page, provide a name and description for the release.
5. Set the status of the release.

#### **Pending**


Releases with pending status are not visible in the drop-down lists displayed when you set **Reported In Release** and **Fixed In Release** fields in an artifact. Use **Pending** status when you have created a release but have not yet finished adding files.

#### **Active**

Releases with active status are visible in the drop-down lists displayed when you set **Reported In Release** and **Fixed In Release** fields in an artifact.

6. Identify the maturity level of this release.
7. Click **Save**.

The release is created. You can begin adding files.

 **Tip:** To facilitate tracking, you may want to match this release to the planning folder that tracks the work that's going into the release. If you do that, the relevant work items will automatically appear on the **PLANNED TRACKER ARTIFACT** tab. See [Create a planning folder](#) on page 28.

#### **Add files to a release**

After you have created a release, you can add one or more files. All files in a release are published as a unit.

Project members can download the entire release in a single .zip file, or download only selected files.

1. Click **FILE RELEASES** from the **Project Home** menu.
2. On the **File Release Summary** page, click the title of the package containing the release.
3. On the **List Releases** page, click the title of the release.
4. On the **View Release** page, click **Add**.

5. On the **Add File** page, choose the desired file, then click **Save**.  
The file is added and you are returned to the **View Release** page.
6. Repeat the last two steps until all files are added.

After you have added your files, you might wish to change the status of the release from pending to active. This will allow users with appropriate permissions to select the release when entering or updating an artifact. You can also change the maturity level if needed.

#### **Update files in a release**

You can replace an existing file in a release with a new file.

1. Click **FILE RELEASES** from the **Project Home** menu.
2. On the **File Release Summary** page, click the title of the package containing the release.
3. On the **List Releases** page, in the **Releases** section, click the title of the release containing the file that you want to update.
4. On the **View Release** page, select the file you want to update, and click **Update**.
5. In the **Update a File** page, go to the new file with which you want to replace the current file.
6. Click **Save**.

The file is now updated.

#### **Delete files from a release**

If a file is no longer needed, it's a good idea to delete it from the release.

1. Click **FILE RELEASES** from the **Project Home** menu.
2. On the **File Release Summary** page, click the title of the package containing the desired release.
3. On the **List Releases** page, in the **Releases** section, click the title of the release containing the file that you want to delete.
4. On the **View Release** page, choose the file you want to delete, and click **Delete**.

The file is deleted.

#### **Update release attributes**

As your release matures, you should update its maturity level and status.

1. Click **FILE RELEASES** from the **Project Home** menu.
2. On the **File Release Summary** page, click the title of the package containing the release.
3. On the **List Releases** page, in the **Releases** section, click the title of the release containing the file that you want to delete.
4. On the **View Release** page, click **EDIT**.
5. On the **Edit Release** page, choose a new status for the release.  
A release can be in Active or Pending status.
6. Select the maturity level of this release from the **MATURITY** field.
7. Click **Save**.

#### **Change a package's name or description**

When the purpose or the audience for a package shifts, you may want to rename the package or describe it differently.


1. Click **FILE RELEASES** from the **Project Home** menu.
2. On the **File Release Summary** page, select the package you want to edit, and click **Edit**.
3. On the **Edit Package** page, provide the new name or description.
4. To make it easier for users to download the package, select **SHOW DOWNLOAD LINK IN PROJECT LIST**.  
This makes the download link show up in the project's entry in the Project Categories system, if the project has been put in one or more categories. See [Categorize a project](#) on page 303.
5. Click **Save**.

## Associate a release with other items

If a file release is related to other TeamForge items, such as documents, tasks, tracker artifacts, integrated application objects, or news items, you can connect the file release to the other item by creating an association.

Creating associations between items helps you define relationships, track dependencies, and enforce work flow rules. Some example uses of file release associations include:


- Associate a release with a requirements document.
- Associate a beta release with a task related to managing the beta program.
- Associate a release with tracker artifacts representing bugs fixed in the release.

 **Note:** You can also associate tracker artifacts such as bugs and feature requests with the related file releases. You do this as part of working with the tracker.

1. Click **FILE RELEASES** from the **Project Home** menu.
2. On the **File Release Summary** page, click the title of the package containing the release with which you want to create an association.
3. On the **List Releases** page, click the name of the release.
4. On the **View Release** page, click the **ASSOCIATIONS** tab and click **Add**.
5. In the **Add Association Wizard** window, select the items with which you want to associate the artifact:
  - **ENTER ITEM ID** - If you know the item's ID, you can enter it directly.

 **Note:**

- To associate an object in an integrated application from within TeamForge, use the [`<prefix_objectid>`] format. Successful associations appear hyperlinked.
  - Each integrated application displays its prefix on moving the mouse over the application name in the tool bar.
  - **ADD FROM RECENTLY VIEWED** - Select one of the last ten items you looked at during this session.
  - **ADD FROM RECENTLY EDITED** - Select one of the last ten items you changed.
6. Click **Next**.
  7. You may add a comment in the **ASSOCIATION COMMENT** text box.
  8. Save your work.
    - Click **Finish and Add Another** to add additional associations.
    - Click **Finish** to return to the **Details** page.

 **Note:** When an association is added to or removed from TeamForge objects such as tracker artifacts, tasks, documents, discussions, and file releases, a notification mail is sent to users monitoring these objects.

An option is provided at site level and user level to make sure whether the notification mail has to be sent or not. For more information on this, see [Configure your site's settings](#) on page 365.

9. Click the **Associations** tab to view a graphical representation of all the associated items.

Through the Association Viewer, you can choose to view associations in the form of a list or flip over to the Trace view to explore the layers of associations (including parent/child dependencies) laid out in a timeline. You can scroll across the Trace view by dragging the mouse over the association layer or use the 'Previous' and 'Next' arrows to view all the objects as events in a timeline.


While the Associations tab shows the count of the total number of associations, you can only view the most recent 500 associations when you click the Associations tab in case the artifact has more than 500 associations. You can, however, browse through the Association Viewer to view older associations.

## Delete a package

If you no longer need a package, you can delete it.

Deleting a package deletes all of the releases and files within it.



 **Important:** Delete a package only if you are sure that you no longer need any of the releases and files within it.


1. Click **FILE RELEASES** from the **Project Home** menu.
2. On the **File Release Summary** page, choose the package that you want to delete.
3. Click **Delete**.

The package and all of the releases and files within it are now deleted.

### Delete a release

If you no longer need a release, you can delete it.

Deleting a release deletes all of the files within it.

 **Important:** Delete a release only if you are sure that you no longer need any of the files within it.

1. Click **FILE RELEASES** from the **Project Home** menu.
2. On the **File Release Summary** page, click the title of the package containing the release that you want to delete.
3. On the **List Releases** page, in the **Releases** section, select the release that you want to delete, and click **Delete**.

The release and all of the files within it are now deleted.

### Watch for changes

When you monitor an item, such as a task or a document, you are notified by email when that item is updated. You can also monitor a discussion forum or project news by subscribing to its Really Simple Syndication (RSS) feed. In addition, TeamForge Activity Stream provides a feed of recent activities and updates from within the TeamForge project context. Activities include Tracker, Source Code, and external activities via EventQ sources.

### Sort emails from TeamForge automatically

When you are involved in a lot of activities on your TeamForge site, you can start receiving many emails from project members and administrators. To save time, you can configure your email application to sort these emails automatically.

For example, you are automatically emailed whenever someone comments on an artifact that you created. You may also want to join the monitoring lists of items you are interested in, and other team members can add you to the monitoring list of items they create. All of these things can add up to a good number of emails after you have been on a project for a while.

Your email program, such as Microsoft Outlook and Mozilla Thunderbird, looks for the X-headers in incoming email messages and acts on them according to your instructions.

1. In your email client, create a message filter.
  - In Mozilla Thunderbird, click **Tools > Message Filters**.
  - In Microsoft Outlook, click **Tools > Rules and Alerts**.

2. Specify the type of information you want to filter.

For example:

- Use `X-TeamForge-ProjectId` to identify the project that the incoming email is related to.
- Use `X-TeamForge-FolderId` to identify messages related to a specific tracker or document folder.
- Use `X-TeamForge-AssignedTo` to flag tasks or artifacts that are assigned to you.
- Use `X-TeamForge-CommitPath` to watch for commits path in a Subversion repository where you are working.
- Use `X-TeamForge-PlanningFolder` to watch for artifacts to be added to or removed from a planning folder you are working with.

3. Tell your email program what to do with the messages with the X-header that you specified.

For example:

- Automatically move emails into a special "To Do" folder if they are about tracker items assigned to you.
- Color emails red if they are about commits to the Subversion repository path you specified.
- Forward all build-related announcements to your mobile phone for reviewing later.

You can use as many X-headers as you need. If you have set them up judiciously, this can decrease the amount of time you need to work through your incoming email.

### Monitor an item


To get notified by email whenever an item changes, monitor the item.

In TeamForge, an item is something you produce with the site's tools. These are some examples of items:


- Documents
- Tasks
- Tracker artifacts
- Discussion forum topics
- Files in a release

1. Go to the folder containing the item or items that you want to begin monitoring.
2. Select the item or items that you want to begin monitoring.
3. Click the **Monitor** down arrow, then choose **Monitor Selected**.

You are now monitoring all selected items. The icon,  is displayed in the item list view and the monitored items appear in the monitored items list on your personal **Monitoring** page.

 **Note:** When you update two or more artifacts at a time, each user who is monitoring any of the changed artifacts gets a single email describing all the updates.

To stop monitoring an item, select it, then roll your mouse over the **Monitor** down arrow and choose **Stop Monitoring Selected**.

 **Note:** You can also stop monitoring any item from the monitored items list on your personal **Monitoring** page.


### Monitor many items


To get notified about a whole class of items whenever one is created or changed, monitor the folder that contains the items.

In TeamForge, a folder is a container for multiple items. Any container can be considered as a folder, even if it is not explicitly called a folder.


- A document folder contains documents.
- A task folder contains tasks.
- A tracker is a folder that contains tracker artifacts.
- A forum is a folder that contains discussion topics.
- A repository is a folder that contains code commits.
- A package is a folder that contains files.

1. Select the folder that you want to begin monitoring.  
For example, in a project where you are a member, click **SOURCE CODE** from your **Projects Home** menu and select one of the code repositories in the project.
2. Click the **Monitor** down arrow, then choose **Monitor Current Folder**.

You are now monitoring the folder. The icon,  is displayed in the item list view and the monitored items appear in the monitored items list under the **MONITORING** menu available in the **My Workspace** page.

 **Note:** You do not receive monitoring notifications for changes that you yourself make to an item in a monitored folder.

To stop monitoring a folder, click the **Monitor** down arrow, then choose **Stop Monitoring Folder**.

 **Note:** You can also stop monitoring any item from the monitored items list under the **MONITORING** menu available in the **My Workspace** page.

## Monitor an application

To monitor the entire application, such as all trackers or all documents in a project, select it from the **MONITORING** tab available in the **My Workspace** page.

Monitoring an application keeps you updated on all items and folders within the application. Unlike individual items, lists of items, and folders, you don't monitor entire applications from within a project.


1. Click **MONITORING** from the **My Page** menu.
2. From the **Edit Monitoring Subscriptions and Preferences** pane, on the **My Workspace** page, choose the project in which you want to monitor an application.
3. Click the **MONITORED APPLICATIONS** tab and select the applications that you want to begin monitoring.
4. Click **Save**.

You are now monitoring the selected applications.


## Monitor discussion forums as RSS feed

To get an update in your RSS feed reader each time there is an exchange of ideas in a discussion forum, subscribe to the forums's RSS feed.

In TeamForge, you can subscribe to discussion forums via RSS feeds.

1. Click **DISCUSSIONS** from the **Project Home** menu.
2. On the **Forum Summary** page, click the name of the forum that you want to subscribe through RSS feed.
3. On the **Topic Summary** page, click the RSS feed icon .

You can now monitor all topics in the selected discussion forum using your RSS feed reader.


 **Note:** You can access the project messages via the RSS reader without logging in to TeamForge, so long as you have the discussion view permissions.

## Monitor project news as RSS feed

To get an update in your RSS feed reader each time something important is announced in your project, subscribe to the project news as RSS feeds.

In TeamForge, you can subscribe to project news for all the projects via RSS feeds but not to the news of any specific project.

On the **NEWS** tab of **My Page**, click the RSS feed icon. 

 **Note:** You must be a member of a project to view its news.

You can now monitor all project news announcements regarding your projects using your RSS feed reader.

 **Note:** You can access the project news via the RSS reader without logging into TeamForge, so long as you have project membership in any project.

## What am I monitoring?

All of your monitored items appear under the **MONITORING** tab available in the **My Workspace** page.

From this list, you can view or stop monitoring any item you are currently monitoring. You can also monitor entire applications from this page.

1. Select **MONITORING** from the **My Page** menu.  
Your personal monitoring page lists all items you are currently monitoring.
2. Specify the filter criteria in one or more filter fields (at the top of each column) and click **FILTER**.
  - You can find a filter field at the top of each column in most of the tables in the TeamForge application.
  - The filter field could be a text box or a drop-down list with multi-select check boxes.


Stories Page 1 of 40 (600 Items) FILTER  COLUM

PRIORITY <input type="text"/>	ARTIFACT ID : TITLE <input type="text"/>	ASSIGNED TO <input type="text"/>	TEAM <input type="text"/>	SUBMITTED BY <input type="text"/>
3 MEDIUM	artf284437 : Create UI for viewing all diffs for entire review on one screen	C. Michael Pilato	SCM	Mark Phippard
3 MEDIUM	artf277197 : Workspace widget for artifacts	Baybora Aksoy	SCM	Mark Phippard
4 LOW	artf246564 : Reports - write protractor tests to cover e2e tests.	None	Reports (Datamart)	Muhammed Ashik
2 HIGH	artf287743 : Include 'Artifact Tags' in Project Templates	Ramkumar Rathinasamy Viveganandan	TF Core > Avengers - Tags	Aishwarya Rajan
2 HIGH	artf287742 : Tags in Edit Artifact Page	Ramkumar Rathinasamy Viveganandan	TF Core > Avengers - Tags	Aishwarya Rajan
2 HIGH	artf283822 : Synchronize SVN Indexer execution	akil roy	Devops/Releng	Johannes Passing
2 HIGH	artf278265 : Cleanup the old EventService	Nithya Pattabiraman	TF Core > Jaguars - Architecture	Sriram
3 MEDIUM	artf288250 : EventQ meta RPM	Sinhuja Harikrishnan	Devops/Releng	Sinhuja Harikrishnan
3 MEDIUM	artf287737 : Update the default/trial license to include DevOps (and perhaps new license types)	Gowthaman Ravindran	Mavericks - CRUS	Mark Phippard
3 MEDIUM	artf287193 : SVN load balancing	akil roy	Devops/Releng	Johannes Passing
3 MEDIUM	artf286210 : Monit Snapshot before restart	h.karthik	Devops/Releng	Johannes Passing

1 2 3 4 5 6 7 8 9 10 >

Monitor Import Export Cut Edit Edit Inline Mass Update

- You can type your filter criteria in the text boxes. The search text is case-insensitive.
- You can also select the filter values from one or more drop-down lists. By default, you can only select up to 10 filter values in a drop-down list. However, you can set a value that suits your requirement for the `FILTER_DROPDOWN_MAX_SELECTION` token in the `site-options.conf` file to increase or decrease the count.
- Filter-as-you-type:** You can find the **Enter keywords** text box in all filter drop-down lists. As you type your filter keyword, instant search results are shown in the drop-down list. For example, in the following illustration, typing "R" instantly shows all statuses having the alphabet "R". The search text is case-insensitive.



- Some search filters may not appear if your site administrator has not enabled them.
- After filtering, if you want to clear the filters, click **FILTER** and select **Clear** from the drop-down list.
  - Use the up-down arrow at the top of any column to sort your list by that column.
    - Your primary sort column is identified by a superscript **1** next to the up-down arrow, and your secondary and third-level sort columns, if any, are likewise marked.
    - Click the up-down arrow again to reverse the sort order.

To stop monitoring an item from your personal monitoring page, select the item you want to stop monitoring, then click **Stop Monitoring**.

### Who is monitoring this item?

To see who is monitoring an item or folder, check the **Users Monitoring This Item** list.

- Go to the page where the item appears.
- Click the item to view the artifact.
- On the item's **View Artifact** page, click **Users Monitoring**.


The **Users Monitoring This Item** window displays a list of all users who are monitoring the item.

### Choose monitoring email frequency

How often do you want to receive monitoring email notifications for all the applications, folders and items you are monitoring?

1. Select **MY SETTINGS** from your **My Page** menu.
2. On your **User Details** section, in the **USER PREFERENCES** tab, choose an email notification preference.
  - **Email Per Change:** Get a separate email notification for each change to a monitored item.
  - **Daily Digest Email:** Get one email notification each day containing a digest of all changes made to monitored items in the preceding twenty-four hours.
  - **Don't Send Email:** Get no email notifications for changes to monitored items. This can be handy when you are on vacation.
3. Click **Save**.

After you have set your global email frequency, you can further customize the frequency of application monitoring emails.

 **Note:** When you update two or more artifacts at a time, each user who is monitoring any of the changed artifacts gets a single email describing all the updates.

### Choose how often to get emails from specific applications


How many email notifications do you want to get about the applications you are monitoring, and how often?

To further personalize your monitoring preferences, you can set the frequency of email notifications for each monitored application on a project level. If you don't set your application monitoring email frequency settings, your global setting applies.

For example, suppose you are contributing code to the "Widgets" project, but your role in the "Gizmos" project is of a more advisory nature. When you monitor an item in the "Widgets" project you'll want more detailed updates than you will want from items you've monitored in the "Gizmos" project.

1. Set the global default email frequency from your **MY SETTINGS** page.  
See [Choose monitoring email frequency](#) on page 133.
2. Click **MONITORING** from the **My Page** menu.
3. From the **Edit Monitoring Subscriptions and Preferences** pane, choose the project in which you want to configure monitoring email frequency.
4. On the **EMAIL NOTIFICATION PREFERENCES** tab, specify how often you want to be notified and click **Save**.

Any email notification preferences you set here will override the default preferences that you set on your **MY SETTINGS** page.

 **Note:** When you update two or more artifacts at a time, each user who is monitoring any of the changed artifacts gets a single email describing all the updates.

### Monitor an item for someone else


To alert another user to an item, add that user to the list of people monitoring that item.

After a user is added to a monitored item, the user can configure their own monitoring preferences for the item, or stop monitoring the item.

1. Go to the item to which you want to add a user.
2. Click the item to view the artifact.
3. On the item's **View Artifact** page, click **Users Monitoring**.
4. In the **Users Monitoring This Item** window, click **Add**.
5. From the list of available project members, select the user or users that you want to add to the monitored item.

 **Tip:** Press and hold the **Ctrl** key to select more than one user.

## 6. Click **Add**.

 **Tip:** You can also click **Add All** to select all users.

## 7. Click **OK**.

The users are now added to the monitored item.

### Activity Stream

The Activity Stream provides a feed of recent activities and updates from within the TeamForge project context. Activities include Tracker, Source Code, and external activities via EventQ sources.

### Overview

The Activity Stream is a feed of recent project activities, listed in reverse chronological order (newest activities on top). There is one activity stream per TeamForge project. The Activity Stream is useful to keep apprised of recent project updates without switching your context.


The Activity Stream includes events from:

- Tracker
- Source Code (Git and Subversion)
- Pull Request and Gerrit code reviews
- All configured EventQ activity sources (for example, Jenkins, JIRA<sup>®</sup>, Chef, Nexus, Artifactory, Testlink, Reviewboard, and so on)

### Usage

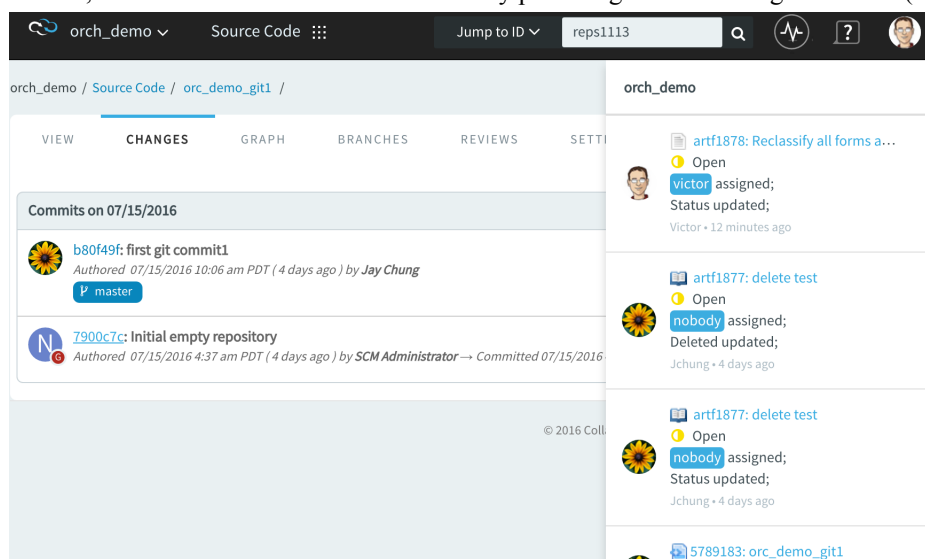
The Activity Stream is activated by clicking the “Activity Stream” icon in the TeamForge header (only available in a project context).




 **Important:** In a distributed TeamForge setup where TeamForge and EventQ run on separate servers, you must make sure that there are no time-synchronization issues. If you see the following error message when you click the "Activity Stream" icon, check if the servers are in sync with network's time.

Please contact your system Administrator. Something went wrong.

When activated, the Activity Stream expands and lays over the page content, anchored to the right side of the browser. Activities appear in reverse chronological order (newest at the top) and are scrollable. As you scroll near the bottom, additional activities load automatically providing more scrolling real estate (and so on).



Each activity represents a distinct event in TeamForge and integrated tools. When available, you can click links and they will open the object in question in a new tab. You will see a notice in your Activity Stream when new activities occur. Click the notice to new activities.



Click the Activity Stream icon in the header a second time to collapse the Activity Stream.

### Permissions

Users must have the EventQ PIPELINE READ permission to access the Activity Stream.

## Lead a project team with CollabNet TeamForge

Managing a project means providing tools and processes to give each project member's contributions their greatest possible impact.

### Create a TeamForge project


Create a new project when you have identified work to be done that has its own distinct character, dependencies or schedule.

What constitutes a project depends on your organization. Some organizations favor a small number of big, centralized projects. Others prefer a larger number of smaller, specialized projects. Your site administrator can help you decide if your work should be part of a larger project or a project of its own.

1. Access the **Projects** page through either of the following ways:
  - Go to **My Workspace > View All Projects**
  - Go to **My Page > Projects**.
2. Click **Create New Project**.
3. On the **Create Project** page, give the project a name and a brief description.
  - The name will appear in project lists and on the project's home page.
  - A terse description is recommended. There will be unlimited room to discuss the project's aims and methods in detail on the project pages themselves.
4. Provide a URL name for the project, if you want the URL for the project to be different from the internal project name.
 

If you do not enter a URL name, the project URL will be the same as the project name.
5. If your site administrator has provided project templates, select the appropriate one for your new project.
 

Project templates give you ready-made artifact types, work flow support, user roles and other start-up content appropriate to the kind of project you are creating.

 **Tip:** The **DEFAULT PROJECT ACCESS** and **PROJECT ACCESS EDITABLE** options are displayed based on the project settings. If you would like to change these settings, ask your site administrator.

6. Click **Create**.

The project is submitted to the TeamForge site administrator for approval. You will receive an email notification when the site administrator approves or rejects your project. When your project is approved, you can get to it from your **MY PROJECTS** tab available under **PROJECTS** menu in the **My Workspace** page or from the **Projects** menu in your navigation bar.

### Build your project web site

Project members and other users need to know how to interact with your project. You can give them the information and tools they need by creating and maintaining a project web site tailored to them.

You can build your project home page from scratch or assemble it from building blocks provided by TeamForge.

### Assemble your project web site


For quick, flexible site building, use the ready-made web page components that come with your TeamForge project site.

There are components for showing text, wiki pages, charts and graphs, and other purposes. These components make it easy to put together a sophisticated project site in little time.

#### *Create a project page*

To provide information and functionality to people viewing your project, build one or more project pages.

1. Go to the page to which your new page will belong.

 **Tip:** Any project page can have subpages belonging to it. A page that belongs directly to the project home page is called a "top-level page."


2. Click **Configure: On**.

3. Choose where your new page will fit in your project's structure.

- To create a page just under the project home page, click **Add Top-level Page**.

A top-level page's title is always visible in the navigation tree at left.


- To create a page under the page you are on right now, click **Add Sub-page**.

 **Tip:** You can create a free-form page only as a subpage.

4. Give your new project page a title.


Keep the title brief and descriptive.


5. Choose who can see this page.

 **Note:** Your choice will apply to all subpages that you create under this page.

- To show this page to anyone with the necessary permissions, select **Visible**. For example, if you have defined a group of users who have access to your project, your new project page is visible only to those users. If your project is open to the public, anyone in the world can see it. Use this option when the information on this page is ready for a wide audience.
- To show this page only to users with the project administrator role, select **Hidden**. Use this option if you are drafting content that you aren't ready to share yet, or want to share only with other project managers.

6. Select the **Freeform** check box to create a free-form page.

 **Note:** A free-form page's title is not visible in the navigation tree at left.

 **Tip:** You can not create free-form subpages when the top-level page is 'Hidden'.

7. Click **Save**.

8. Click **Configure: Off**.

Now you are ready to build the functionality into your project page with components such as text, news or tracker queries.

#### *Create a project page component*

Put information and resources in your users' hands with project page components.



For example, to let people know about important new developments, create a news component. To enable project members to find tracker items quickly, create a query component.

1. On the project page, click **Configure: On**.
2. Click **Add New Component**.
3. On the **Create Component** page, give your new component a title.  
Keep the title brief and descriptive.
4. Select one of the following component types that suits your need.

<b>Text</b>	Write free-form messages, reports or rants, in plain text or HTML.
<b>Reports</b>	Add various reports and charts to your project page.
<b>Documents</b>	Let project members exchange and review documents from the project page.
<b>Wiki Page</b>	Open up your project page to two-way communication.
<b>Tracker Search Results</b>	Make saved search results available from the project page.
<b>Tracker Metrics</b>	Add charts about tracker metrics to your project page.
<b>Project News</b>	Maintain a journal or blog about your project, share information and make announcements.
<b>Project Statistics</b>	Show visual measures of your project activities on the project page.
<b>Subprojects</b>	Add a list of subprojects to your project page.

5. Choose who can see this component.
  - To show this component to anyone with the necessary permissions, select **Visible**.  
For example, if you have defined a group of users who have access to trackers in your project, a query component will be visible only to those users. If your project's trackers are open to anyone, all users who view this project page will see the query component. Use this option when you are sure the component is ready for general use.
  - To keep this component under wraps until you are ready to show it, select **Hidden**.  
Now only users with the project administrator role can see this component. Use this option if you are drafting content that you aren't ready to share yet, or want to share only with other project managers.
6. Select one of the locations, Top of page or Bottom of page, where the component shows up on the project page.
7. Depending on the component type you selected, set the properties of the component.

<b>Component type</b>	<b>In the Properties of this component section...</b>
<b>Text</b>	<ol style="list-style-type: none"> <li>1. Type your free-form messages, rants or announcements in the text box.</li> <li>2. Click <b>Save</b>.</li> </ol>
<b>Reports</b>	For more information about adding reports, refer to the topics: <a href="#">Add reports to a project page</a> on page 138 and <a href="#">Reporting in TeamForge</a> on page 202.
<b>Documents</b>	<ol style="list-style-type: none"> <li>1. Select a folder from the list to display its contents on the project page.</li> <li>2. Click <b>Save</b>.</li> </ol>
<b>Wiki Page</b>	<ol style="list-style-type: none"> <li>1. Type a title for the wiki page.</li> <li>2. Click <b>Save</b>.</li> </ol>

<b>Component type</b>	<b>In the Properties of this component section...</b>
<b>Tracker Search Results</b>	<p>Make sure one or more shared tracker searches are available to add to the project page. For more information about sharing saved tracker searches, see <a href="#">Share a saved tracker search</a>.</p> <ol style="list-style-type: none"> <li>1. Click <b>Add Saved Tracker Searches</b>.</li> <li>2. Select one or more shared tracker searches from the <b>Select from Shared Tracker Searches</b> window.</li> <li>3. Click <b>Add Selected</b>.</li> <li>4. Select the number of rows of the search results to display from the <b>Display Rows</b> drop-down list.</li> <li>5. Click <b>Save</b>.</li> </ol>
<b>Tracker Metrics</b>	<ol style="list-style-type: none"> <li>1. Select the number of charts from the drop-down list. You can add up to three tracker charts to your project page.</li> <li>2. Select one of the chart types: Burndown, Open by Priority or Open vs Closed.</li> <li>3. Select a data source for your chart, a tracker or a planning folder.</li> <li>4. Click <b>Save</b>.</li> </ol>
<b>Subprojects</b>	<ol style="list-style-type: none"> <li>1. Select the number of subprojects to be displayed on the project page.</li> <li>2. Click <b>Save</b>.</li> </ol>

#### 8. Click **Configure: Off**.

##### *Add reports to a project page*

Add one or more reports to your project home page to publish your project status to other project members. You can add only reports of type 'Public' to your project home page.

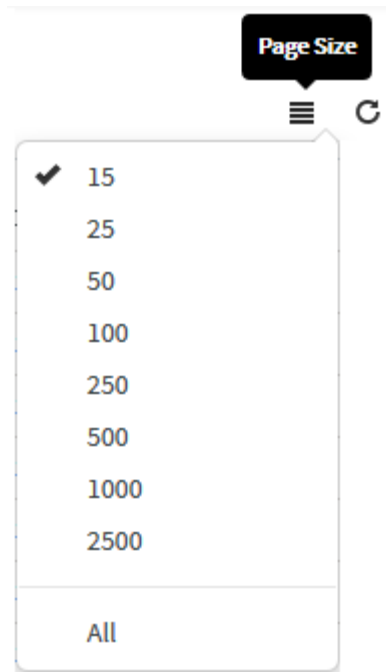
For more information about TeamForge reports, 'Public' and 'Private' reports, see [Reporting in TeamForge](#) on page 202.

1. On the project page, click **Configure: On**.
2. Click **Add new component**.  
The **Create Component** page appears.
3. Type a title for the report component.
4. Select Reports component type.  
A list of 'Public' reports for the project in context is shown.
5. Select **Visibility** and **Location**.
6. Select one or more 'Public' reports to add to the project page.




#### **Note:**

- You can add up to three reports per Reports component.
- The 'Public' reports list may span over multiple pages if the number of reports exceeds **Page Size**. You can navigate through pages by clicking the page numbers at the bottom of the list.
- However, you can change the **Page Size**:



7. If you would like to add reports from other project(s) of which you are a member, select the required project from the **Select Public Reports** drop-down list.



The list of reports existing in the selected project is displayed from which you can select the relevant reports.

 **Note:**



- Private reports and Table reports are not displayed in this list.
- You can add a maximum of three reports per Reports component.

8. Click **Save**.

The selected reports are added to the project page.

*Reorder project page components*


To make things easy for your project members, lay out your page components in a useful order.

1. On the project page, click **Configure: On**.
2. In the title bar of the project page component you want to move, click the up arrow  or the down arrow .
3. Click **Configure: Off**.

*Hide a project page component*

While you are preparing a project page component, you may wish to keep it from being seen.


You can hide a page component from everyone except other project administrators.

1. On the project page, click **Configure: On**.
2. In the title bar of the project page component you want to hide, click the edit icon .
3. Select the **Hidden** option for **VISIBILITY**, and click **Save**.
4. Click **Configure: Off**.

Now no one who does not have the Project Admin role can see the component, even if they have access to the project and to this page. Users with the Project Admin role can see this component only when **Configure** is set to **On**.


#### *Rename a project page component*

As a page component's focus evolves, it's a good idea to rename it to match its changing function.

1. On the project page, click **Configure: On**.
2. In the title bar of the project page component you want to rename, click the edit button .
3. Update the title as appropriate to its current function, and click **Save**.
4. Click **Configure: Off**.

#### *Edit a project page component*

You may want to make some changes to your project page component to make it more useful.

1. On the project page, click **Configure: On**.
2. In the title bar of the project page component you want to edit, click the edit icon .
3. Update the title as appropriate to its current function, if required.
4. Select the **Hidden** option for **VISIBILITY** to hide the page while you work on it.

You can hide a page component from everyone except other project administrators.

5. Change the following settings depending upon the project page component that you are editing:
  - For a **Sub-projects** component, you can change the number of projects to be displayed.
  - For a **Documents** component, you can change the folder location.
6. Click **Configure: Off**.

The project page component is displayed with modified content and/or settings.


#### *Delete a project page component*

When a component is no longer useful, remove it from the page to avoid distracting users.


1. On the project page, click **Configure: On**.
2. In the title bar of the project page component you want to delete, click the **Delete this Page** button.
3. Click **Configure: Off**.

#### *Control access to a project page*

Before you put information or functionality on your project page, make sure it is accessible to the people it is intended for.


 **Note:** This is only relevant if your page is not hidden. If your page is hidden, users who are not project administrators cannot see the page even if their role-based permissions would allow it.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Settings** page, click **Permissions**.
3. Click the role to which you want to give access to your project page.

 **Tip:** If the appropriate role does not exist, you must create it first.

4. On the **Edit Role** page, click **Project Pages**.
5. Under **Project Pages Permissions**, select the pages that users with this role can see and edit.
  - To enable users with this role to create, read and modify all project pages, select the **Project Pages Admin** permission.
  - To allow users with this role to see pages but not edit them, select the appropriate page in the **View** section.


- To allow users with this role to modify the contents of a text component project pages, select the appropriate pages in the **Edit Text Content** section.

 **Note:** The project home page is always visible to any user who is authorized to see the project.

### *Reorder project pages*

Facilitate your users' experience in your project by putting your project pages in an order that matches their needs.

1. Click **Configure: On**.
2. Select a project page and click **Edit Structure**.
3. Rearrange the structure of this page in any of these ways:
  - Select a page, and use **Cut** and **Paste** to place it in a different location.

 **Note:** When you move a page, all of its sub-pages will also be moved.

- Drag and drop a page to a different position.
  - Click **Add New Page** to add a sub-page to the current page.
  - Select one or more pages and click **Delete** to remove them.
4. Click **Save Changes**.
  5. Click **Configure: Off**.

### *Hide a project page*

While you are preparing a project page, you may wish to keep it from being seen.

You can hide a page from everyone except other project administrators.

1. Click **Configure: On**.
2. Click **Edit Page Properties**.
3. For **Visibility**, select **Hidden** and click **Save**.

Now no one who does not have the Project Admin role can see the page, even if they have access to the project and to this page's parent page.

When you are finished building your page, don't forget to switch the page to "Visible" so that the intended users can see it.

### *Rename a project page*

As a page's focus evolves, it's a good idea to rename it to match its changing function.

1. Click **Configure: On**.
2. Click **Edit Page Properties**.
3. Change the **Title** as appropriate to its current function, and click **Save**.
4. Click **Configure: Off**.

### **Hand-code your project web site**

To fully custom control your project web site pages, build your own HTML and check it into the project's Publishing repository.


You can link to these pages as you would to any normal HTML page.

### *Create a custom project home page*


If you decide to hand-code your project home page, the page works like any HTML page.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Settings** page, select **Show custom web page** from **PROJECT HOME OPTIONS**.
  - **Show default project pages and components** is selected by default.
  - Site administrators can restrict public access to Publishing Repositories at the site level and you cannot view or modify **PROJECT HOME OPTIONS** in such case. See [DISABLE\\_REMOTE\\_PUBLISHING](#) on page 740 for more information.

3. Add the `index.html` file to the `publishing/www` Subversion repository. The content of your `index.html` shows up immediately.

 **Note:** When you switch to a hand-coded home page, you can still access any project pages you have created.

4. Click **Save**.  
Your custom **Project Home** page is now active.

 **Note:** If your project has any subprojects, the subprojects are also listed. To see them, make sure your `index.html` file exists and is in the right repository location.

See *what's in your publishing repository*

You can view the files in your TeamForge site's publishing repository through relative URLs.

For example:

1. Enter the URL: `https://forge.collab.net/sf/projects/sampleproject/index.html` in the address bar to display contents of the `index.html` file.

In this case, your `index.html` file must be in the publishing repository under the `www/sampleproject` path.

The contents of the `index.html` file are displayed in the **Project Home** page.

2. Enter the URL: `https://forge.collab.net/sf/projects/sampleproject/roles/roledetails.html` in the address bar to display contents of the `roledetails.html` file.

In this case, your `roledetails.html` file must be in the publishing repository under the `www/sampleproject/roles` path.

The contents of the `roledetails.html` file are displayed in the **Project Home** page.

## Set up trackers

A tracker is a collection of related artifacts that describe work to be done or issues to be resolved. Every project should have one or more trackers.


When you start a tracker, you decide which fields will be used, who will use them, and how they will use them.

### Create a tracker

Create a tracker whenever you need to report and track bugs, feature requests, support requests, or any other type of issue where ownership, status, and activity must be managed.

Individual tracker entries are referred to as tracker artifacts. A tracker is a set of tracker items with a common purpose, such as bug reports, feature requests, or tasks.


1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Tracker Settings**.
3. Click **Create**.
4. On the **Create Tracker** page, provide a name and description for the tracker.

 **Tip:** Descriptions help users learn how best to provide the information you want from them. To maximize your chances of getting useful data, make your description as informative as you can.

5. Select an icon that suggests the type of work the tracker is handling.

This icon will appear with any artifact in this tracker, wherever it is viewed on the site. For example, if someone brings an artifact from this tracker into a planning folder, users of the planning folder can glance at the artifact's icon to see where it comes from.

6. Select the relevant unit from the **DISPLAY EFFORT IN** field. The units displayed here are configured based on the size of the artifacts in the tracker. Eg. Select the unit as **HOURS** for a tracker of small defects, **DAYS** for a tracker of tasks, and **WEEKS** for a tracker of epics.

 **Note:** Configure the units at the project level and not at the planning folder level.

7. Select **INCLUDE FOREIGN CHILDREN** to include points and efforts from children artifacts across the projects in TeamForge.



**Note:** In a parent artifact, enabling **CALCULATE POINTS** field sums and rolls up the points from all its children artifacts within the project. In this total, if you want to include children artifacts from other projects across TeamForge, have the **INCLUDE FOREIGN CHILDREN** option enabled.

**8. Click **Create**.**

The new tracker appears at the bottom of your list of trackers.

**9. If necessary, drag the tracker to a place in your tracker list that makes sense.**

The order you set here controls the order of every tracker list in your project.

**10. You'll probably need some custom fields to capture information that's specific to your project.**

See [Create your own tracker fields](#) on page 54.

**11. To speed up the team's work, you may want to set up some rules for automatically reassigning artifacts when their contents change.**

See [Create a tracker workflow](#) on page 55.

### Create your own tracker fields

To track data that is not captured by the default set of fields, create new fields that fit your project's purposes.

You can create the following user-defined fields in each tracker:

- Up to 30 text entry fields.
- Up to 30 date fields.
- Up to 30 single-select fields.
- An unlimited number of multiple-select fields.



**Caution:** Creating a large number of user fields and multiple-select fields may affect site's performance.

#### Create a text field

To let users type in data, create a text entry field.

A tracker can have up to 30 text fields.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Tracker Settings**.
3. Click the tracker to which you want to add a text field.
4. On the **TRACKER FIELDS** tab, click **Add Field**.
5. On the **Create Field** page, provide a name for the field.
6. Configure the shape of the field with the **Field Width** and **Field Height** fields.
7. To help users enter the right text values, select **Use Text Validation** and supply a regular expression that describes the appropriate values.

This can help reduce errors and keep your team's data as meaningful as it can be. For more detailed instructions, see [Validate text entries in a tracker artifact](#) on page 151.

**8. Click **Save Field**.**

The new field is created.

#### Create a "Select" field

To let users choose values from a list that you define, create a "Select" field.

You can create up to 30 single-select fields and an unlimited number of multiple-select fields in a tracker.



**Note:** Creating a large number of multiple-select and user fields may affect TeamForge's performance.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Tracker Settings**.
3. From the list of existing trackers, click the tracker to which you want to add a "Select" field.
4. On the **TRACKER FIELDS** tab, click **Add Field**.
5. On the **Create Field** page, provide a name for the field.

6. Use the **Input Type** menu to specify whether users will be able to select one value or more than one.

If you're going to make this a required field, pick one of the values to be the default value. This value is applied to existing artifacts and artifacts that are moved from another tracker.

7. Decide whether users *must* choose a value.

- Required fields automatically appear on the **Submit Artifact** page.



**Note:** If you make the field required, you must specify a default value.

- If you make a **User** field required, specify one or more default users.
- If you make a **Date** field required, the default is "today."
- For optional fields, select **DISPLAY ON SUBMIT** if you want the field to appear when a user first creates an artifact.
- To prevent the field from being used at all, select **DISABLED**. (By default, new fields are enabled.)

8. Use the **Values** section of the **Create Field** page to add more values for the user to choose from.

9. Keep adding values until you have the list of options you want, then click **Save Field**.

### *Create a people-picker field*

To let users choose other users from a list, create a "people picker" field.

The **Assigned to** field is a people-picker field that is present in every tracker. An artifact can be assigned to only one user at a time. You can give yourself more flexibility by adding any number of custom people-picker fields.

For example, your QA team may want to assign a person to track progress on an artifact while it is assigned to a developer. You might create a people-picker field called "QA monitor" to specify who that person should be.

In people-picker fields you create, users can select multiple users.



**Note:** Creating a huge number of user fields can slow down the site's performance.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Tracker Settings**.
3. From the list of existing trackers, click the tracker to which you want to add a people-picker field.
4. On the **TRACKER FIELDS** tab, click **Add Field**.
5. On the **Create Field** page, provide a name for the field.
6. On the **INPUT TYPE** menu, select **Select User(s)**.
7. In the **DEFAULT FILTER** field, choose whether the list of people available in your new field will include members of this project or everyone who is registered on the site.
8. Configure the size of the field with the **FIELD WIDTH** field.
9. Click **Save Field**.  
The new field is created.

### **Organize tracker fields**

Most tracker artifacts ask the user for a lot of information. You can arrange the input fields in columns and rows to make it easier for users to find the fields they need.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Tracker Settings**.
3. From the list of current trackers, click the tracker whose fields you want to organize.  
Click the **TRACKER FIELDS** tab if it isn't already showing.
4. If some fields seem to be logically connected to each other, create a section to bring them together.
  - a) Click **Add Separator** and select **Section**.  
Give the section a short but descriptive label.
  - b) In the list of fields, drag your new "Section Separator" row to a position that makes sense.
  - c) Drag the appropriate fields under the Section Separator that you just created.
5. Within a section, arrange fields logically into columns.



- a) Click **Add Separator** and select **Column**.

Give the column a short but descriptive label.

- b) In the list of fields, drag your new "Column Separator" to a position that makes sense.
- c) Drag the appropriate fields under the Column Separator that you just created.
- d) Create as many columns as you need.


Drag a column separator above another column separator to move it to the left in the artifact entry form. Drag it below to move it to the right.

6. Within a column, group fields into rows if appropriate.

- a) Click **Add Separator** and select **Row**.

Give the row a short but descriptive label.

- b) In the list of fields, drag your new "Row Separator" to a position that makes sense, then drag the appropriate fields under the Row Separator.


 **Note:** You can have rows and columns without sections.

### Enable or disable fields


If a tracker field is disabled, it does not appear on the **Artifact** page. Most fields can be disabled.

Disabled fields are accessible only to tracker administrators on the **TRACKER FIELDS** tab.

You can enable or disable any field that is user-defined or configurable. However, you can't disable all configurable fields. For example, the **Status**, **Priority**, **Category**, and **Planning Folder** fields can't be disabled.

 **Note:** If your goal is to prevent users from entering data into a field when submitting an artifact, but still display the field on the **Edit Artifact** and **Tracker Search** pages, don't disable the field. Instead, clear the **DISPLAY ON SUBMIT** option on the **Edit Tracker Field** page.


1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Tracker Settings**.
3. From the list of current trackers, click the tracker you want to configure.
4. On the **TRACKER FIELDS** tab, select the fields you want to enable or disable.
  - Click **Disable** to remove them from the **Artifact** page.
  - Click **Enable** to allow them to be configured and displayed.

 **Note:** Data in disabled fields is still search-able, but disabled fields do not appear as inputs on the **Search** pages.

### Require users to complete a tracker artifact field

If a field is set as required, users cannot submit artifacts without completing it. Most tracker fields can be required or optional.


Each tracker can have its own required and optional fields. Required fields are marked with a red asterisk \* on the **Submit Artifact** page.

 **Note:** When you make a field required, any field whose values depend on that field's values is also required. See [Help users select options in a tracker artifact](#) on page 149 for more information about dependent field values.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Tracker Settings**.
3. From the list of current trackers, click the one you want to configure.
4. On the **Tracker Fields** tab, click the name of the field you want to set as required or optional.
 

By default, only the **Title**, **Description**, and **Status** fields are set as required.
5. On the **Edit Field** page, select or clear the **Required** option to make a field required or optional.
 

Required fields automatically show up on the **Submit Artifact** page.

 **Note:** System-defined fields and the **Status** field are always required.

6. For optional fields, select or clear the **DISPLAY ON SUBMIT** option.

This specifies whether the field will appear on the **Submit Artifact** page.

7. Click **Save Field**.

### Configure tracker "Select" field values

To help users provide meaningful information, supply them with useful field values to choose from in the input fields in the tracker entry form.

 **Tip:** Once a tracker has been created, you may:

- Create one or more user-defined single-select or multiple-select fields.
- Add predefined values to the fields. Remove values, if required.
- Enable or disable fields.
- Change the default values for fields.

1. Click **PROJECT ADMIN** from the **Project Home** menu.

2. Click **Tracker Settings**.

3. From the list of existing trackers, click the name of the tracker that you want to configure.

4. On the **TRACKER FIELDS** tab, click the name of the field whose values you want to edit.

5. On the **Edit Field** page, set up the field values you want users to see when they create a tracker artifact.

- To define a new value, click **Add**.
- To rename a value, edit the existing text. If you rename a value, the value is renamed in all existing artifacts.
- To remove a value, check the box and click **Delete**. If you delete a value, the value is changed to **None** in all existing artifacts.
- Select **DEFAULT VALUE** to set which option will be chosen if the user makes no selection. When you move a tracker artifact from one tracker to another, the default field value is the value that comes along.

 **Note:**

- When you edit the values of the **Status** field, you are also asked to describe what each value's status means, as shown in the **Values** section of the **Edit Field** page. This status meaning is used in **Advanced Search** to define which values are returned when searching for open or closed artifacts.
- As always, when you create a new tracker, the default value for the **'Priority'** field is set as '4 - Low'. However, you can change the default value by editing this configurable single select field, **'Priority'**. You cannot delete or disable the **Priority** tracker field.
- When you change the tracker fields, the values in the existing artifacts remain unchanged.

6. Click **Move Up** or **Move Down** to order the list the way you want it.


7. Click **Save Field**.

All values are now available in the selection menu for the field.

### Configure tracker units

You can estimate the value of efforts meaningfully in the form of units using the TeamForge tracker.

You can set any unit as the default for a tracker or planning folder and create any number of units. The default base unit is 'Hours' which you can rename, but not delete. You can also toggle between the burndown chart and effort values for any unit.

 **Note:** You cannot enter decimal values in the **Conversion** field. Deleting a unit will cause all the associated artifacts to express effort in the form of the base unit.

1. Click **PROJECT ADMIN** from the **Project Home** menu.

2. Click **Tracker Settings**.

3. On the **UNITS** tab, click **Add**.

4. On the **Add Unit** page, enter the **UNIT NAME** and **CONVERSION** value for the unit.

5. Click **Save**.

### Assign a tracker artifact automatically

You can configure the tracker to automatically assign newly submitted artifacts to specific project members.

You can assign artifacts to individuals based on the values in the **Category** or **Release** field.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Tracker Settings**.
3. From the list of existing trackers, click the tracker you want to configure.
4. On the **AUTO ASSIGNMENT** tab, select **Auto Assign By Category** or **Auto Assign By Release**.
5. For each category or release value, choose a project member from the **AUTOMATICALLY ASSIGN TO** drop-down menu.

This menu contains all project members with the tracker edit permission. Click the Search icon to display a list of project members to whom you can auto-assign artifacts. Choose **None** if you do not want artifacts with a specific value automatically assigned.

6. Click **Save**.

Whenever a new artifact is submitted with a value in the relevant field, the assignee receives an email notification and the artifact appears on the assignee's **My page**.

### Create a tracker workflow

To channel project members' work on tracker items, set up rules for how a tracker item can move forward.

Before creating a tracker workflow, see that these criteria are met:

- You have a tracker to work with.
- The tracker has a set of statuses defined, such as "In progress" and "Ready for QA."
- Roles exist, and you can assign project members to them.

A workflow is a sequence of changes from one status to another. You can define status transitions for any combination of tracker statuses in the tracker.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Tracker Settings**.
3. From the list of existing trackers, select a tracker.
4. Click the **WORKFLOW** tab.  
The **Workflow** page lists all of your status values and the tracker workflow that you have configured.
5. On the **Workflow** page, click the status value for which you want to create a workflow.
6. On the **Edit Field Transition** page, select a status value from the **Create Transition to Status** drop-down menu.



**Note:** By default, the **Any** workflow is shown. This means that you are defining workflow rules for the transition from the status that you selected to *any* other status.

7. Click **Add**.  
A new workflow is added. The **Any** workflow is changed to **Remaining Statuses**.
8. In the **ROLES** section, specify which users can make this change.  
For example, only users with the QA Engineer role are allowed to change artifacts from **Open** to **Cannot Reproduce**.
9. Click the **Advanced Transition** link.



**Note:** For any new unsaved transition, an alert is shown asking you to save the transition so that you can view it in the **Transition Status To** drop-down list. Click **OK** to configure already saved transitions or click **Cancel** and then click **Save and View** in the **Edit Field Transition** page to save the unsaved transitions.

### Advanced Transition ✕

Transition Status To Under Consideration -> Completed ▼

NAME	TYPE	REQUIRED	VISIBLE
Group	Single Select	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Status	Single Select	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Category	Single Select	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Customer	Single Select	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Priority	Single Select	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Team	Folder	<input type="checkbox"/>	<input checked="" type="checkbox"/>

Cancel
Save And Close
Save

- a) Select the transition workflow for which you want to apply Advanced Transition settings from the **Transition Status To** drop-down list.
- b) Select the **REQUIRED** check box against the fields for which the user *must* provide values. For example, the user must assign the tracker item to someone and enter a comment.

 **Note:**

Fields whose values depend on a required parent field are automatically required. (See [Help users select options in a tracker artifact](#) on page 149 for more information on parent and child fields.)

- c) Select or unselect the **Visible** check box for showing or hiding fields respectively for the selected status transition.

Points to note:

- **Required** fields are always visible.
- Advanced Transition rules are applied when you create or edit artifacts in a tracker and only when you edit artifacts in Planning, Task, and Kanban boards.
- Hidden field values, if updated via SOAP/REST APIs, are ignored.

**10. Click Save.**



The workflow is now saved. When a user submits or edits the status of a tracker artifact, he or she sees only the options that are allowed by the workflow.

**Set default tracker columns**

Choose what information is visible by default to users of your project when they open a tracker or planning folder.

Users can also set their own personal default column configurations, which override the project defaults.



1. Click **TRACKERS** from the **Project Home** menu.
2. Select a tracker, planning folder or team and click **COLUMNS > Configure**.
  - If you've already saved a column configuration, click it and skip the rest of these steps.

- To go back to the default column configuration, click **System (default)** and skip the rest of these steps.
  - To set up a new configuration, click **Configure**.
3. Choose your columns.
    - a) Move the columns you want from **Available Columns** to **Selected Columns**.  
**Artifact ID : Title, Priority and Status** are required columns.
      -  **Note:** Selecting more columns can increase the time required to load the listing page.
    - b) Remove any columns you don't need from **Selected Columns**.
    - c) Use the move up and move down arrows to change the display order of the columns.
  4. Apply your choices to your view of the tracker.
    - To use this arrangement this time only, click **Apply**. The next time you log in, you'll start with the default view again.
    - To save your column layout for repeated use, click **Apply and Save**, then give your arrangement a name. The next time you log in, you'll see the column arrangement you just selected. (If you've sorted the records in your view, that sort order is saved too.)
      -  **Tip:** If you are editing a column configuration that already exists, you can rename it by saving it under a new name.
  5. To make the same set of columns appear every time you come to this tracker, planning folder or team, click **COLUMNS > Save** and from **Save Column Configuration** page, select **Make this my default view**.
  6. To make the same set of columns appear for every user the first time they see any tracker, planning folder or team in the project, click **COLUMNS > Save** and from **Save Column Configuration** page, select **Make this the default view for all project members**.

The project default configuration you set is now the default configuration for all project members, unless they have created their own personal default column configuration.

### Change a tracker

A tracker's real-world uses often outgrow the name or description you gave it when you created it. When that happens, it's a good idea to update the tracker to reflect its changing role.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Tracker Settings**.
3. From the list of existing trackers, click the tracker you want to edit, and click **EDIT**.
4. On the **Edit Tracker** page, provide a new name or description for the tracker, and update the icon.
5. Update the units from **DISPLAY EFFORT IN** and click **Save**.
  -  **Note:** These units are configured in the **Units** page at the project level, and not at the planning folder level.
6. Select **INCLUDE FOREIGN CHILDREN** to include points and efforts from children artifacts across the projects in TeamForge.
  -  **Note:** In a parent artifact, enabling **CALCULATE POINTS** field sums and rolls up the points from all its children artifacts within the project. In this total, if you want to include children artifacts from other projects across TeamForge, have the **INCLUDE FOREIGN CHILDREN** option enabled.
7. If necessary, drag the tracker to a place in your tracker list that makes sense.  
 The order you set here controls the order of every tracker list in your project.

### Help users select options in a tracker artifact

You can help users cope with complex information by guiding them to eligible values in single-select fields.

Any tracker that manages real-world information will quickly become very complex. Users can be confused by a proliferation of "Select" fields. Confusion can lead to inconsistent data, which makes your job harder.

You can help relieve the complexity by showing users their eligible options in a given field based on values they have already selected in another field. You can create overlapping sequences of dependent fields, with as many levels as you need.

This simplifies things for the user, but for the tracker administrator it can quickly get complicated. So let's look at an example.

1. Click **PROJECT ADMIN** from the **Project Home** menu.

2. Click **Tracker Settings** and create a tracker.

For this example, let's call it the **Lunch Planning** tracker.

3. Create a single-select field and call it **Lunch type**.

(For this example, we'll ignore the built-in fields, such as **Status** and **Assigned to**. We're just working with fields that you create.)

4. Create some values for the **Lunch type** field. Let's call them Buffet, Picnic, and Banquet.

Each type of lunch will make sense in some kinds of locations and not others. For example, you would not normally plan a banquet lunch in a park. We are now going to make it easy for users to avoid making such a mistake.

5. Create a single-select field called **Location type**, with Lunch type as the parent field, and give it some plausible values.

a) Start by adding an option called Beach. In the **Parent Values** column, choose Picnic, because that's the kind of lunch you would have at the beach.

(The **Parent Values** column lists all the values in the parent field you selected.)

b) Add a Restaurant. In the **Parent Values** column, hold down the **Ctrl** key and choose Banquet and Buffet, because either of those could be held at a restaurant.

6. Create a single-select field called **Location**, with Location type as the parent field, and give it some values to choose from.

a) Start with an option called Happy Food Restaurant. In the **Parent Values** column, select Restaurant, because that's the type of location that Happy Food Restaurant is.


b) Add another option, Hanalei Cove. Under **Parent Values**, select Beach.

7. Save your work and go to the tracker whose settings you have been editing.

Try selecting from the interdependent values you have just created.

Observe that the value you choose in the **Lunch Type** field controls which values are available to you in the **Location Type** field, and that selecting a value in **Location type** in turn controls the values that appear in the **Location** field.

- When a user selects Banquet for a lunch type, they can select Restaurant but not Beach in the **Location type** field. You will have less error correction to do, and users will avoid confusion.
- When a user selects Picnic for their lunch type, the **Location Type** field offers only Park and Beach. Now you will not have to go through and clean up after users who mistakenly choose to plan a picnic at Happy Food Restaurant, and the doorman at Happy Food Restaurant will not have to turn away users who mistakenly show up with picnic baskets.

 **Important:** If you are used to defining your own Tracker fields, the ability to make field values depend on other values may change how your trackers work in ways you didn't anticipate. Keep these points in mind:

- Linking fields in this way doesn't modify existing data, but when users later modify fields that are linked, they will have to adhere to the relationships you set here.
- If a field has a parent, and that parent field also has a parent, the top-most parent field must have at least one value.
- When a field has a parent field that is required, the child field's default value is set to None. If that required parent field is deselected, the child field no longer has to be required, but Required remains the default.
- If you require a specific field value before an artifact can be placed in a given status, that field's children are subject to the same requirements. See [Create a tracker workflow](#) on page 55 for more about controlling what status an artifact can be in.


- If you delete a field that contains values that another field's values depend on, the dependent field becomes a standard single-select field on its own.
- When you cut and paste an artifact from one tracker to another, only those field values that also exist in the new tracker come along with the artifact. If those values aren't valid under the dependency rules of the new tracker, they are still brought along.

### Validate text entries in a tracker artifact

You can help users contribute useful information by testing their text entries against rules you configure.

Text fields can be error-prone because they invite free-form input. You can help users provide usable information by automatically rejecting values that don't match the needs of the tracker.

This simplifies things for the user, but for the tracker administrator it can be complicated. So let's look at an example.

 **Note:** If your goal is to require users to enter *some* value, whatever the value is, don't use text field validation. Select the **Required** option instead.

1. Click **PROJECT ADMIN** from the **Project Home** menu.

2. Click **Tracker Settings** and create a tracker.

For this example, let's call it the **Bugs** tracker. It will be used to record entomological specimens in a collection.

3. Create a text field and call it **Legs**.

This is for users to record the number of legs each specimen displays.

4. Select **USE FIELD VALIDATION** and supply a validation rule that requires the user to enter a number.

For example, if a given insect has six legs, you'll want the user to enter the numeral **6**, and not a string such as "six" or "several."

Try this regular expression:

```
\d{1,3}
```

This rule requires the user to enter a number with one, two or three digits. Now, a user who means to record a centipede with 100 legs but enters 1000 by mistake will not be able to save the artifact until the error is corrected.

5. Enter a sample string to test your regular expression.

Any part of the sample string that matches your regular expression appears under **Match Results**. If nothing appears, rework your regular expression until you get a match.

6. Create another text field and call it **Location**.

This is where users will record the geographical spot where they collected the bug.


7. Select **Use Text Validation** and supply a validation rule that requires the user to enter a pair of geographical coordinates.

For example, if a given insect was found outside CollabNet's California headquarters, you'll want the user to enter a string like `37.674689, -122.384652`, and not something like "Brisbane" or "Out on the lawn."

Try this regular expression:

```
[ - ] ? [ 0 - 9 ] * [ . ] { 0 , 1 } [ 0 - 9 ] { 0 , 4 }
```

This rule requires the user to enter two numbers, separated by a comma, in the general format of a pair of mapping coordinates.

 **Note:** This particular regular expression does not guarantee that the coordinates are valid, just that they look like coordinates.

8. Save your work. In the tracker whose settings you have been editing, try entering a number greater than 999 in **Legs**, or a street address in **Location**

The red **X** next to the field indicates that the text entry is incorrect. A green check indicates that the value meets the requirements.

Notice that any field in which you are validating text entries is identified by **Text Entry (with Field Validation)** when listed on the **TRACKER FIELDS** tab.

## Document your work

The document management system is a centralized repository for creating, storing, and managing information about a project.

Project members with the Document Admin permission can create, edit and administer documents and document folders. In addition to the Document Admin permission, individual permissions to create, edit and delete documents, and document folders can also be set so that project members who do not have the Document Admin permission can still perform these tasks.

## Document settings

Managing document settings is possible starting from TeamForge 17.1. At the project level, document administrators can set up document fields, document workflow and mandatory document reviewers.

In addition to the system defined fields Document ID, Title and Description, a new configurable single select field, **Status**, has been added in TeamForge 17.1. In addition to the default values, DRAFT and FINAL, you can now add custom statuses for documents.

You can also select one or more TeamForge users, for example project administrators and scrum masters, as mandatory document reviewers. This ensures that all project documents go through them when document reviews are initiated.

1. To configure document settings for a project, select the project from **My Workspace** menu and select **Project Home > Project Admin > Document Settings**.

NAME	INPUT TYPE / SIZE	ENABLED / DISABLED	FIELD TYPE	REQUIRED	PARENT FIELD
Document ID	Read Only	Enabled	System Defined	Yes	
Title	Free Text / 50	Enabled	System Defined	Yes	
Description	Free Text / 50	Enabled	System Defined	Yes	
Status	Single Select	Enabled	Configurable	Yes	

2. Add one or more default document reviewers. Click the **Default Document Reviewers** field's search icon. Search and select users from the **Find a User** dialog box and click **OK**.

3. Set up custom document statuses. In addition to the default document statuses, DRAFT and FINAL, you can add one or more custom document statuses.
  - a) Select the **DOCUMENT FIELDS** tab and click the **Status** field. The **Edit Field** page appears.



**Edit Field**

FIELD NAME: Status      REQUIRED:  (Will display on Document Submit and Document Edit.)

INPUT TYPE: Single Select (Specify Field Values)      HELP TEXT:

FIELD TYPE: Configurable

DISPLAY ON SUBMIT:

**Values:**  
Rename values with caution. Renaming a value causes all documents with the current value to be changed to the new value. You can only delete values that are not in use by any document(s).

VALUES	STATUS TYPE	DEFAULT VALUE	MOVE UP	MOVE DOWN
<input type="text" value="Draft"/>	Open ▾	<input checked="" type="radio"/>	<a href="#">Move Up</a>	<a href="#">Move Down</a>
<input type="text" value="Final"/>	Close ▾	<input type="radio"/>	<a href="#">Move Up</a>	<a href="#">Move Down</a>

[Alphabetize](#) [Delete](#) [Add](#)

[Cancel](#) [Save Field](#)

- b) Click **Add**. A new row is added in the **Values** section.
  - c) Type the name of the new document status, select a status type (open or close) from the drop-down list and select the **Default Value** option if you want to make this the default status when a document is created. You can move the row up or down by clicking the **Move Up** and **Move Down** links, if required.
  - d) When done, click **Save Field**.  
A confirmation message appears.
  - e) Click **OK** to save the new status.
  - f) Repeat the above steps to add more custom document statuses.
4. Add custom document fields (flex fields), if required. For more information, see [Create your own document fields](#) on page 153.

### *Create your own document fields*

To track data that is not captured by the default set of fields, create new fields that fit your project's purposes.

You can create the following user-defined fields for your documents:

- Up to 30 text entry fields.
- Up to 30 date fields.
- Up to 30 single-select fields.
- Up to 30 multiple-select fields.

#### Create a "Text" field

To let users type in data, create a text entry field.

You can have up to 30 text fields.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Document Settings**.
3. On the **DOCUMENT FIELDS** tab, click **Add Field**.  
The **Create Field** page appears.

4. On the **Create Field** page, provide a name for the field.
5. Configure the shape of the field with the **Field Width** and **Field Height** fields.
6. Select Text Entry from the **Input Type** drop-down list.
7. To help users enter the right text values, select **Use Text Validation** and supply a regular expression that describes the appropriate values.

This can help reduce errors and keep your team's data as meaningful as it can be. For more detailed instructions, see [Validate text field values](#) on page 154.

8. Click **Save Field**.


The new field is created.

#### Validate text field values

You can help users contribute useful information by testing their text entries against rules you configure.

Text fields can be error-prone because they invite free-form input. You can help users provide usable information by automatically rejecting values that don't match the needs of the document.

This simplifies things for the user, but for the document administrator it can be complicated. So let's look at an example.

 **Note:** If your goal is to require users to enter *some* value, whatever the value is, don't use text field validation. Select the **Required** check box instead.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Document Settings**.
3. On the **DOCUMENT FIELDS** tab, click **Add Field**.  
The **Create Field** page appears.

4. On the **Create Field** page, provide a name for the field. For example, name the text field **Legs**.  
This is for users to record the number of legs each specimen displays.
5. Configure the shape of the field with the **Field Width** and **Field Height** fields.
6. Select Text Entry from the **Input Type** drop-down list.

7. Select the **USE FIELD VALIDATION** check box and supply a validation rule that requires the user to enter a number.

For example, if a given insect has six legs, you'll want the user to enter the numeral **6**, and not a string such as "six" or "several."

Try this regular expression:

```
\d{1,3}
```

This rule requires the user to enter a number with one, two or three digits. Now, a user who means to record a centipede with 100 legs but enters 1000 by mistake will not be able to save the document until the error is corrected.

8. Enter a sample string to test your regular expression.  
Any part of the sample string that matches your regular expression appears under **Match Results**. If nothing appears, rework your regular expression until you get a match.

9. Create another text field and call it **Location**.

This is where users will record the geographical spot where they collected the bug.


10. Select the **USE FIELD VALIDATION** check box and supply a validation rule that requires the user to enter a pair of geographical coordinates.

For example, if a given insect was found outside CollabNet's California headquarters, you'll want the user to enter a string like 37.674689, -122.384652, and not something like "Brisbane" or "Out on the lawn."

Try this regular expression:

```
[ - ]? [ 0 - 9 ] * [ . ] { 0 , 1 } [ 0 - 9 ] { 0 , 4 }
```

This rule requires the user to enter two numbers, separated by a comma, in the general format of a pair of mapping coordinates.

 **Note:** This particular regular expression does not guarantee that the coordinates are valid, just that they look like coordinates.

11. Save your work. In the document whose settings you have been editing, try entering a number greater than 999 in **Legs**, or a street address in **Location**  
The red **X** next to the field indicates that the text entry is incorrect. A green check indicates that the value meets the requirements.


Notice that any field in which you are validating text entries is identified by **Text Entry (with Field Validation)** when listed on the **DOCUMENT FIELDS** tab.

Create a "Select" field

To let users choose values from a list that you define, create a "Select" field.

You can create up to 30 single-select and 30 multiple-select fields for documents.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Document Settings**.
3. On the **DOCUMENT FIELDS** tab, click **Add Field**.  
The **Create Field** page appears.
4. On the **Create Field** page, provide a name for the field.
5. Use the **Input Type** menu to specify whether users will be able to select one value or more than one.  
If you're going to make this a required field, pick one of the values to be the default value. This value is applied to existing documents and documents that are moved from another project.
6. Decide whether users *must* choose a value.
  - Required fields automatically appear on the **Submit Artifact** page.

 **Note:** If you make the field required, you must specify a default value.

- If you make a **User** field required, specify one or more default users.
- If you make a **Date** field required, the default is "today".

- For optional fields, select **DISPLAY ON SUBMIT** if you want the field to appear when a user first creates a document.
  - To prevent the field from being used at all, select **DISABLED**. (By default, new fields are enabled.)
7. Use the **Values** section of the **Create Field** page to add more values for the user to choose from.
  8. Keep adding values until you have the list of options you want, then click **Save Field**.

#### Create a "People-picker" field

To let users choose other users from a list, create a "People-picker" field.

For example, you may want to create a "Document Owner" field that can be used to identify the user who owns the document. You might create a people-picker field called "Document Owner" to specify who that person should be.

In people-picker fields you create, users can select multiple users.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Document Settings**.
3. On the **DOCUMENT FIELDS** tab, click **Add Field**.  
The **Create Field** page appears.
4. On the **Create Field** page, provide a name for the field.
5. On the **INPUT TYPE** menu, select **Select User(s)**.
6. In the **DEFAULT FILTER** field, choose whether the list of people available in your new field will include members of this project or everyone who is registered on the site.
7. Configure the size of the field with the **FIELD WIDTH** field.
8. Click **Save Field**.  
The new field is created.

#### Create a document work flow

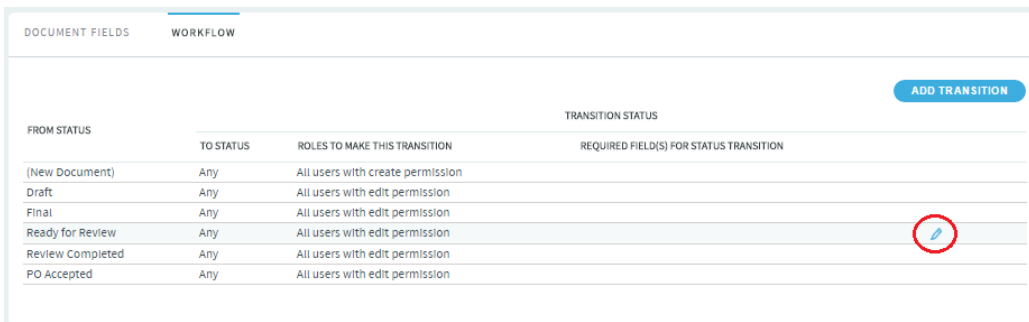
To channel project members' work on documents, set up rules for how a document can move forward.

Before creating a document work flow, see that these criteria are met:

- You have configured a set of statuses, such as "Draft", "Ready for Review", "Review In Progress" and so on.
- Roles exist, and you can assign project members to them.

A work flow is a sequence of changes from one status to another. You can define status transitions for any combination of document statuses in Document Settings.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Document Settings**.
3. Click the **WORKFLOW** tab.  
The **WORKFLOW** tab lists the transition rules for all the document statuses that you have configured already. You can also add new transition rules, if required.
4. Move your mouse over the status rows to view the edit transition icon.



FROM STATUS	TRANSITION STATUS		
	TO STATUS	ROLES TO MAKE THIS TRANSITION	REQUIRED FIELD(S) FOR STATUS TRANSITION
(New Document)	Any	All users with create permission	
Draft	Any	All users with edit permission	
Final	Any	All users with edit permission	
Ready for Review	Any	All users with edit permission	
Review Completed	Any	All users with edit permission	
PO Accepted	Any	All users with edit permission	

5. Select one or more roles that can make the status transition and the **Required Fields for Status Transition** from the drop-down lists.
6. Click **Save**.

- To add a new transition rule, click **Add Transition**.

A new row is added for the work flow.

- Select the **From Status** and **To Status** from drop-down lists, select the **Select Roles** option and select one or more roles that can make the status transition, select the **Required Fields for Status Transition**, and click **Save**.

The work flow is now saved. When a user submits or edits the status of a document, he or she sees only the options that are allowed by the work flow.

### Find and view a document

You can find and view a document by navigating to it, by searching for it, or by viewing it from a URL or unique identifier. You can also open a document from TeamForge directly into Microsoft Office applications.

#### *Go to a document*

If you know where a document is posted, or would like to browse to see what documents are available, you can find a document by navigation.

You can get to documents from your project home page or by searching the project's documents page.

- Choose a way to find documents:

- Look for a **Documents** window on your project home page. If your project manager has added a document component, you will find all your project documents there (see [Create a project page component](#) on page 136 to know how to add a "Documents" component to the project home page).
- If there is no Documents component, click **DOCUMENTS** in the project navigation bar. The **List Documents** page appears.

- Select a documents folder.

If a documents folder contains sub folders, you can expand the document tree by clicking the arrow to the left of the folder name.

- To view or save a document, click the file type icon to the left of the document name.

Use the filter options to locate the document you want. Documents with file types that can be viewed in a browser, such as HTML and text, are displayed in a TeamForge window. For other file types, use the appropriate viewing software.

- To view additional information about the document, click the document name.

#### *Search for a document*

Search for a document by its title, its content, the project or folder in which it is located, its file type, its creation date, or its status.

If you don't know which project to search, or want to search for documents and other types of items simultaneously, use global search.

- Click **DOCUMENTS** from the **Project Home** menu.
- Click **SEARCH DOCUMENTS**.
- On the **Document Search Criteria** page, enter the desired search criteria.



**Tip:** You can use wildcards.

- To search by location: If you know the document is located in the current folder, check the **Current Folder** in the **SEARCH CATEGORY** field. Otherwise, the default **Root Category** searches all document folders in the project.
- To search by name or description, enter the text in the **SEARCH KEYWORDS** field and select **Search Name and Description**.
- To search by document content, enter the text in the **SEARCH KEYWORDS** field and **Search Within Document**.
- To search by name, enter the text of the document name in the **DOCUMENT NAME** field.
- To search all document versions, select **SEARCH ALL VERSIONS**.
- To search by date range, enter the start and end dates for the search. Click the Calendar icon to select dates from a calendar.

- To search by author, enter the user name in the **CREATED OR EDITED BY** field. Click the search icon to display a list of project members.
- To search by document status, select one or more document statuses from the **Status** list box.


#### 4. Click **Search**.

TeamForge lists the documents matching your search criteria.

#### **Create a document**

You can create a document by uploading a file or by entering the text directly into TeamForge.

You can also create a document in TeamForge. TeamForge automatically detects the file extension for many common file types.

 **Note:** To create a document, you must have the Document Admin permission, or 'Create/View' permission for document folders or documents.


1. Click **DOCUMENTS** from the **Project Home** menu.
2. Using the document tree, find the folder in which you want to create the document, and click **Add Document**.  
You can create a document in the root document folder or any other document folder.
3. On the **Create Document** page, enter a name and description for the document.
4. You can tag a document if you want to. Select one or more tags from the **TAGS** field.

Your Project Administrator must have set up tags for you to add tags to your documents. For more information see [Set up Tags](#) on page 313. You can also create new tags if you have Document CREATE/EDIT permissions. To create new tags, just type a new tag title and click **Create a new tag** link.


TAGS:

STATUS:\*

5. Identify the status of the document.
  - **Draft** - The document is not yet completed.
  - **Final** - The document is finished.
6. If you want to prevent others from editing or downloading the document, select appropriate check boxes in the **LOCK DOCUMENT** section.
  - To lock a document for editing, select **Prevent others from editing**.
  - To lock a document for downloads, select **Prevent others from downloading**. The download lock prevents other users from downloading the document. You must have enabled the edit lock to make this selection.
  - To enable the Document Admin to edit the document, select **Enable Document Admin to edit**. You can choose to allow the Document Admin to unlock and edit a document after enabling the edit lock.

 **Note:**

  - By default, all the locks are disabled. You must have the edit lock enabled to select the download lock. Selecting the edit lock automatically enables the Document Admin to edit the document.
  - To choose locks in the **LOCK DOCUMENT** section, you must have the Document Edit permission.
7. Select the document creation method.
  - To upload a document, select **Upload From File** and browse to the desired file.
  - To enter the content as text, select **Create Online**.
  - To upload a document from a specific URL, select **Enter Document URL** and provide the URL.
 

 **Note:** You cannot prevent others from editing or downloading the document if you choose this URL option. Edit and download locks are not effective in this case.

8. Write any comments you may have about this version of the document.

9. Save your changes:

- Click **Save** to save the document and return to the document list view.
- Click **Save and View Details** to save the document and display the document details.

### **Edit a document**

When you are called on to help develop a document, you can edit the document online or by posting a new version.

To edit a document that is currently locked for editing, the user must have the Document Edit permission. To download a document that is currently locked for downloading, the user must have the Document View permission. A Document Administrator that have the option, **Allow Document Admin to edit** enabled in the **LOCK DOCUMENT** section or a Site Admin can also remove the edit and download locks.

### *Update a document*

To update a document, you can upload a new document version or change the document's name, description, or lock status.

You can also update a document in TeamForge directly from one of the following Microsoft Office applications:

- Word
- Excel
- Project
- PowerPoint

1. Click **DOCUMENTS** from the **Project Home** menu.
2. Using the document tree, find the folder containing the document that you want to edit.
3. Select the document you want to edit by clicking the document name.
4. On the **Document Details** page, click **EDIT / UPDATE**. You can see this button at the top right corner of the page only if you have Document Edit Permission.
5. On the **Edit Document** page, make your changes.
  - Update the name or description by entering new text in the appropriate field.
  - Change the status value, if appropriate.
  - To set locks, select appropriate check boxes in the **LOCK DOCUMENT** section.
  - Update a binary document by selecting **Upload From File**, then browsing for the updated file.
  - Update a text document by selecting **Create Online**, then editing the document text.
6. Enter a version comment to describe the changes made, if appropriate.
7. Click **Save**.


The document is now updated.

### *Change the active document version*

You can specify which version of a document is the active version at any time.

TeamForge stores every document version that you or any other user posts. When you open a document from the document list view, you open the active document version. When you open a document's **Document Details** page, the information you see is about the active document version.

The **Versions** section of the **Document Details** page lists all document versions and indicates which one is currently active. By default, the most recently posted document version is the active version.

1. Click **DOCUMENTS** from the **Project Home** menu.
2. Using the document tree, find the folder containing the document.
3. Select the document by clicking its name.
  -  **Tip:** On the **Document Details** page, in the **Versions** section, the active document is indicated with a check in the **Active** column.
4. Select the version that you want to make the active version, and click **Mark As Active**.

The selected document version is now the active version.

### *Lock a document*

While you are working on a document, you can lock it to prevent others from editing or downloading it and then posting new versions to TeamForge .

You can lock a document when you create it or at any time afterward. When a document is locked, any user with the Document View permission can view or download the document. You can edit or download a locked document, if you are a Site Admin or a Document Administrator that have the option, **Allow Document Admin to edit** enabled in the **LOCK DOCUMENT** section.

1. On the **Document Details** page, confirm that the lock status is currently Unlocked.
2. Click **EDIT / UPDATE**. You can see this button at the top right corner of the page only if you have Document Edit Permission.
3. On the **Edit Document** page, select appropriate check boxes in the **LOCK DOCUMENT** section.
4. Click **Save**.

The document is now locked. It is represented by the Padlock icon in the document list view.

### *Unlock a document*

To work with a document that someone else has locked, you must unlock the document.

You can unlock a document that you have locked before. If you want to unlock a document that was locked by someone else, you must be a Site Admin or a Document Administrator that has the option, **Allow Document Admin to edit** enabled in the **LOCK DOCUMENT** section.

1. On the **Document Details** page, confirm that the document is currently locked.
2. Click **EDIT / UPDATE**. You can see this button at the top right corner of the page only if you have Document Edit Permission.
3. On the **Edit Document** page, clear the check boxes in the **LOCK DOCUMENT** section.
4. Click **Save**.

The document is now unlocked. All project members with the Document Edit permission can now update the document.

### **Get a document reviewed**

There are a few things you have to do to get the most out of a document review.

#### *Start a document review*

To get a document reviewed, select the appropriate reviewers, set a due date, and explain to the reviewers what you need from them.

Before starting a document review, you must create the document in TeamForge .

1. Click **DOCUMENTS** from the **Project Home** menu.
2. In the document tree, click the title of the document you want input on.
3. On the **Document Details** page, click the **REVIEW** tab, then click **Start a Review**.



**Note:** If a review is already in process, **Start a Review** does not appear, because you can only have one review at a time.

4. On the **Start a Review Cycle** page, give the review a descriptive name.
5. Enter a due date for reviewers.

Click the Calendar icon to select dates.



**Tip:** Comments are due at 12:00 midnight on the date selected. Therefore, if you want all document reviews to be completed by end of day Tuesday, select Wednesday as the review end date.


6. Select the required and optional reviewers.

Click the Search icon to pick from a list of project members whose access permissions enable them to see this document.

7. Write a message to your reviewers describing the document and the type of input you need from them.

Your reviewers will get this message in an automatic email that shows you as the sender.




 **Tip:** If you add yourself to the reviewers list, you will get a copy of this email for future reference.

## 8. Click **Save**.

The document is now submitted for review.

- All reviewers receive an email with the details of the review, your message, and a link to the **Review Details** page where they can download the document and post their responses.
- The document status is changed to **Review** and remains so until the review is closed.

 **Tip:** To help keep track of the development of your document, associate a tracker artifact with your document review.

### *Read review responses*

After starting a document review, you receive email notifications whenever a reviewer posts review responses.

You can read the review responses from the email or from the **Review Details** page. The **Review Details** page is the starting point for managing all document review activities.

1. Click **DOCUMENTS** from the **Project Home** menu.
2. Using the document tree, find the folder containing the desired document.
3. Click the title of the document.
4. On the **Document Details** page, click the **REVIEW** tab.  
A summary of all open and closed reviews appears.
5. Click the title of the desired review.  
Review comments are listed in the bottom section of the page.

You can also reach the **Review Details** page by clicking the link provided in any related email notification or by clicking the title of the review in the **My Document Reviews** section of your **My page**.

### *Edit review details*

At any time after a review is started, you can change the review name, the required and optional reviewers, and the due date.

When you edit review details, affected reviewers get an email notification with the details of the change.

1. Click **DOCUMENTS** from the **Project Home** menu.
2. Using the document tree, find the folder containing the desired document.
3. Click the title of the document.
4. On the **Document Details** page, click the **REVIEW** tab.  
A summary of all open and closed reviews appears.
5. Click the title of the desired review.  
The **Document Review** page appears.
6. Click **EDIT**.
7. On the **Edit Review Cycle** page, make your changes.
  - Click the Calendar icon to select dates from a calendar.
  - Click the Search icon to display a list of project members from whom you can select.
8. Click **Save**.
  - When you add reviewers, new reviewers receive an email with the details of the review.
  - When you remove reviewers, removed reviewers receive an email telling them that their reviews are no longer necessary.
  - When you change the due date, outstanding reviewers receive an email with the new due date.

### *Send a reminder email*

To help project members keep track of their review obligations, you can send a reminder email to reviewers who have not yet submitted their responses.

1. On the **Document Review** page, find the outstanding required and optional reviewers.
2. Click **SEND REMINDER EMAIL**.
3. On the **Send Reminder Email** page, write a message to include in the email.
4. Choose whether to send the document as an attachment to the email.



**Note:** The URL to the document is included in the email regardless of whether the document is attached.

5. Click **Send**.

The email is sent to all reviewers who have not yet submitted responses.

#### *Close a document review*

After a document review is completed, close the review.

Reviews are never closed automatically. You must close any outstanding reviews on a document before you can start a new one on the same document.



**Note:** You can close a review even if the due date has not yet passed.

1. Click **DOCUMENTS** from the **Project Home** menu.
2. Using the document tree, find the folder containing the desired document.
3. Click the title of the document.
4. On the **Document Details** page, click the **REVIEW** tab.  
A summary of all open and closed reviews appears.
5. Click the title of the desired review you want to close.
6. On the **Document Review** page, click **CLOSE REVIEW**.
7. On the **Close Review Cycle** page, change the document status to **Final** or **Draft**.
8. Click **OK**.

The review is now closed. Any reviewers with outstanding reviews receive email notifications that the review is closed. No additional review responses can be submitted.

#### **Review a document**

To contribute to a review, you look at the document, consider responses from other reviewers, and post your own responses.

When you are identified as a reviewer of a document, two things happen:

- You receive an email about the review. The details include the due date, whether you are a required or optional reviewer, and any message text entered by the review submitter. The email also includes a link to submit your response.
- A new entry in the **Documents Awaiting Review** section of your **My Page** provides a summary of the review, including the review due date, the project in which the document is posted, and the person who is requesting your review.

1. Go to the **Document Review** page by clicking the link provided.
  - Click the link in the email notification you received with the review details.
  - Click the title of the document in the **Documents Awaiting Review** section of your **My page**.
2. On the **Document Review** page, in the **Submit a Response** section in the center of the page, enter your response.
3. If appropriate, use the **Browse** button to attach a marked-up copy of the document, or another attachment.
4. Click **Submit**.

Your response is now submitted.

- The review submitter receives an email notification containing your response.
- Your response comments appear in the **Responses** section of the **Document Review** page.

## Associate a document with other items

When a document is related to other CollabNet TeamForge items such as tasks, tracker artifacts, file releases, or news items, you can connect the document to the other item by creating an association.


Creating associations between items enables you to define relationships, track dependencies, and enforce work flow rules. Some example uses for document associations include:

- Associate a requirements document with supporting feature requests or bugs in the tracker.
- Associate a document, or a document review, with the task that requires it as a deliverable.
- Associate a master document, such as a product plan, with other supporting documents such as a test plan or beta plan.
- Associate a product announcement document with a file release.

1. Click **DOCUMENTS** from the **Project Home** menu.
2. Using the document tree, find the folder containing the document to which you want to add an association. Select the document by checking the box, then click **Details**.
3. On the **Document Details** page, click the **ASSOCIATIONS** tab.
4. Click **Add**.
5. In the **Add Association Wizard** window, select the items with which you want to associate the artifact:
  - **ENTER ITEM ID** - If you know the item's ID, you can enter it directly.

 **Note:**

- To associate an object in an integrated application from within TeamForge, use the [`<prefix_objectid>`] format. Successful associations appear hyperlinked.
  - Each integrated application displays its prefix on moving the mouse over the application name in the tool bar.
  - **ADD FROM RECENTLY VIEWED** - Select one of the last ten items you looked at during this session.
  - **ADD FROM RECENTLY EDITED** - Select one of the last ten items you changed.
6. Click **Next**.
  7. You may add a comment in the **ASSOCIATION COMMENT** text box.
  8. Save your work.
    - Click **Finish and Add Another** to add additional associations.
    - Click **Finish** to return to the **Details** page.

 **Note:** When an association is added to or removed from TeamForge objects such as tracker artifacts, tasks, documents, discussions, and file releases, a notification mail is sent to users monitoring these objects.

An option is provided at site level and user level to make sure whether the notification mail has to be sent or not. For more information on this, see [Configure your site's settings](#) on page 365.

9. Click the **Associations** tab to view a graphical representation of all the associated items.

Through the Association Viewer, you can choose to view associations in the form of a list or flip over to the Trace view to explore the layers of associations (including parent/child dependencies) laid out in a timeline. You can scroll across the Trace view by dragging the mouse over the association layer or use the 'Previous' and 'Next' arrows to view all the objects as events in a timeline.

While the Associations tab shows the count of the total number of associations, you can only view the most recent 500 associations when you click the Associations tab in case the artifact has more than 500 associations. You can, however, browse through the Association Viewer to view older associations.


## Manage documents

Project members with the Document Admin permission can copy and move documents. They can also create and administer documents. However, they can delete documents, only if they have the specific delete permission.

### *Copy a document*

If a given document belongs logically to two or more document folders, keeping a copy in each folder can help users locate it.

Copying a document can also be useful if you want to use an existing document as a template for creating a new document. When making a copy of a document, reference to the source document is maintained in the document title and the **Associations** tab.

 **Note:** Each copy must be maintained separately. Edits made to one copy are not propagated to other copies.


1. Click **DOCUMENTS** from the **Project Home** menu.
2. Using the document tree, find and select the folder containing the document that you want to copy.
3. Select the document you want to copy, then click **Copy**.  
The **Copy Files** dialog box appears.
4. Select the project from the **Select Project** drop-down list and select the document folder where you want to copy the selected document.
5. Click **Copy**.

The document copy is now placed in the selected folder, with the name "Copy of <Source Document Name>". The reference to the source document is also maintained in the **Associations** section of the **Document Details** page.

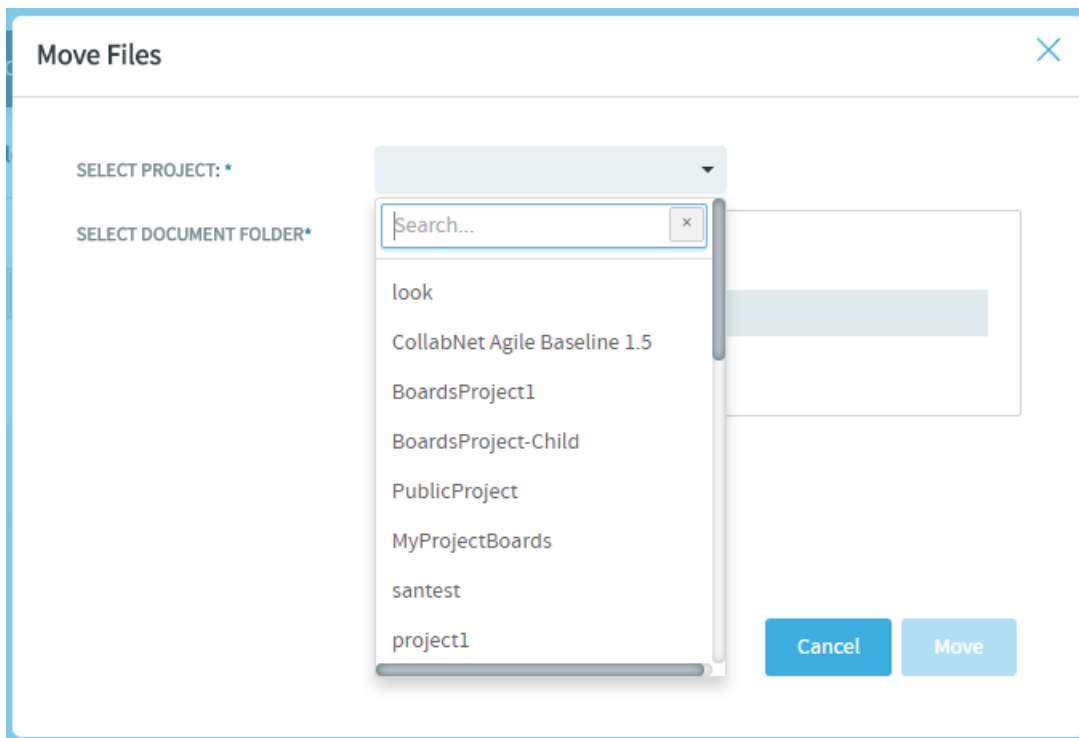
### *Move a document*

You can move it to another document folder within the project or to a document folder in another project at any time. This allows you to reorganize documents as needed.

You can move one or more documents in the same operation. You cannot move a document while it has an open review.

 **Note:** To move a document between projects, you must have document administration permission in both the source and destination projects.

1. Click **DOCUMENTS** from the **Project Home** menu.
2. Using the document tree, find and select the folder containing the document that you want to move.
3. Select the document you want to move, then click **Move**.  
The **Move Files** dialog box appears.
4. Select the project from the **Select Project** drop-down list and select the document folder where you want to move the selected document.




5. Click **Move**.

The document is now moved to the selected folder.

*Delete a document*

If you no longer need a document, you can delete it.

Deleting a document deletes all of the document's versions, plus the document's details and review history.

 **Note:** To delete documents, you must specifically have the delete permission ('Delete/View') for document folders or documents. This is irrespective of whether you have the Document Admin permission or not.

1. Click **DOCUMENTS** from the **Project Home** menu.
2. Using the document tree, find the folder containing the document that you want to delete.
3. Select the document you want to delete.
4. Click **Delete**.

The document is deleted.

*Delete document versions*

Document versions can grow in number over time. You may choose to delete one or more versions of a document if you no longer need them. This saves a lot of disk space and maintenance overheads.

You can delete a document's version only if:

- you have the delete permission ('Delete/View') for documents. This is irrespective of whether you have the Document Admin permission or not.
- the document has more than one version available.
- the selected document version is not marked as an active version.



**Caution:** Deleting a document's version permanently removes the version and cannot be undone. Exercise caution while deleting document versions.

1. Click **DOCUMENTS** from the **Project Home** menu.
2. Using the document tree, find the folder containing the document that you want to delete.
3. Click the document's name link to see its versions.
4. Select one or more document versions you want to delete from the **VERSIONS** tab and click **Delete**.

A confirmation message is displayed.

5. Click **OK** to delete.

The document versions are deleted.

### Organize documents

Project members with the Document Admin permission can create, edit, move and reorder document folders.

Similarly, project members with the relevant individual permission can create, edit, delete or view document folders.

#### *Create a document folder*

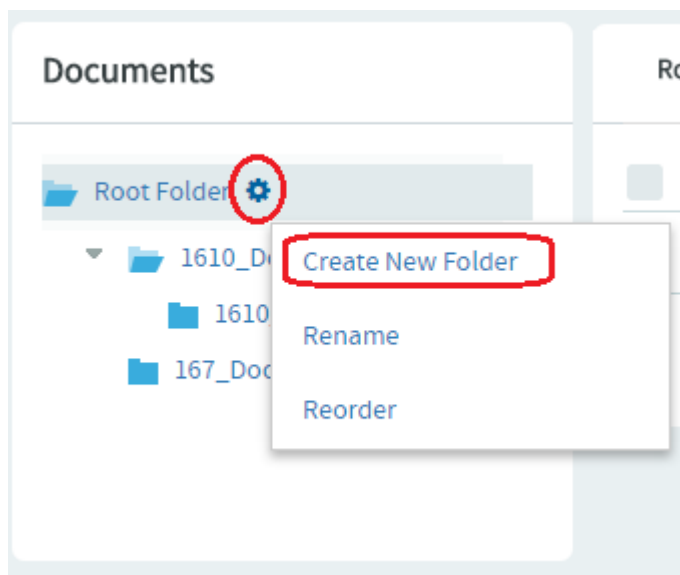
Document folders, like file directories, allow you to organize your documents in a logical, hierarchical manner using labels that are meaningful to you and your organization.

You can create folders in the root folder or as subfolders in any other folder. There is no limit to the number of folders you can create.

1. Click **DOCUMENTS** from the **Project Home** menu.
2. Using the document tree, find and hover your mouse over the folder where you want a new document folder created. For example, hover your mouse over the **Root Folder**.

A settings icon shows up when you hover your mouse over a folder.

3. Click the settings icon and select **Create New Folder** from the menu.



The **Create Folder** dialog box appears.

4. In the **Create Folder** dialog box, enter a name for the folder.
5. Click **Save**.

The folder is created.

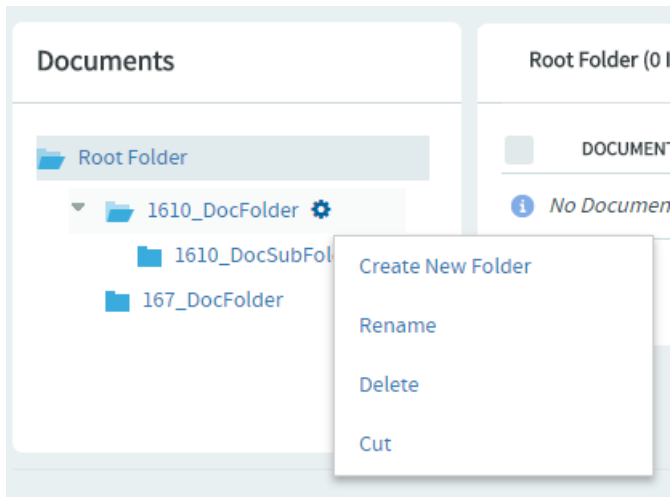
#### *Rename a document folder*

When the function of a document folder has changed, it's a good idea to rename the folder.

1. Click **DOCUMENTS** from the **Project Home** menu.
2. Using the document tree, find and hover your mouse over the folder which you want to be renamed.

A settings icon shows up when you hover your mouse over a folder.

3. Click the settings icon and select **Rename Folder** from the menu.



The **Edit Folder** dialog box appears.

4. In the **Edit Folder** dialog box, enter a new name for the folder.
5. Click **Save**.

The folder is renamed.

#### *Move a document folder*

You can move a document folder within a project.

You can move a folder in the following ways:

- From the root folder to a subfolder.
- From a subfolder to the root folder.
- From a subfolder to another subfolder.


When you move a document folder, any documents and subfolders that it contains are also moved to the destination folder.

Project members monitoring the destination folder will continue to receive monitoring notifications after a folder is moved. Project members not monitoring the destination folder will not receive further monitoring notifications.

1. Click **DOCUMENTS** from the **Project Home** menu.
2. Using the document tree, find and hover your mouse over the folder that you want to move.

A settings icon shows up when you hover your mouse over a folder.

3. Click the settings icon and select **Cut** from the menu.



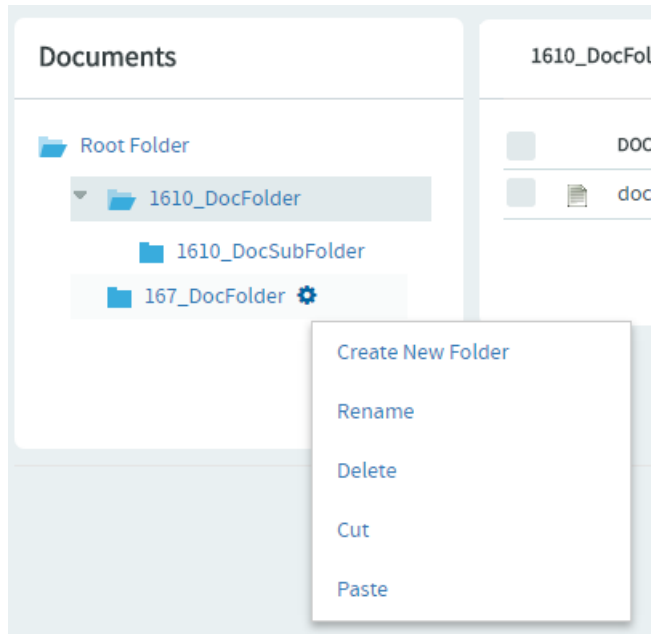
- Using the document tree, find and hover your mouse over the folder into which you want to move the selected folder.

A settings icon shows up when you hover your mouse over a folder.



**Note:** You can move a folder into the root folder or into any other folder.

- Click the settings icon and select **Paste** from the menu.



The folder is moved to the destination folder.

### *Reorder document folders*

To change the order in which document folders are displayed, you can alphabetize them or reorder them arbitrarily.

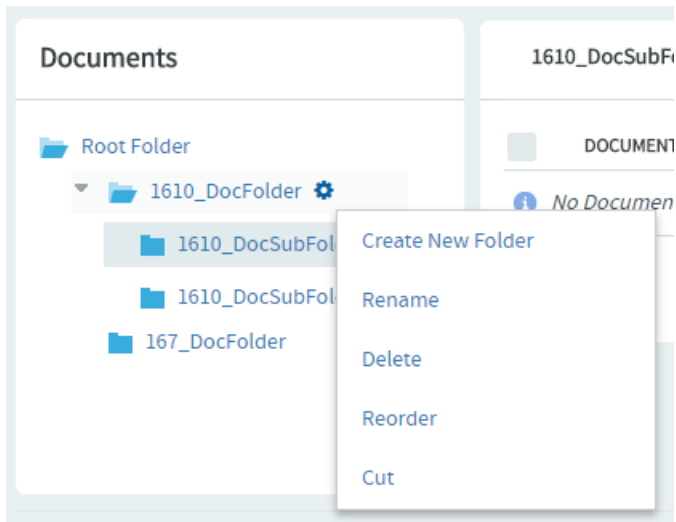
By default, document folders are ordered as they were created. New document folders are added to the end of the document tree.

- Click **DOCUMENTS** from the **Project Home** menu.
- Using the document tree, find and hover your mouse over the folder containing the subfolders that you want to reorder.

A settings icon shows up when you hover your mouse over a folder.

- Click the settings icon and select **Reorder** from the menu.





The **Reorder Document Subfolders** dialog box appears.

### Reorder Document Subfolders

Parent Folder: 1610\_DocFolder

Folders:

1610\_DocSubFolder  
1610\_DocSubFolder01



Cancel

Save

#### 4. Organize the subfolders:

- To sort them alphabetically, click alphabetize icon.
- To reorder a specific folder, select it by clicking the title, then click **Move Up** or **Move Down** until the folder is where you want it.

#### 5. Click **Save**.

#### *Delete a document folder*

If you no longer need a document folder, it's a good idea to delete it.

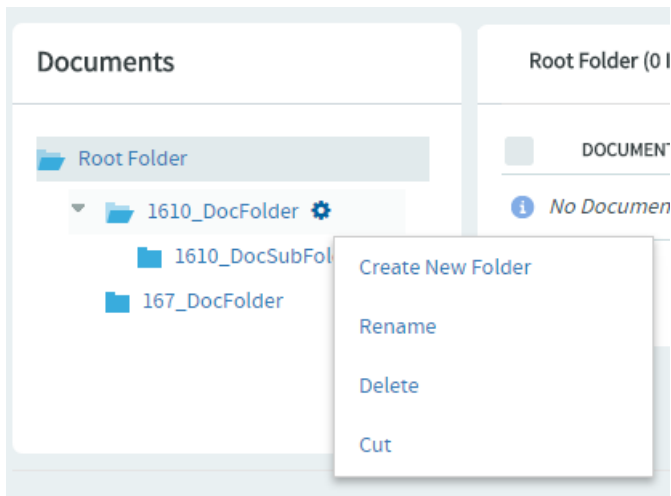
Deleting a document folder deletes all of the documents within it.

#### **Important:**

- Delete a document folder only if you are sure that you no longer need any of the documents within it. Move any documents that you want to save.

- Even if you have the Document Admin permission, you cannot delete a document folder unless you specifically have the delete permission.

1. Click **DOCUMENTS** from the **Project Home** menu.
2. Using the document tree, find and hover your mouse over the folder you want to delete.  
A settings icon shows up when you hover your mouse over a folder.
3. Click the settings icon and select **Delete** from the menu.



A confirmation message appears.

4. Click **OK** to delete the document folder.



**Attention:** If the folder contains one or more documents that you do not have permission to delete, you cannot delete the folder.

### Work with documents from Microsoft Office

You can use the **File** menu of your Microsoft Office application to work with documents that are stored on your TeamForge site.

#### *Open a TeamForge document in Microsoft Office*

You can get to your Microsoft Word, Excel, Project, and PowerPoint documents in CollabNet TeamForge even if the CollabNet TeamForge web interface is not open.

1. Open your Microsoft Office application.
2. From the **File** menu, choose **Open from CollabNet TeamForge**.
3. In the **Login** window, provide the following information:
  - Your TeamForge user name and password.
  - The URL of your organization's TeamForge installation.



**Note:** Do not enter the `http://` or `https://` before the server name.

4. Click **Log In**.  
The **Selection Dialog** window shows all of the projects, document folders, and documents that you have permission to view.
5. Select the document you want to open.

The document opens in your Microsoft Office application.


#### *Save a document to TeamForge*

You can store your Microsoft Office documents in CollabNet TeamForge even if the CollabNet TeamForge web interface is not open.


CollabNet TeamForge works with Microsoft Word, Excel, Project, and PowerPoint.

1. Open your Microsoft application.


2. Open the document you want to save.
3. From the **File** menu, choose **Save to CollabNet TeamForge**.
4. In the **Log in** window, enter the following information:
  - Your TeamForge user name and password.
  - The URL of your organization's TeamForge installation.

 **Note:** Do not enter the `http://` or `https://` before the server name.

5. Click **Log In**.  
The **Selection Dialog** window displays all of the projects and document folders in which you have permission to create documents.
6. Select the document folder in which you want to save the document.
 

 **Note:** Do not click on an existing document. This will prompt you to update that document.
7. In the **Document Properties** window, clear **Use Filename**, then enter a name in the **CollabNet TeamForge Document Name** field.  
By default, the file name is used as the document name.
8. Write a description for the document.
9. Choose whether the document is in draft or final status.
10. To lock the document and prevent others from editing it, select **Lock Document**.
11. Write a version comment, if appropriate.
12. Click **OK**.

The document is now saved in CollabNet TeamForge. You can see the document by going to the CollabNet TeamForge document folder that you selected.


 **Note:** After you save a document to CollabNet TeamForge, future saves will also save in CollabNet TeamForge. If you want to save a document to a local drive, choose **Save As** from your **File** menu.

## Manage a product with CollabNet TeamForge


A product owner can use TeamForge to help define the scope of a product and contribute to its planning and implementation.

### Define the scope of your project

Defining scope is an iterative, interactive process. As you go through it, you'll find elements of your scope expanding, shrinking or changing shape in response to feedback from analyzing and planning out the work.

 **Note:** It's a good idea to get the feature definition process under way *before* setting up planning folders (see [Create a planning folder](#) on page 28), because defining features gives you the raw material for your planning process.

1. Before you get started, you'll do some kind of user research, even if it's only a few phone calls, to get an idea of the needs of the customers you hope to satisfy. Express these needs and desires as stories about what a user can do with your product. Then create an artifact for each user story you identify.  
A user story describes the situation after your product has been launched: What can the user do now that they couldn't do before?
 

 **Tip:** For best results, your user story artifact should be limited to a clear description of the capability you want in the product. Remember that a user story is not an implementation plan. Details about the implementation will be recorded in the user story artifact, but when you write a user story, try to leave the implementation up to the developer who will be responsible for it.
2. You may want to define one or more fictional users who resemble the real-life users of your product. This can help simplify and focus your thinking.
3. Use the Priority field to express your opinion of how important the story is to the user. In general, the most eagerly desired capabilities will be addressed first. (During implementation planning, your Priority setting will be used as one input in summing up the effort involved in each priority level.)

4. Make it as clear as you can at the outset what degree of functionality is acceptable. For example, if your team is creating an airplane, how high must it be able to fly? How far must it go before refueling? How many passengers must it carry? Stating your acceptance criteria concretely helps reduce the time needed for ongoing reviews and changes.

### Monitor project output

Regularly test and measure the features your team produces, involving real users as much as possible.

Accurately representing the needs and responses of real-life users is a key part of the product owner's role.

1. Schedule acceptance meetings throughout the iteration, either at regular intervals or whenever a given user story is completed.

Use the latest working build of the product, and exercise the functionality live.


2. Test that the functionality implemented matches what is specified in the user story.

Also review the user interface and user documentation associated with the user story, if any.

3. Enter any issues in the artifact's **Comments** field.

4. Update the user story artifact to reflect the outcome.

- If a user story is accepted, change the artifact status to "Accepted" or the equivalent.
- If a user story is not accepted, change the artifact status back to "In Development," or the equivalent, for further work. You can review it again in another acceptance meeting.

 **Note:** Sometimes a feature turns out not to meet the user's needs even though it has been implemented as specified. This is normal. It is not uncommon for the user's needs to change after the user story has been captured, or for the user research process to miss one or more important details. Just note the discrepancy in the user story and route the artifact back into the appropriate part of the process.

### Organize your project with planning folders

Create and populate all the planning folders you need to capture the work you are planning.

A planning folder is a way to organize work into feasible chunks and monitor its progress. When you've thought about the general categories the work falls into, you can to set up planning folders that reflect those ideas.

A planning folder can represent:

- A set of tasks, such as Iteration 3, or "Initial infrastructure development."
- A period of time, such as "April," or "Q2-2010."
- A phase of development, such as "Testing" or "Deployment."
- A component of the product, such as "Chapter 12" or "Rear stabilizer."

 **Tip:** It often makes sense to set up planning folders *after* you have outlined and analyzed the features you plan to deliver. See [Define the scope of your project](#) on page 171.

When you've set up your planning folders, you have four views available to work with them:



- **LIST:** The list view.
- **PLAN:** The planning board view.
- **TRACK:** The task board view.
- **KANBAN:** The kanban board view.

For more information, see the related concept: [What is a planning folder?](#) on page 990

#### Create a planning folder

A planning folder is a way to organize work into feasible chunks and monitor its progress. As a project admin or as a user with the appropriate permissions, create and populate all the planning folders you need to capture the work you are planning.

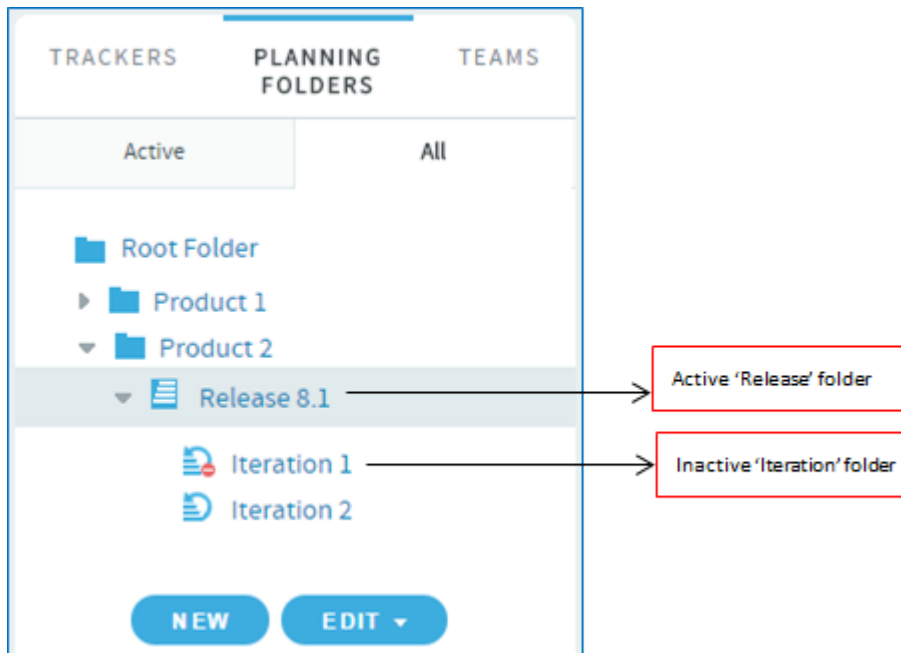
When you've thought through your plan, express it in one or more planning folders.

- 👉 **Tip:** It often makes sense to set up planning folders *after* you have outlined and analyzed the features you plan to deliver. See [Define the scope of your project](#) on page 171.

For your agile projects, you have the option to create planning folders specific to iterations and releases, or you can create generic planning folders which you may customize later. When you've thought about the general categories the work falls into, you are ready to create planning folders that reflect those ideas.

A planning folder can represent:

- A set of tasks, such as Iteration 3, or "Initial infrastructure development."
- A period of time, such as "April," or "Q2-2010."
- A phase of development, such as "Testing" or "Deployment."
- A component of the product, such as "Chapter 12" or "Rear stabilizer."



1. Click **TRACKERS** from the **Project Home** menu.
2. Under **PLANNING FOLDERS**, click the folder in which you want to create a new planning folder. Click **NEW**. An appropriate planning folder creation page is displayed.
3. Select the type of planning folder you want to create: **Release**, **Iteration** or **Folder**.

While "Release" and "Iteration" are specific planning folder types, "Folder" is your standard, generic planning folder. Though having multiple planning folder types help you logically organize your planning folders, there's no hierarchical restrictions on these different planning folder types. For example, as a user, you have the liberty to create a "Release" planning folder as a child of an "Iteration" planning folder, which may not be the ideal use case in an Agile project.

4. Enter a brief and descriptive name for your planning folder.

For example:

- In an agile project, depending on your requirement, you can create two or more iteration folders called "Iteration 1", "Iteration 2" and so on within the Release planning folder.
- In a phased, waterfall-style project, you might name your first planning folder "Design," the next "Build," and so on.

- 👉 **Tip:** Don't worry if you don't have anything of interest to put into your planning folders yet. The parallel process of filling out the feature tree will provide plenty of material for this.


- 👉 **Note:** This field is mandatory.

5. Use the **DESCRIPTION** to briefly signal the kind of work that will be contained in this planning folder.

Include enough information to help people get up to speed quickly when they join your project. But save most of the detail for the individual tracker artifacts, where project members will spend most of their time.


6. For **STATUS**, (available since TeamForge 5.4), select a value that communicates where the planning folder is in its life cycle.

For a new planning folder, you'll probably want to select Not started.


 **Tip:** You can create more values to choose from in your Project Settings.

7. If you have set up a file release in the File Releases tool to deliver the work you are tracking here, you can identify it in the **FILE RELEASE** field (available since TeamForge 5.4).


Any artifact you add to this planning folder will also appear in the Planned Tracker Artifacts tab when you look at that file release.

 **Note:** If you are creating an Iteration folder, this field displays the default file release value set at its parent folder level ("Release" folder) which you can modify.

8. For **CAPACITY**, (available since TeamForge 5.4), provide a number that expresses how much work your team thinks it can do in the period represented by this planning folder. This field is available only for "Iteration" and "Folder" types.

 **Tip:** When you create a new planning folder, it can be a good idea to set its capacity to zero until you get a feel for the artifacts and resources you are working with.

9. For **POINTS CAPACITY**, (available since TeamForge 7.1), provide a number that represents the amount of work (expressed as number of story points) that you think can be handled in that specific planning folder. This field is available only for "Release" and "Folder" types.

 **Note:** To assist project managers gauge the planning folder's points capacity, when a release or standard planning folder is selected in one of the *planning board* swim lanes, this information is displayed as a ratio of planned number of story points versus the actual number of story points. This information does not apply to Iteration folders and therefore does not appear on the planning board swimlanes when you select an Iteration planning folder. For more information, see *Use the Planning Board* on page 182.

10. Select or enter a start and end date for the work that this planning folder will contain.

You can change these dates at any time. However, it's important to specify dates even if they're not firm, because useful visual aids and reports depend on them.

 **Note:**

- The start and end dates of the child planning folder should be within that of its parent folder.
- These dates are mandatory for an "Iteration" planning folder because the burndown chart for Iteration planning folders is enabled only when you specify the start and end dates of an Iteration planning folder. Also, these dates are not mandatory for a Release planning folder because the burndown chart of a Release planning folder takes into account the start and end dates of its child planning folders (Iterations) and not its own.


11. Select the unit that is suitable for the planning folder from the **DISPLAY EFFORT IN** field. Eg. if the planning folder represents a sprint, then select the unit as **Hours** and if it represents a collection of sprints, then select the unit as **Points**.

 **Note:** Units are configured at the project level and not at the planning folder level.


12. Select either **Effort** or **Points** from the **BURNDOWN VIEW BY** drop-down list to set the View Mode in the burndown charts (in the planning folder **List Artifacts** page). This field is available only for "Folder" type.

Depending on the planning folder type, the burndown chart is displayed.


- Burndown chart for the standard planning folder type which includes the View Mode:



- Burndown chart for "Iteration" planning folder type:



- Iteration-wise burndown chart for a "Release" planning folder type:



13. Click **Create**. The planning folder is created.

*Update a planning folder*

A planning folder is a dynamic representation of a changing situation. From time to time, you'll want to update its name, description and other parameters to reflect changes in the underlying work.

1. Click **TRACKERS** from the **Project Home** menu.


2. In the planning folder tree, click the planning folder you want to update.
3. Use the **EDIT** menu to make your changes.

### *Manage statuses for a planning folder*

To help team members understand how to work with a planning folder, create statuses for it that correspond to the planning folder's life cycle.


When you have created a planning folder status, you can apply it to any planning folder in your project.

For example, at a given moment you might have one planning folder in "Under development" status, while another is in "Finished" status and another is in "Preliminary scoping" status.

 **Tip:** You may also want to move the planning folder to a relative position in the planning folder list that communicates where it stands in the order of work.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Planning Folder Settings**.
3. Under **Statuses**, click **Add**.
4. Configure your new planning folder status.
  - a) Give the status an unambiguous name that others will find easy to understand.
  - b) Specify whether the status counts as active, inactive, or freeze.

Project members may want to exclude inactive planning folders from their view, to avoid overload.


 **Note:** If you have selected the status as "Freeze" for any planning folder, you cannot move any artifact into or out of that planning folder. However, you can update any artifact within that planning folder.

- c) Select which value is the default value for new planning folders and planning folders that originated in earlier TeamForge versions.
- d) Arrange your planning folder statuses in the order that makes sense for you.

### *Reorder planning folders*

When the sequence of work in your project changes, it's a good idea to reorder the planning folders so that they accurately reflect the real order of work.

For example, it's not unusual for product owners, who must constantly weigh the changing business value of features, to ask for one set of user stories to be put aside so that another set can be delivered first. When this happens, a project manager may want to move the corresponding planning folders so that their relative positions visually confirm their relationship.

 **Tip:** You may also want to set the planning folder's status to a value that communicates the change in the order of work.

By default, planning folders are ordered according to when they were created. New planning folders are added to the end of the folder tree.

1. Click **TRACKERS** from the **Project Home** menu.
2. In the folder tree, click the folder containing the subfolders that you want to reorder.
3. Choose **Reorder Subfolders** from the **EDIT** menu.
4. Organize the subfolders:
  - To sort them alphabetically, click Alphabetize icon.
  - To reorder a specific folder, select it by clicking the title, then click move up or move down arrow until the folder is where you want it.


### *Reorder the contents of a planning folder*

You can move artifacts in a planning folder up and down to reflect the order in which they should be addressed.

A set of user stories often includes dependency relationships that require that one user story be addressed before another. For example, suppose a backend database modification must be in place before a user interface component can be created to access it. Both user stories may have the same degree of priority, but one has to be done before the other.



One way to manage this is by ranking the user stories in your planning view so that they reflect the sequence you want. When you do this, team members have a better chance of selecting work efficiently and reducing the need for rework.

 **Remember:** Rank is not the same as priority. Two artifacts may have the same priority but different ranks, or sequence positions.

1. If you are not an administrator on this project, ask your project administrator or site administrator to assign you sequencing permission.
2. Click **TRACKERS** from the **Project Home** menu.
3. In the folder tree, open the planning folder containing the artifacts that you want to reorder.
4. Click **RANK**.  
You can now drag and drop the artifacts in your planning folder into whatever relative positions you want.

To go back to the sort order you had before (for example, sorted by priority), click **SORT**.


#### *Move a planning folder*

A planning folder's position may reflect scheduling changes in relation to the work contained in other planning folders.

For example, when plans change and the work represented in a later planning folder moves ahead of an earlier planning folder, move the folders around in the hierarchy to reflect the new relationship.

When you move a document folder, any documents and subfolders that it contains are also moved to the destination folder. Before promoting a planning folder to a higher level in your hierarchy of folders, make sure its member artifacts (and its members' parent artifacts, if any) are all assigned to the same planning folder, or to a planning folder that belongs to the same folder hierarchy.

1. Click **TRACKERS** from the **Project Home** menu.
2. In the planning folder tree, find the folder that you want to move. Click the title to select it.
3. Choose **Cut** from the **EDIT** menu.
4. In the folder tree, find the folder into which you want to move the selected folder. Click the title of the folder to select it.

 **Note:** You can move a folder into the root folder or into any other folder, as long as no member artifact ends up having a parent belonging to a higher-level folder in the same tree.


5. Choose **Paste** option from the **EDIT** menu.  
The folder is moved to the destination folder.

#### *Export planning folder artifacts*

To use the contents of artifacts from a planning folder in other applications, export them to a .csv, .xml, .xlsx or tab-delimited file.

For example, sometimes it can be useful to look at the status of a project by sharing a spreadsheet using Microsoft Excel or Google Spreadsheets.

1. Click **TRACKERS** from the **Project Home** menu.
2. On the **List Trackers, Planning Folders and Teams** page, click **PLANNING FOLDER**.
3. On the **Summary** page, click the desired planning folder.
4. From the **List artifacts** page, click **SORT** or **RANK** as required. You can also select specific set of artifacts to be exported.

 **Note:** When you export artifacts in the rank mode, the parent artifacts are listed along with the children as ranked in the **List Artifact** page. Similarly the sort order in the list view is retained in the exported list as well.

5. Click **Export**.
6. In the **Export Data** window, select an export format that you can import into the other application.  
For example, to use the data in a spreadsheet program, select CSV.

7. Move the fields you are interested in from the **AVAILABLE COLUMNS** list to the **SELECTED COLUMNS** list, then click **Export** to complete the process.

### Manage your project with Teams


The Team feature, as the name implies, enables you to create logical groups of team members to carry out project activities more efficiently in an agile environment.

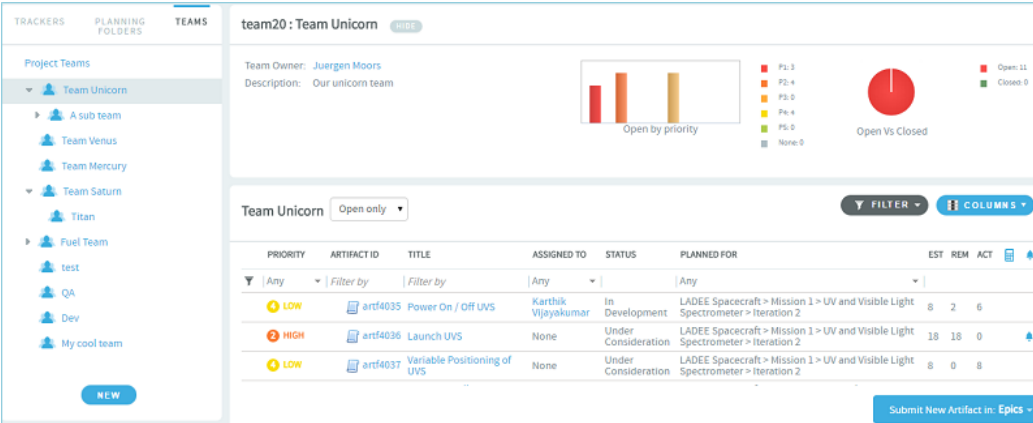
#### Teams - an overview

Using Teams, project activities can be planned, tracked, collaborated, reported and executed in a more organized and structured manner.

This feature facilitates effective communication among the team members who need to be aware of the changes and latest updates occurring in their projects and act accordingly. For example, using the team's view of backlogs, any project impediment can be communicated to the team and resolved quickly.

The Team list view consists of a tree structure of all the teams in a project, a summary and a filter table of the selected team. The filter table allows you to assign artifacts to a specific team. You can configure your team filter table and filter work items depending on the information you want. The filter table lists only the artifacts for the selected team. If you have selected a parent team, the filter table does not include the artifacts of the subteams (child teams).

 **Important:** To access the Team feature, you must have the Tracker 'view' permission.



The screenshot displays the 'Team Unicorn' view. On the left, a sidebar shows a tree of teams including 'Team Unicorn', 'A sub team', 'Team Venus', 'Team Mercury', 'Team Saturn', 'Titan', 'Fuel Team', 'test', 'QA', 'Dev', and 'My cool team'. The main area shows team details: 'Team Owner: Juergen Moors', 'Description: Our unicorn team', and two charts: 'Open by priority' (a bar chart) and 'Open Vs Closed' (a pie chart). Below the charts is a table of artifacts for the selected team, filtered to 'Open only'. The table has columns: PRIORITY, ARTIFACT ID, TITLE, ASSIGNED TO, STATUS, PLANNED FOR, EST, REM, ACT. Three artifacts are listed with their respective details.

### Team roles and permissions

Other than the Tracker 'view' permission, there are no role permissions that you need to set specifically for Team on the **Permissions** page of **Project Admin**. Depending upon your project role, you can view the appropriate icons on the Team list view and perform Team related tasks.

- **Project administrator:** You can create, edit, delete and view teams.
- **Team owner:** Any team member can be designated as a team owner. You can be a team owner for more than one team. As a team owner, you can only edit and view your team details.
- **Team member:** Any project member can be added to a team. As a team member, you are only allowed to view the team members.

### Permissions

The following table lists the various roles and their specific permissions associated with Team:

Permissions	Team member	Team owner	Project admin
Create a team			✔

Permissions	Team member	Team owner	Project admin
Delete a team			✓
Add team members while creating a team			✓
Designate a team owner while creating a team			✓
Edit your team details		✓	✓
Add / delete team members while editing your team(s)		✓	✓
Designate a team member as the team owner while editing your team		✓	✓
View a team tree	✓	✓	✓
Assign artifacts to any team / team member	✓	✓	✓

### Manage Teams

Using the Team list view, you can create, edit, delete and view teams at the project level.

You can access Teams only if you have the Tracker 'view' permission.

### Create a parent team

As a project administrator, you can create a parent team at the root level, **Project Teams**.

1. Click **TRACKERS** from the **Project Home** menu.
2. On the left pane of the default list view, click **TEAMS**.  
All the teams in your project are displayed on the pane.
3. Select **Project Teams** and click **NEW** to create a new team.

Alternatively, you can hover the mouse on **Project Teams** and click the **Create Team** button.



The **Create Team** window is displayed.

4. Give a name and description to the new team.

- The team name is a required field. You are allowed to enter a maximum of 64 characters including spaces.
- No two teams on the same level in the Team tree hierarchy can have the same name within a project; however, a team can have the same name as that of a deleted team.

5. Add team members by selecting the users from the list displayed.

Alternatively, you can search for specific users by typing the relevant alphabets and selecting from the list of matching names.

- The user who creates a team becomes the team owner automatically.
- A team member can be a part of more than one team.

6. Designate a team member as the team owner by hovering the mouse on the specific name and clicking the **Team Owner** button.



7. Click **Create**.

A new team is created.

### Create a sub team

A team can have many sub teams.

8. On the Team tree view, select the team for which you want to create a sub team and click **NEW** or hover your mouse on the specific team and click the **Create Sub Team** button.



The **Create Sub Team** window is displayed.

9. Provide all the required information, and click **Create**.

A sub team is created.

### Delete a team

As a project administrator, you can delete any team from your project. If a team has sub teams (child teams), you can delete the parent team only after deleting all its child teams.

10. On the Team tree, hover your mouse on the team you want to delete and click the **Delete Team** button.



### Edit a team

As a project administrator, you can edit any team in your project. Similarly, if you are a team owner, you can edit the team(s) for which you are the owner.

11. On the Team tree, hover your mouse on the team you want to edit and click the **Edit Team** button.



The **Edit Team** window is displayed.

12. Make the required changes:

- Edit the team name and description.
- To add more team members, select the names from the list displayed.
- To delete team members, hover the mouse on the specific name and click the **Remove Member** button.



- To designate another team member as the team owner, hover the mouse on the specific name and click the **Team Owner** button.



**Note:** As a team owner, you may decide to designate a different user as your team owner. In that case, once you update your changes, you will only be a team member and you will no longer have the permission to edit your team.

13. Click **Save**.

#### View team members

14. On the Team tree, hover your mouse on the team whose members you want to view and click the **View Team Members** button.



The **Team Members** window is displayed.



### Manage your product release with the Planning Board

The Planning Board is an important tool for your TeamForge project's Agile planning activities. It enables you to plan and monitor the features that are required in each sprint (or iteration), and assign them from the product backlog to specific sprints.

When you've set up your planning folders, you have four views available to work with them:



- **LIST**: The list view.
- **PLAN**: The planning board view.
- **TRACK**: The task board view.
- **KANBAN**: The kanban board view.

The planning board view complements the list view. While the latter offers you capabilities to accomplish various actions such as create, edit, and delete artifacts, planning folders and teams, the former offers product owners (or similar users) the ability to view, rank and move artifacts across the three planning folders (swimlanes) in a physical board-like user interface. In the Planning Board, planning folders are represented as swimlanes. In each swimlane, the tracker artifacts for the selected planning folders are represented as cards. You can also have a team's view of artifacts (backlogs and tasks), which is a swimlane representation of artifact cards for the selected team in the selected planning folder.

TeamForge user roles and permissions that are in place for planning folders apply to all the four views.

#### *Use the Planning Board*

You start populating the Planning Board by selecting a planning folder for each swim lane.

The drop-down list for each swim lane displays the hierarchy of active planning folders for all the sprints and scrum teams that are under the project's root planning folder. By default, inactive planning folders are not shown in the swim lane drop-down list. You can organize the Planning Board the way you want by populating the swim lanes with planning folders of interest to you: for example, you may select the planning folder corresponding to the product backlog in the leftmost swim lane and various teams working on the release in the other two swim lanes, or a different sprint in each swim lane. Your selections are remembered for the current session.


In the Planning Board, artifact cards are displayed in ranked sequence within each swim lane. The tasks that you can accomplish while you work with the Planning Board include:

- Adding artifacts (quick add) in select planning folders with minimum required artifact information.
- Editing artifacts.
- Moving artifact cards within a swim lane and ranking them in select planning folders.
- Reassigning (move) artifacts from one planning folder to the other.

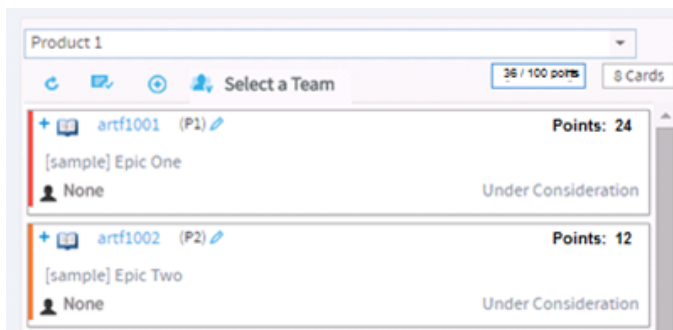
1. Click **TRACKERS** from the **Project Home** menu.
2. In the **List Trackers, Planning Folders and Teams** page, click **PLAN**.  
A Planning Board for the current project context is displayed.

3. For each swim lane in the Planning Board, select a planning folder from the drop-down list. The tracker artifacts contained by the planning folders are displayed as story cards within the corresponding swim lanes.
  - If a selected planning folder has subfolders, the artifacts within the subfolders are not displayed.
  - Cards have a color-coded bar to visually identify the priority. In addition, the open and closed cards' background is color-coded to uniquely identify the status.
  - For each card, the artifact ID, title, priority, status, assigned to and points (story points) are displayed. For an artifact, the estimated effort (in hours) is shown only if its points = 0. Relevant tooltips appear when you hover your mouse over these data elements.

Assistive information such as the number of cards in the selected planning folder and planning folder points capacity are displayed at the top of the swim lane.

 **Note:** The planning folder points capacity is not displayed for "Iteration" folder types.

For example, in the following screenshot, the planning folder points capacity is shown as **36 / 100 Pts.** Which means, you planned for 100 points, whereas the current points capacity (the sum of all open and closed artifacts) is 36 points. In other words, you have room for taking up more work in that particular planning folder.




You can also see **2 Cards** at the top of the swim lane, which is the count of the number of cards in the selected planning folder (by default, only open cards are counted).

Click the Show/hide closed artifacts toggle button to see the count of both open and closed artifacts.

#### Show/hide closed artifacts toggle button



 **Note:** You must have set the **Points Capacity** while creating planning folders to be able to see the current points capacity against the planned points capacity. For more information, see [Create a planning folder](#) on page 28.

4. Click the swim lane refresh button to refresh individual swim lanes.

#### Swimlane refresh button




5. To know more about an artifact, click the artifact ID link. The **View Artifact** page is displayed.
6. By default, closed artifacts are not shown in swim lanes. Click the Show/hide toggle button to show or hide closed artifacts in swim lanes.

#### Show/hide closed artifacts button



7. To expand an artifact that has child artifacts, click the "+" symbol in the artifact card. The first, second and third-level child artifacts are displayed.

When you expand an artifact, the parent and child artifacts are visually outlined by means of colored boxes as shown in the following screenshot.



### Work with artifact cards

In the Planning Board, artifact cards are displayed in ranked sequence within each swim lane. Using Planning Board, you can quickly add new cards (quick add), edit cards, move cards within a swim lane to rank them and move cards between swim lanes to reassign them to other planning folders.


### To quickly add artifacts:

You can add artifacts (quick add) in select planning folders with minimum required artifact information.

#### **Important:**

- While you quickly add artifacts with data for just three fields such as the tracker type, title and description, the artifacts are, however, saved with default values for other required fields, which you may choose to update later.
- You cannot add artifacts from the Planning Board if a work flow is configured for the meta status change from **New** to **Open**.

1. Select a planning folder in one of the swim lanes. The quick add (+) button at the top of the swim lane is enabled.

 **Note:** To be able to add artifacts in Planning Board, you need to have the required permissions.

#### Quick add button




2. Click the quick add (+) button. The **Create Artifact** window appears.

#### Create Artifact window

3. Select a tracker type from the **TRACKER** drop-down list.
4. Type a title and description. You can select the **USE TITLE AS DESCRIPTION** check box to copy the title you type to the **DESCRIPTION** text box.



**Artifacts support @mentions:** Artifact description and comments now support @mentions and users called out via @mentions are added to the monitoring list. Include usernames with "@" as prefix (for example, @mhippard) to add users to the monitoring list.


 **Note:** Users called out via @mentions must have "Artifact View" permission to be added to the monitoring list.

**5. Click Submit.**

The artifact is added as the bottommost artifact in the swim lane (planning folder). You may choose to rank the artifact (if you have rank permission) or update the artifact with more meta data later. To update the artifact, click the artifact ID link.

**To rank a card:**


**6. If you have the requisite permission, you can drag and drop a card above or below the other cards within a swim lane to position it in the order you want the artifact addressed.**

 **Note:** The positioning of cards is retained when you switch between the list and board views (planning, task and kanban).

For related information on ranking artifacts in the list view, see [Reorder the contents of a planning folder](#) on page 176.

**To reassign (move) a card:**

**7. In the course of planning your release, you may want to assign an artifact from the product backlog to a specific sprint or team, or move an artifact from one sprint to another. To do this, just drag and drop the artifact card from one swim lane to the other.**

 **Note:** When you drop the card, its rank is relative to its position in the new swim lane. However, if a child artifact is moved to a swim lane (planning folder) where its parent is already present, the child artifact is filed under its parent and its rank is relative to the parent's position.

To be able to move cards between swim lanes, you need to have the permission to update planning folders and rank artifact cards.

**To edit an artifact using the artifact ID link:**

**8. Click the artifact ID link.**




The **View Artifact** page appears.

**9. Make the necessary changes and click Update to update the artifact and return to the Planning Board.**

**To quickly edit an artifact:**

**10. Click the edit artifact button in a card. The Edit Artifact window appears.**

 **Note:** To be able to edit artifacts, you need to have the required permissions.

**Edit artifact button**




**Edit Artifact window**

11. You can edit the following in the **Edit Artifact** window:

- **Artifact title**
- **Artifact description**
- **Primary and secondary attributes**

You have two expandable frames, **Primary Attributes** and **Secondary Attributes**, that list fields you can edit. Primary attributes include Priority, Status, Assigned To, Team, Estimated Effort, Remaining Effort, Actual Effort and Points. The expandable frames and the fields listed in them depend on your application configuration that you set in Tracker Settings. For example:

- Configurable fields which are enabled. For more information, see [Enable or disable fields](#) on page 53.
- User-defined fields (flex fields) with the field type Text Entry and Single Select and set as **Required** (Tracker Settings > Tracker Administration > Create Field) or workflow configured. For more information, see [Create your own tracker fields](#) on page 54 and [Create a tracker workflow](#) on page 55.

 **Note:** You cannot change the status of an artifact in the Planning Board view, if 'Assigned To' is set as a mandatory field for status change per your tracker work flow settings (Tracker Settings). However, you can edit the status of an artifact in the List view with the same 'Assigned To' work flow constraint.


By default, the following field types which you select while creating a user-defined field (Tracker Settings), are non-editable from the planning board's **Edit Artifact** window: Multiple select, Date picker and User picker.

#### Non-editable field types

- **Comments**

In this expandable frame, in addition to adding comments, you can view the last five comments for an artifact.

**Artifacts support @mentions:** Artifact description and comments now support @mentions and users called out via @mentions are added to the monitoring list. Include usernames with "@" as prefix (for example, @mhippard) to add users to the monitoring list.

 **Note:** Users called out via @mentions must have "Artifact View" permission to be added to the monitoring list.

12. After editing an artifact, click **Update**.


### *Team view in Planning Board*

The Team view in Planning Board provides an option to filter artifacts for a specific team within the selected planning folder.

You can select the same planning folder in all three swimlanes. This enables an efficient and faster way of planning a sprint simultaneously for three different teams residing in the same planning folder. You can also select the same planning folder and same team in all three swimlanes for easy handling of artifact cards.


### **View artifact cards:**

1. Click **TRACKERS** from the **Project Home** menu.
2. In the **List Trackers, Planning Folders and Teams** page, click **PLAN**. A Planning Board for the current project context is displayed.
3. Select a planning folder from the drop-down list.

 **Note:** You can select the same planning folder in all three swimlanes.


4. View the artifact cards based on your requirement:

- Select a team from the **Select a Team** drop-down list to view artifacts of a specific team.
- Select **None** to view unassigned artifacts.
- Select **All artifacts** to view all the artifact cards residing in the selected planning folder.

 **Note:** If there are no teams created for a project, then the **Select a Team** drop-down list is not displayed at all.

### **Assign artifacts to a team:**

5. You can select an unassigned artifact from one swimlane and drag it to another where you have a team selected.

 **Note:** Whenever you move an artifact card from one swimlane to another, all its child artifacts (if any) are also moved along with it.

### **Manage tasks with Task Board**

Task board is an important tool in the Agile process. It helps the team to focus on the work at hand in the current sprint and feed progress data back into the system.

The Task Board is not a sprint-planning tool; it is more of a sprint-tracking tool. Use the [Planning Board](#) for planning the stories that will be dealt-with in the sprint. During the sprint, team members can use the Task Board to break down the stories into tasks and track them to completion.

For more information, see [What are Planning, Task and Kanban boards?](#) on page 991

### *Set up the Task Board*

TeamForge project administrators can configure one Task Board per project. Once that is done, project members can use the Task Board.

- During Task Board configuration:
  - Select one or more backlog<sup>#</sup> trackers.
  - Select a task tracker. You can select only one task tracker for a project.
  - Select at least two and up to seven statuses. The number of task swimlanes in the Task Board is equal to the number of statuses you select.
  - Select the unit of effort for backlog and task trackers.

---

<sup>#</sup> Backlog trackers are tracker items such as epics, stories, defects and so on, for which tasks can be created.

- For more information, see [What are Planning, Task and Kanban boards?](#) on page 991

1. Click **TRACKERS** from the **Project Home** menu.
2. In the **List Trackers, Planning Folders and Teams** page, click **Track**.



A Task Board for the current project context is displayed with the following message.

Task Board is not configured for the project.

3. Click the Task Board configuration icon.



The **Settings** window appears.

4. Select one or more backlog trackers and a task tracker.

Trackers	Backlog	Task	Trackers
<input checked="" type="checkbox"/>	<input type="checkbox"/>		Epics
<input checked="" type="checkbox"/>	<input type="checkbox"/>		Stories
<input checked="" type="checkbox"/>	<input type="checkbox"/>		Tasks
<input type="checkbox"/>		<input checked="" type="checkbox"/>	Tests
<input checked="" type="checkbox"/>	<input type="checkbox"/>		Defects
<input type="checkbox"/>	<input type="checkbox"/>		Impediments

Buttons: Cancel, Next



**Note:** You can always revisit your backlog and task tracker settings for Task Board and change them if required at any future point in time.

5. Click **Next**.
6. Select at least two and up to seven statuses from the **Available** list box. Press and hold the **Ctrl** key to select multiple items.
7. Click the forward arrow to add the selected statuses to the **Selected** list box.

Available: Open, In Progress, Passed, Failed

Selected: (Empty)

Buttons: Cancel, Previous, Next



**Note:**

- You can always revisit your status settings for Task Board and change them if required at any future point in time.
- Use the forward and backward arrows to move statuses to and fro the **Available** and **Selected** list boxes.
- Use the up or down arrows to move statuses up or down within the **Selected** list box.

8. Click **Next**.

9. Select the task effort field and backlog item size field.



**Note:** If the effort field is disabled for the backlog and task trackers, it is not possible to select the task effort and backlog item size fields and the value is set to 'None'.

10. Click **Finish**.

The following success message appears: `Task Board preferences saved successfully.`

You have now successfully configured the Task Board for your project. Project members can start using the Task Board to manage their tasks.

*Use the Task Board*

During a sprint, TeamForge project members can use the Task Board to view tasks, create tasks for backlog items, edit tasks and drag and drop tasks across swimlanes as they progress.

With the Task Board, you can:

- View backlog items of a selected planning folder in the **Backlog Items** swimlane, which is the leftmost swimlane of the task board.
- View tasks (belonging to backlog items) pinned to swimlanes based on the status.
- Edit backlog items and tasks by clicking the artifact ID link.
- Move tasks (cards) from one swimlane to the other as tasks progress.
- Add new tasks for backlog items.

1. Click **TRACKERS** from the **Project Home** menu.

2. In the **List Trackers, Planning Folders and Teams** page, click **TRACK**.



A Task Board for the current project context is displayed.


3. Select a planning folder from the drop-down list.

The Task Board is populated with backlog items and their tasks.

- If a selected planning folder has subfolders, the artifacts within the subfolders are not displayed.
- Cards have a color-coded bar to visually identify the priority. In addition, the open and closed cards' background is color-coded to uniquely identify the status.
- For each card, the artifact ID, title, priority, assigned to and effort (in terms of 'Points' for backlog items and 'Hours' for tasks) are displayed. Status information is shown on backlog cards only. Relevant tooltips appear when you hover your mouse over these data elements. You can also hover your mouse over the tracker icon in backlog and task cards to know the tracker name.


**Add new tasks**

4. Click the '+' icon.




The **Create Task** window appears.

5. Type a title and description and click **Submit**.

 **Tip:** To use the title you type as the description, select the **Use Title as Description** check box.

#### To edit a backlog item or task using the task ID link:

6. On the **List Trackers, Planning Folders and Teams** page, select the planning folder.
7. From the listed backlog items and tasks, click the artifact ID link you want to edit.




The **View Artifact** page appears.

8. Modify the backlog item or task and click **Update** to update the artifact and return to the Task Board.

#### To quickly edit a task

9. Click the edit artifact button in a task. The **Edit Task** window appears.

 **Note:** To be able to edit tasks, you need to have the required permissions.

#### Edit task button



#### Edit Task window

Edit Task: artf6211
✕

Title \*

Description \*

**Primary Attributes**

Priority	4 - Low ▼	Estimated Effort	0	Hours ▼
Status *	Not Starte ▼	Remaining Eff...	0	Hours ▼
Assigned ...	None ▼	Actual Effort	0	Hours ▼
Points	0			

**Secondary Attributes**

Comments

Cancel
Submit

10. You can edit the following on the **Edit Task** window:

- **Artifact title**
- **Artifact description**
- **Primary and secondary attributes**

You have two expandable frames, **Primary Attributes** and **Secondary Attributes**, that list fields you can edit. Primary attributes include Priority, Status, Assigned To, Estimated Effort, Remaining Effort, Actual Effort, Points and Team. The expandable frames and the fields listed in them depend on your application configuration that you set in Tracker Settings. For example:

- Configurable fields which are enabled. For more information, see [Enable or disable fields](#) on page 53.
- User-defined fields (flex fields) with the field type Text Entry and Single Select and set as **Required** (Tracker Settings > Tracker Administration > Create Field). For more information, see [Create your own tracker fields](#) on page 54.
- Fields which are workflow configured. For more information, see [Create a tracker workflow](#) on page 55.

By default, the following field types, which you select while creating a user-defined field (Project Admin > Tracker Settings), are non-editable from the task board's **Edit Task** window: Multiple select (**Multiple Select**), Date picker (**Select Date**) and User picker (**Select User(s)**).


- **Comments**

In this expandable frame, in addition to adding comments, you can view the last five comments for an artifact.


11. After editing an artifact, click **Update**.

#### Move tasks across swimlanes

12. To move tasks across swimlanes, drag a task card from the source swimlane and drop it on the destination swimlane.

 **Note:** You need to have task tracker 'edit' permission to move cards across swimlanes.

#### View both 'open' and 'closed' backlog items

13.  **Note:** When you select a planning folder, the Task Board shows only 'open' backlog items by default.

Click the 'Show/hide closed artifacts' toggle button to view both 'open' and 'closed' backlog items.

#### Show/hide closed artifacts toggle button



#### Refresh the Task Board

14. Click the following 'Refresh' button to refresh the Task Board.



#### Auto assign task


If a task is not assigned to anyone, that is, if it is assigned to **None**, then it can be assigned to the user who has logged in using the **Auto Assign Task to Me** feature. This check box appears only if:

- The Task Board is configured.
- The user who logged in has the edit permission to the task tracker.
- The project is not locked.

15. Select the **Auto Assign Task to Me** check box and drag the artifact from one swimlane to another (from one status to another) and the artifact is automatically assigned to you.

#### Position artifact cards

16. If you have the requisite permission, you can drag and drop a card above or below the other cards within a swim lane to position it in the order you want the artifact addressed.

 **Note:** The positioning of cards is retained when you switch between the list and board views (planning, task and kanban).

For related information on ranking artifacts in the list view, see [Reorder the contents of a planning folder](#) on page 176.

#### Team view in Task Board

The Team view in Task Board provides an option to filter backlog items for a specific team within the selected planning folder.

This team's view of artifacts applies only to backlog items. When you filter and view the backlog items for a specific team, the tasks (if any) within the filtered backlog items are also displayed in the relevant status swimlanes. However, you cannot filter tasks alone for a team.

You can move the artifact cards between swimlanes. If the artifacts are workflow configured, appropriate error messages are displayed when a workflow rule is broken.

1. Click **TRACKERS** from the **Project Home** menu.
2. In the **List Trackers, Planning Folders and Teams** page, click **TRACK**.  
A Task Board for the current project context is displayed.
3. Select a planning folder from the drop-down list. All the backlog items and their tasks in the selected planning folder are displayed.
4. Select a team from the **Select a Team** drop-down list.





**Note:** If there are no teams created for a project, then the **Select a Team** drop-down list is not displayed at all.

All the backlog items and their tasks in the selected team are displayed.

### Manage your project workflow with Kanban Board

The Kanban Board is a project management tool, which gives you a snapshot of the work items, in which states they are, how they are progressing and if there is any bottleneck to be cleared for a smooth delivery of the product on time.

#### Set up Kanban Board

As a project administrator, you can configure your Kanban Board based on your project requirements.

Create Kanban Board states for the different stages of your project development process, specify workflow constraints for each state and map them to your tracker statuses.

1. Click **TRACKERS** from the **Project Home** menu.
2. In the **List Trackers, Planning Folders and Teams** page, click **Kanban**.



The following message is displayed prompting you to configure a Kanban Board for the current project.

Kanban Board is not configured for the project.

3. Click the Kanban Board configuration icon to open the **Settings** window.



**Note:** Use the configuration icon when you create a Kanban Board for the first time or when you want to modify the settings of your current Kanban Board.

Alternatively, you can use **Manage Boards** to create a new board.



For more information about its usage, see [Manage Boards](#).

4. Give a name to the new Kanban Board.
5. Select the trackers whose artifacts you want to view on your Kanban Board. (These are the trackers you created for your project. For more information on how to create a tracker, see [Create a tracker](#) on page 27.)
6. Click **Next**.

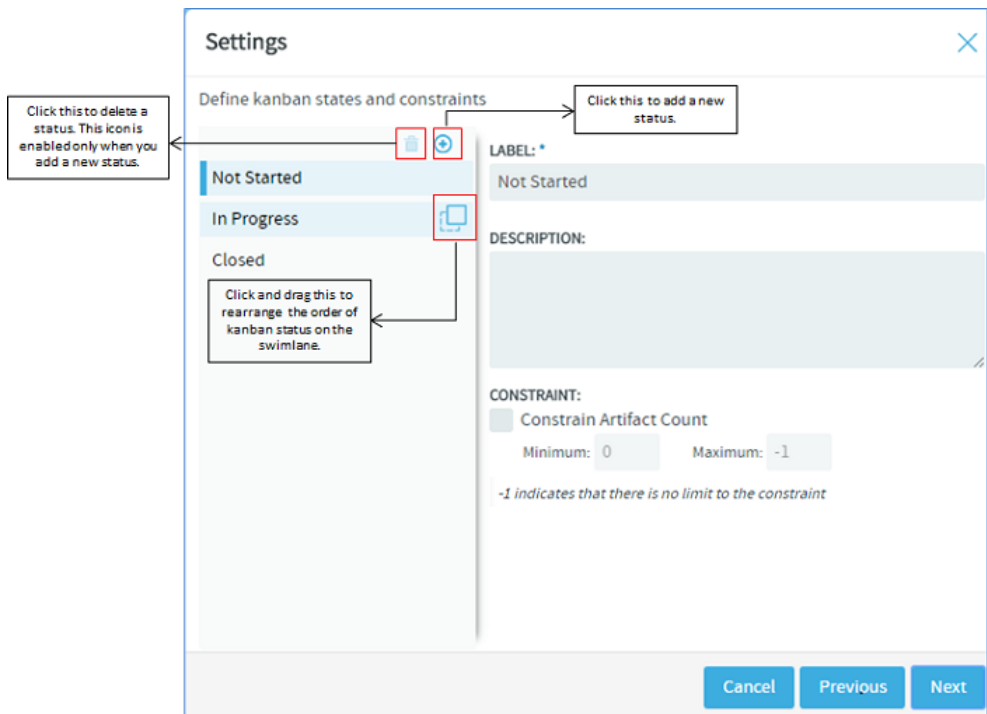
#### Kanban states and constraints

A kanban state is the status of a work item in a value stream (swimlanes). The constraints or the limits that you specify here dictates the behaviour of the board. These are checkpoints which eventually help you track the progress of a work item, identify bottlenecks and fix them. The default kanban statuses are 'Not Started', 'In Progress' and 'Closed'.

You can create your own kanban statuses, delete any state including the default one and rearrange the order in which they would appear as swimlanes on your Kanban Board.



**Note:** If there are only three kanban statuses, the **Delete Status** icon is disabled disallowing you to delete any more because you must maintain a minimum of three statuses.



7. Create and configure kanban statuses.

- a) Click the **Add Status** icon. You can have a maximum of 20 statuses.
- b) Give a label and description to the new status.

**Note:** The kanban status label can have a maximum of 64 characters and the description, a maximum of 256 characters. Alpha numeric and special characters are allowed.

- c) Select the **Constrain Artifact Count** check box to specify the minimum and maximum workflow constraints, that is, specify the minimum and maximum number of work items (artifact cards) that ought to be present in each status. You may specify these values based on the number of work items in the queue versus the number of resources available.

Remember the following when you set the minimum or maximum value to -1.

Constraint value	Logic	Example
Minimum value = -1	This is treated as zero. Therefore, though the maximum value should be greater than the minimum value, the following value range is allowed:  Minimum = -1  Maximum = 0 and above	Let us assume that you are in the last iteration of your release. As a project manager, you would not want to see any artifact in the 'Impeded' status. So you set the minimum to -1 and maximum to 0, which translates to zero artifacts. When this constraint is applied, your Kanban Board flags a violation, if artifacts show up in the 'Impeded' status, thereby drawing your attention to address the issue immediately.
Maximum value = -1	This indicates there is no limit to the maximum value.	You may choose to set -1 as the maximum constraint for the 'Closed' status because you would want to see as many closed artifacts as possible.



**Important:** These constraints are applicable only at the planning folder level and not at the team level because the Kanban Board is expected to give an overall visibility to the work items for an iteration or release within a planning folder and not just within a specific team.

### Map kanban states with tracker status

Next, make a logical mapping of the kanban states with the tracker statuses to view those tracker artifacts on your Kanban Board. For example, to view defects, you must first make a meaningful mapping of the kanban states with the defect tracker status(es). The mapped kanban states appear as swimlanes on the Kanban Board.

The trackers (you had selected in [Step 5](#)) and their statuses appear on the left. All the kanban statuses appear at the top. The tracker statuses are unmapped by default.

- You need to map a minimum of 3 kanban statuses.
- You can map a single kanban status to one or more tracker statuses, but you cannot map a single tracker status to more than one kanban status.

Settings


Map the kanban states to tracker statuses.

Minimum of three Kanban statuses are required to be mapped

TRACKER/STATUS	UNMAPPED	NOT START...	IN PROGRE...	CLOSED
<b>Epics</b>				
Under Consideration	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
In Progress	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Completed	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Rejected	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<b>Stories</b>				
Under Consideration	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
In Development	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Ready for UAT	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Ready for QA	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Cancel Previous Finish

8. Click the appropriate check box. For example, for an Epic tracker, map 'Completed' and 'Rejected' tracker statuses with the Kanban statuses 'Closed' by selecting the appropriate check boxes as shown in the following screenshot.




9. Click **Finish** to complete the process.

### Manage Boards

You can create and manage multiple kanban boards per your requirement. As a project administrator, you may want to create multiple boards at the iteration level, at each unit level, (Dev, QA and so on), at the hierarchical level or at each project team level, depending upon the needs of the users. For example, to get the big picture and make executive decisions you may want to create a kanban board at the management level; at the same time, at a project member level, you can create a board to focus on the immediate tasks to be completed.

10. Use the various icons available in **Manage Boards** to maintain multiple kanban boards in a project.

To do this...	Click this icon
To use <b>Manage Boards</b> .	
To create a new board	
To edit a board	
To delete a board	
<b>Note:</b> You cannot delete the default Kanban Board. When you delete your current active board, the default one shows up automatically.	

To do this...	Click this icon
To set a default board	

### Use Kanban Board

After the kanban board is configured, you can use it to view, plan and track work items for the selected planning folder or the selected team within that planning folder.

1. Click **TRACKERS** from the **Project Home** menu.
2. In the **List Trackers, Planning Folders and Teams** page, click **Kanban**.




3. Select the required kanban board from **Manage Boards** (available only for project administrators). If you are a project member, select your kanban board from the **Select Board** list. The selected kanban board for the current project context is displayed.
4. Select the planning folder for which you want to view the status of the artifacts.


If your project contains teams, the **Select a Team** drop-down list appears. From this drop-down, select **All artifacts** or if you want to view artifacts specific to a team within the selected planning folder, select the relevant team.

Depending upon your selection, artifacts pertaining to the mapped trackers are displayed in the appropriate swimlane.

- You can view a maximum of 5 swimlanes at a time on your Kanban Board. If there are more, use the carousel scroll to slide across the swimlanes. You can expand or collapse each swimlane. When there are no results to display in a swimlane, the **Expand/Collapse** button does not appear.
- Based on the configuration values, if the constraints are violated, the relevant status headers are highlighted appropriately. A violation of minimum constraint is highlighted in orange indicating that resources are underutilized whereas that of maximum constraint is highlighted in red indicating resources being overloaded and bottlenecks to be fixed. For more information on how a kanban board is configured, see [Set up Kanban Board](#) on page 193.



- Each swimlane header displays the status label, the total number of artifacts within the selected planning folder, and the minimum and maximum constraints that were configured for the kanban state.



In the above scenario, for the 'In Progress' status:

'3' is the total number of artifacts in the 'In Progress' status within the selected planning folder.

'2' indicates the minimum constraint and '\*' indicates the maximum constraint '-1'. (-1 as the minimum constraint translates to 0 whereas -1 as maximum indicates there is no limit and so represented by an \*.)

- 👉 **Important:** These values and constraint validation are planning folder specific and not team specific. For example, in the above scenario, the total number of artifacts is the overall count of artifacts which are in the 'In Progress' status within the selected planning folder and not within the selected team.
- Each artifact card displays its points (story points) or estimated effort if these fields are enabled in your tracker settings. If both are enabled, only the estimated effort is displayed hovering the mouse on which the artifact's points is displayed. Similarly, if all the effort fields (estimated, remaining and actuals) are enabled, you can view them by hovering the mouse on the estimated effort.
- Cards have a color-coded bar to visually identify the priority. In addition, the open and closed cards' background is color-coded to uniquely identify the status.
- A parent artifact card displays the count of its child artifacts (**Show/Hide child artifacts** button). Click this button to view or hide the child artifacts. If there are more than 5 child artifacts, when you expand the parent artifact, you will see pagination arrows ('Previous'/Next') at the bottom of the parent artifact card.

The screenshot displays a Kanban board interface. At the top, an orange header reads 'IN PROGRESS 4/[5..9]'. Below this, a parent artifact card for 'artf1002' is shown with '100 Points' and a circular button containing the number '6'. The card title is 'Kanban board configuration' and it shows 'Open: 6, Closed: 0'. A list of child artifacts follows: 'artf1009' (Code browser set up), 'artf1010' (Task board configuration), 'artf1014' (Kanban board swimlanes), 'artf1015' ([sample] Story Six), and 'artf1085' (Doc task). At the bottom of the parent card are left and right navigation arrows. Below the parent card, an individual child artifact card for 'artf1014' is visible, showing '20 Hours' and the title 'Kanban board swimlanes'. Red boxes highlight the '6' button, the 'artf1014' card in the parent list, and the 'artf1014' card below it.

- If the tracker types of both the parent and child artifacts have been mapped with the kanban statuses, the child artifacts appear within their parent artifact card and also as an individual card. In the above scenario, Epic (parent artifact) and Story (child artifact) trackers have been mapped to the kanban status 'In Progress'. Story 'artf1014' is the child artifact of the epic 'artf1002'. So you will see 'artf1014' within the parent artifact card as well as outside of it as an individual artifact card.
- When a child artifact is closed, a 'Closed' tag appears next to the child artifact ID within the parent artifact card.


## Edit artifacts

Use Kanban Board to edit an artifact using the Edit icon on the artifact card or move artifacts from one status to another appropriately. Based on the status changes you make, the swimlane headers get updated appropriately.

### Remember:


- You can edit an artifact only if you have the tracker edit permission. Otherwise, you can only view it.

**Artifacts support @mentions:** Artifact description and comments now support @mentions and users called out via @mentions are added to the monitoring list. Include usernames with "@" as prefix (for example, @mhippard) to add users to the monitoring list.

 **Note:** Users called out via @mentions must have "Artifact View" permission to be added to the monitoring list.

5. When you drag the artifacts from one swimlane to another, note the following validations:


- The changes you see in swimlane headers do not apply to any specific team but takes into account the total count of artifacts within the selected planning folder.
- You cannot move an artifact card to a status that is not mapped to that particular tracker type. For example, a 'Ready for QA' kanban status may not have been mapped to any of the epic tracker statuses. So when you attempt to move an epic to that status, you will get an `Invalid status configuration` error.
- You can move a parent artifact to the 'Closed' status only if all its child artifacts are closed. However, a child artifact is independent of its parent with regard to the status change, that is, when you move a parent artifact from one status to another, the status of the child may still remain unchanged. For example, an epic may have many stories, the status of some may change from 'Not Started' to 'In Progress' whereas some may not. So when the status of an epic as one single unit may change, it need not apply to all of its child artifacts.
- When your move violates a constraint (minimum or maximum), a warning is displayed; but you can still move the artifacts.
- If a kanban status is mapped to more than one tracker status, when you move an artifact, you have an option to choose a status as shown in the following screen shot.



## Create artifacts and child artifacts


Using the Create Artifact icon on the top right of your kanban board, you can quickly create an artifact. The Create Artifact icon appears only when you configure a kanban board and select a planning folder.

Similarly, using the Create Child Artifact icon available on the artifact card, you can quickly create a child artifact for any artifact card displayed on Kanban Board.

 **Note:** The Create Child Artifact icon does not appear on closed artifact cards.

You can create an artifact or a child artifact only if you have the required tracker permission.

**Artifacts support @mentions:** Artifact description and comments now support @mentions and users called out via @mentions are added to the monitoring list. Include usernames with "@" as prefix (for example, @mhippard) to add users to the monitoring list.

 **Note:** Users called out via @mentions must have "Artifact View" permission to be added to the monitoring list.

6. Click the required icon: Create Artifact or Create Child Artifact.



7. Enter the required information in the relevant window and click **Submit**.

- Only trackers configured for the kanban board appear in the Trackers drop-down list.
- While you quickly add artifacts with data for just three fields such as the tracker type, title and description, the artifacts are, however, saved with default values for other required fields, which you may choose to update later. If the default state of the selected tracker is not mapped for the newly created artifact or child artifact, the artifact card does not show up on the kanban board.

### Show / Hide closed cards

8. When you have a large volume of closed cards in a planning folder, you can restrict the number of closed cards you want to view by toggling between the 'Show all closed artifacts' icon:



Or the 'Hide artifacts older than 60 days' icon using which you can hide closed cards older than 60 days. This is the default option:



 **Note:**

- The Show/Hide toggle icon appears only on configured kanban boards.
- Your selection is saved for the subsequent sessions as well.

### Position artifact cards

9. If you have the requisite permission, you can drag and drop a card above or below the other cards within a swim lane to position it in the order you want the artifact addressed.



**Note:** The positioning of cards is retained when you switch between the list and board views (planning, task and kanban).

For related information on ranking artifacts in the list view, see [Reorder the contents of a planning folder](#) on page 176.

### Manage task work flow

As a task administrator, you can control how users move tasks along under various conditions.

#### Alert project members about overdue tasks

If a task is not completed by its proposed end date, you may want to send an automatic reminder to the task's owner, its administrator, or other project members.

You can choose which of the people involved with the overdue task get the alert, and how late the task must be to trigger an alert.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin Menu**, click **Task Manager Settings**.
3. Click the **ALERTS** tab.
4. Specify the number of days after a task's due date that an overdue task alert is sent.
5. Specify the recipients of overdue task alerts.


- **Project Role Members** - All project members who are assigned that role.



**Note:** If you select roles or project members that do not have the task view permission, these project members will receive Permission Denied errors when attempting to access overdue tasks.

- **Specific Project Members** - Click the Search icon to select from a list of project members.
- **Task Creator** - The person who originated the task.



 **Note:** The person to whom the task is assigned receives overdue task alerts automatically.

## 6. Click **Save**.


Overdue task alerts are now enabled for all tasks in the project.

### **Require approval for changing tasks**

During times when changes to tasks carry a high risk, task administrators can prevent users' changes from taking effect until the administrator approves them.

Requiring change approval keeps you informed about changes to task status, dates, dependencies, and other changes that can affect the project plan.

You can select the fields for which you want require change approval. Whenever a user makes a change to a field for which change approval is required, the change request is placed in the **Change Requests** section of the **Task Summary** page, with a link in the **Items Pending My Approval** section of each task administrator's **My Page**.

 **Note:** Changes made by task administrators do not require approval.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin Menu**, click **Task Manager Settings**.
3. Click the **CHANGE APPROVAL** tab.
4. Select the fields for which you want to require change approval.
5. Click **Save**.


Change approval is now in force for all future changes made to the selected fields.

### **Handle a change request**

When change approval is in force, changes made by project members to selected task fields must be approved or rejected by a task administrator.

Change requests are placed in the **Change Requests** section of the **Task Summary** page. Change requests are also displayed in the **Items Pending My Approval** section of each task administrator's **My Page**.

1. Click **TASKS** in the project navigation bar.
2. On the **Task Summary** page, click the **CHANGE REQUESTS** tab.
3. Respond to a change request in one of these ways:
  - If you do not need to view the request details, select the task and click **Approve Selected** or **Reject Selected**.
  - To view the request details or add a comment before responding to the change, click the task name and write your comments, then click **Approve Selected** or **Reject Selected** in the **Change Requests** section of the **Task Details** page.

 **Note:** If you approve a task end date change that impacts dependent successor tasks, you must confirm your changes for each of the successor tasks individually.

You can also approve or reject change requests from the **ITEMS AWAITING MY APPROVAL** section of your **My Page**. In the **Tasks Awaiting Approval** section, click the task name to view the request details, add a comment, and approve or reject the request.

### **Indicate task status with colors**

The status of each task is represented by a color. The status color of a task folder is an accumulation of the status colors of the tasks in the folder.

When a certain percentage of tasks have turned a given status color, the color "rolls up" to the task folder. The status colors of the task folders also "roll up" to the project according to the same criteria.

The percentage in each case is chosen by the project admin.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin Menu**, click **Task Manager Settings**.
3. On the **STATUS ROLLUP** page, set your color-coding criteria.

For example:

- A task folder is represented by red when at least 10 percent of its tasks are red.
- A task folder is represented by green when at least 10 percent of its tasks are green.
- A task folder is represented by gray only when at least 80 percent of its tasks are complete.

When one or more criteria are met simultaneously, the task folder is represented by the most critical of the applicable status values. For example, using the criteria above, if 10 percent of the tasks in a task folder are yellow and 75 percent of the tasks in the same folder are gray, the task folder is represented by yellow.

### Measure task workloads

For help in planning and allocating resources, look at the project's task workloads.

You can measure the tasks assigned to each project member by status, due date, and estimated and actual effort.


1. Click **TASKS** in the project navigation bar.
2. On the **Task Summary** page, click the **WORKLOAD** tab.
3. Find the following information for each project member:
  - The number of assigned tasks, organized by status value.
  - The total number of incomplete tasks.
  - The total estimated and actual effort for all tasks.

### Set the default task calendar

For each task, you can specify whether work is tracked on a five-day or a seven-day calendar.

For each TeamForge project, you can specify a default task calendar. This selection will be the default calendar for new tasks. You can override this value for individual tasks.

Five-day calendars count only Monday through Friday. Seven-day calendars also count Saturday and Sunday.

 **Note:** Changing the default calendar does not update existing tasks.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin Menu**, click **Task Manager Settings**.
3. Click the **DEFAULT TASK CALENDAR** tab.
4. Choose a five-day or seven-day calendar as the default for the project.
5. Click **Save**.

Your default task calendar will be applied to all new tasks.

### Reporting in TeamForge

Generate a report to get a snapshot of what is going on in a project. You can generate reports on data stored in both TeamForge's production (operational) database or datamart. Datamart, also known as the Reporting database, is built by extracting, transforming and loading (ETL) TeamForge's production data to a separate database (datamart) at regular intervals.

 **Important:**

Unless otherwise stated, you must have datamart enabled on your site to create reports in TeamForge. Note that a few *Distribution reports* use data from TeamForge's operational database and a few *Activity Reports* use data from the EventQ's event data store.

You can specify the time at which the reporting data is refreshed from the production database. By default, the extraction takes place daily at 2:30 a.m. in the TeamForge application server's time zone. See *Schedule data extraction for reporting* on page 712.

You can use reports to display data and group relevant information appropriately and specify intervals at which the datamart extracts TeamForge data from the datamart. For advanced reporting options and datamart information, see *Advanced reporting and datamart access* on page 248.

You can also use external reporting tools to connect to the datamart and generate customized reports. See *Datamart access using external tools* on page 259.

## Reporting framework

Here's a list of some of the advantages of the new reporting framework:

- Reporting under one umbrella with a central dashboard of reports.
- Cross-project reporting capability.
- Tracker custom defined fields are in data-mart and can be used to filter.
- All events are captured via EventQ event data-store.
- You can query and write custom reports.
- All data, including event, associations, and traceability data, can be queried through an elegant API for custom reporting and data extraction.
- High charts-based interactive data visualization charts. You can hover over the charts to see data points, click legends in the chart to toggle specific data point in and out of the chart and so on.
- You can drill one level down on some of the activity reports using column charts to get more clarity on the data points of your interest.
- Categorization of reports.
- Improved usability: Hassle free report creation with new widgets for selecting report criteria such as planning folders, trackers and repositories.
- Ability to save 'Public' and 'Private' reports: While 'Public' reports are visible to all project members with view reports permission, 'Private' reports are only for your consumption. You cannot publish 'Private' reports in project pages.

## Reports: Role based access control

In general, TeamForge site and project administrators, and users with tracker and task view permissions can see the **REPORTS** button on the project navigation bar.

- **Object-based access permissions:** If you are a TeamForge user, your ability to create, edit, preview or view reports depends on whether you have permission to access specific objects such as trackers and repositories in TeamForge. For example, to generate a report on the number of SCM commits in a repository, you must have permission to access the repository. Otherwise, a message such as `You do not have sufficient permission to perform this operation` is displayed.
- **Task and Tracker reports (Table reports):** **Task and Tracker View** permission is required to generate task and tracker reports. You must have **View Activity** permission or object-level permission to view Activity reports such as SCM Commits, Build Activity, Build and Test Activity, Artifact Created and Artifact Closed reports.
- **Context-sensitive access permission:** Reports are shown or hidden from users based on the context. For example, a user must have **View Project Page** permission to view reports published on a project home page.
- **Deleting reports:** Site and project administrators can delete reports. In addition, you can delete your own reports.

## Reports available in TeamForge

Reports in TeamForge are grouped under the following categories.

### *Activity reports*


Here's a list of TeamForge activity reports such as the SCM commits, artifact created and closed reports.

These are reports to track activities such as SCM commits, artifact creation and closure, build and test activities and so on. Some of these reports are powered by data from EventQ's event data store.

- 👉 **Note:** You can also drill one level down on some of the activity reports using column charts to get more clarity on the data points of your interest.

### SCM Commits

Shows the number of commits in one or more selected repositories in a given time period.



1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Scm Commits** from **Activity Reports**.
4. Type a report title and description.
5. Select one or more repositories from the **SELECT REPOSITORY(S)** drop-down list.

Click here to Select Repository(s) ▾

Select All
  Select None

Search...

Publishing

svnr

6. Select date range and click **Apply**.

To set the date range, select the start date first and then the end date. Use the month/year navigation arrows to select the month/year you want. If you don't have the exact dates at hand, you can also select one of the time periods from the **Select Range: Custom** drop-down list.

Select date range

January 2015							February 2015						
Su	Mo	Tu	We	Th	Fr	Sa	Su	Mo	Tu	We	Th	Fr	Sa
				1	2	3	1	2	3	4	5	6	7
4	5	6	7	8	9	10	8	9	10	11	12	13	14
11	12	13	14	15	16	17	15	16	17	18	19	20	21
18	19	20	21	22	23	24	22	23	24	25	26	27	28
25	26	27	28	29	30	31							

Select range:

Custom ▾


7. Select a chart type such as 'Line' or 'Bar' or 'Column' from the **CHART DISPLAY TYPE** drop-down list.
8. Select report visibility: Public or Private.
9. Click **Preview**.

**10. Click Create.**


The report is created and the **View Report** page appears.

**Print or download charts**

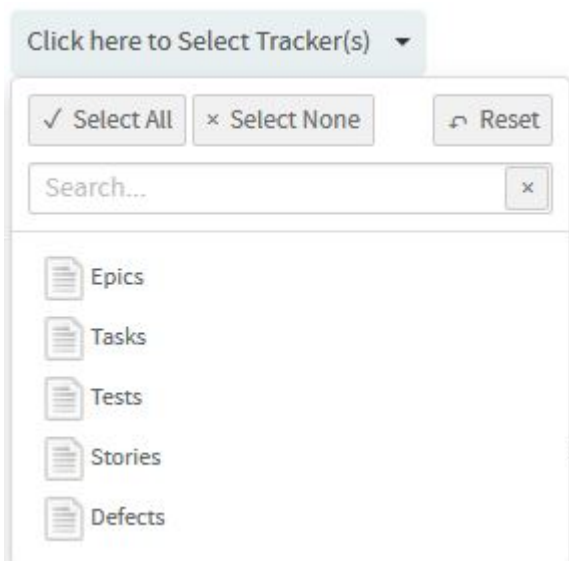
You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.

**11. Click Back to Reports List** to go back to the Reports dashboard.**Artifact Created**

Shows the number of artifacts created in one or more selected trackers in a given time period.




1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Artifact Created** from **Activity Reports**.
4. Type a report title and description.
5. Select one or more trackers from the **SELECT TRACKER(S)** drop-down list.



6. Select date range and click **Apply**.

To set the date range, select the start date first and then the end date. Use the month/year navigation arrows to select the month/year you want. If you don't have the exact dates at hand, you can also select one of the time periods from the **Select Range: Custom** drop-down list.



7. Select a chart type such as 'Line' or 'Bar' or 'Column' from the **CHART DISPLAY TYPE** drop-down list.

8. Select report visibility: Public or Private.


9. Click **Preview**.

10. Click **Create**.

The report is created and the **View Report** page appears.

### Print or download charts


You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.



11. Click **Back to Reports List** to go back to the Reports dashboard.

### Artifact Closed

Shows the number of artifacts closed in one or more selected trackers in a given time period.



1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Artifact Closed** from **Activity Reports**.
4. Type a report title and description.
5. Select one or more trackers from the **SELECT TRACKER(S)** drop-down list.

Click here to Select Tracker(s) ▾


✓ Select All   × Select None   ↺ Reset

Search... ×

- 📄 Epics
- 📄 Tasks
- 📄 Tests
- 📄 Stories
- 📄 Defects

6. Select date range and click **Apply**.

To set the date range, select the start date first and then the end date. Use the month/year navigation arrows to select the month/year you want. If you don't have the exact dates at hand, you can also select one of the time periods from the **Select Range: Custom** drop-down list.

Select date range 


January 2015							February 2015						
Su	Mo	Tu	We	Th	Fr	Sa	Su	Mo	Tu	We	Th	Fr	Sa
				1	2	3	1	2	3	4	5	6	7
4	5	6	7	8	9	10	8	9	10	11	12	13	14
11	12	13	14	15	16	17	15	16	17	18	19	20	21
18	19	20	21	22	23	24	22	23	24	25	26	27	28
25	26	27	28	29	30	31							

Select range:

7. Select a chart type such as 'Line' or 'Bar' or 'Column' from the **CHART DISPLAY TYPE** drop-down list.
8. Select report visibility: Public or Private.
9. Click **Preview**.
10. Click **Create**.  
The report is created and the **View Report** page appears.

**Print or download charts**


You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.



11. Click **Back to Reports List** to go back to the Reports dashboard.

**Activity Over Time**

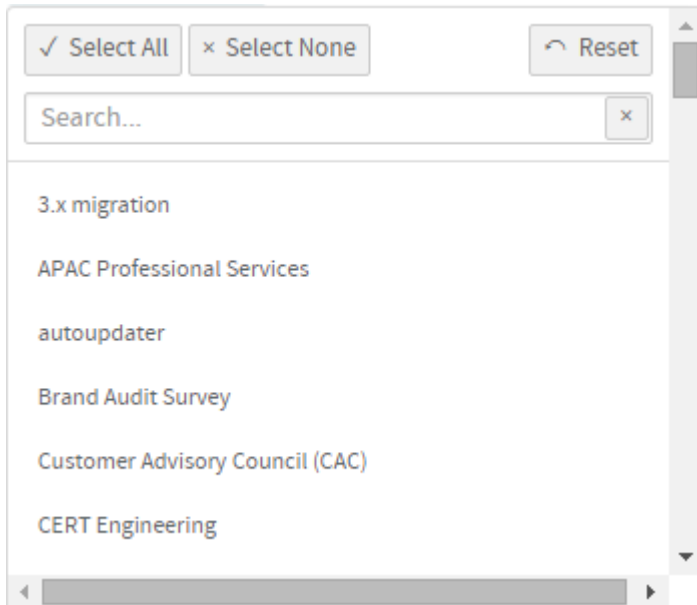
Shows the count of various project activities over a period of time such as commits, reviews, builds, and so on. This report is powered by data from the EventQ's event data store.



1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.



3. Select **Activity Over Time** from **Activity Reports**.
4. Type a report title and description.
5. Select one or more projects from the **PROJECT(S)** drop-down list.




6. Select the number of days from the **LAST N DAYS** drop-down list.
7. The default **Display Type** is **EventsDual**.
8. Select report visibility: Public or Private.
9. Click **Preview**.
10. Click **Create**.

The report is created and the **View Report** page appears.

#### **Print or download charts**


You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.



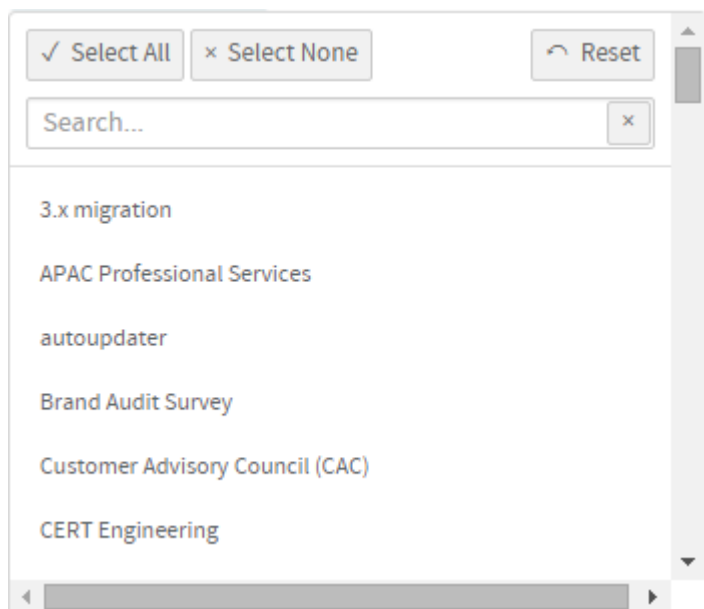
11. Click **Back to Reports List** to go back to the Reports dashboard.

#### Build Activity

Shows the build activity over a period of time including the count of builds and the average build duration. This report is powered by data from the EventQ's event data store.



1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Build Activity** from **Activity Reports**.
4. Type a report title and description.
5. Select one or more projects from the **PROJECT(S)** drop-down list.




6. Select the number of days from the **LAST N DAYS** drop-down list.

7. The default **Display Type** is **EventsDual**.
8. Select report visibility: Public or Private.
9. Click **Preview**.
10. Click **Create**.  
The report is created and the **View Report** page appears.

**Print or download charts**


You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.



11. Click **Back to Reports List** to go back to the Reports dashboard.

**Build Activity by Project**


Shows the number of builds over a period of time. This report is powered by data from the EventQ's event data store.



You can drill down this report.




1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Build Activity by Project** from **Activity Reports**.
4. Type a report title and description.
5. Select one or more projects from the **PROJECT(S)** drop-down list.



6. Select the number of days from the **LAST N DAYS** drop-down list.
7. The default **Display Type** is **EventsDrilldown**.
8. Select report visibility: Public or Private.
9. Click **Preview**.
10. Click **Create**.  
The report is created and the **View Report** page appears.

#### Print or download charts


You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.



11. Click **Back to Reports List** to go back to the Reports dashboard.

#### Build and Test Activity


Shows the total number of CI tests—including passed, failed, and ignored—performed for each build over a period of time. Also shows the total test duration for each build. This report is powered by data from the EventQ's event data store.



1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Build and Test Activity** from **Activity Reports**.
4. Type a report title and description.
5. Select one of the projects from the **PROJECT(S)** drop-down list.
6. Select the number of days from the **LAST N DAYS** drop-down list.
7. The default **Display Type** is **EventsDual**.
8. Select report visibility: Public or Private.
9. Click **Preview**.
10. Click **Create**.  
The report is created and the **View Report** page appears.

**Print or download charts**


You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.



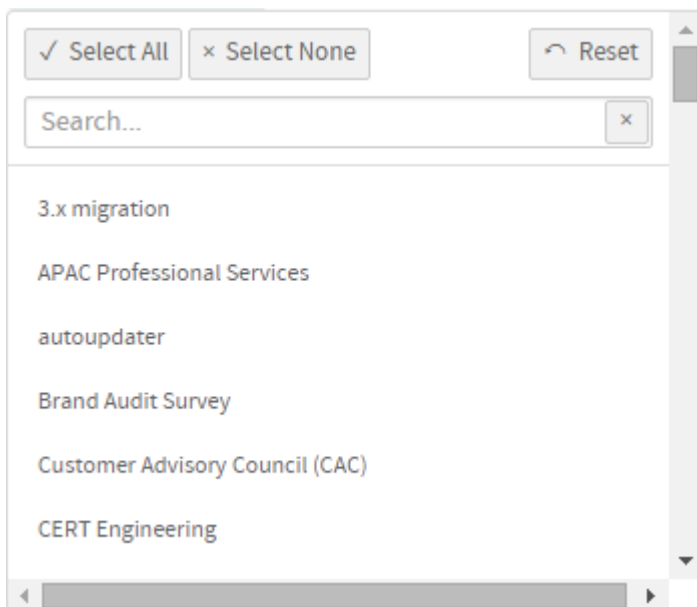
11. Click **Back to Reports List** to go back to the Reports dashboard.

**Test Activity**

Shows the count of tests over a period of time including the average build duration. This report is powered by data from the EventQ's event data store.



1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Test Activity** from **Activity Reports**.
4. Type a report title and description.
5. Select one or more projects from the **PROJECT(S)** drop-down list.




6. Select the number of days from the **LAST N DAYS** drop-down list.
7. The default **Display Type** is **EventsDual**.
8. Select report visibility: Public or Private.
9. Click **Preview**.
10. Click **Create**.

The report is created and the **View Report** page appears.

#### **Print or download charts**


You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.




11. Click **Back to Reports List** to go back to the Reports dashboard.

#### Commit Activity

Shows the count of daily commits over a period of time. You can group this report by the review status, if required. This report is powered by data from the EventQ's event data store.



1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Commit Activity** from **Activity Reports**.
4. Type a report title and description.
5. Select one or more projects from the **PROJECT(S)** drop-down list.




6. Select the number of days from the **LAST N DAYS** drop-down list.



7. By default, the **GROUP BY: Review Status** check box is selected. Clear this check box if you do not want to group the report by review status.
8. The default **Display Type** is **EventsDual**.
9. Select report visibility: Public or Private.
10. Click **Preview**.
11. Click **Create**.  
The report is created and the **View Report** page appears.

#### Print or download charts

You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.




12. Click **Back to Reports List** to go back to the Reports dashboard.

#### Commits by User

Shows the number of commits by each user over a period of time. This report is powered by data from the EventQ's event data store.




1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Commits by User** from **Activity Reports**.
4. Type a report title and description.
5. Select one or more projects from the **PROJECT(S)** drop-down list.



6. Select the number of days from the **LAST N DAYS** drop-down list.
7. By default, the **GROUP BY: Review Status** check box is selected. Clear this check box if you do not want to group the report by review status.
8. **EXCLUDE USERS (COMMA-SEPARATED LIST)**: Optionally, you can type a list of comma-separated user names that you want to exclude from the report.
9. The default **Display Type** is **EventsDual**.
10. Select report visibility: Public or Private.
11. Click **Preview**.
12. Click **Create**.  
The report is created and the **View Report** page appears.

#### Print or download charts


You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.




13. Click **Back to Reports List** to go back to the Reports dashboard.

#### File Changes Over Time

Shows the file change statistics over time, including the total count of changes and the number of add, delete, copy and modify operations. This report is powered by data from the EventQ's event data store.



1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **File Changes Over Time** from **Activity Reports**.
4. Type a report title and description.
5. Select one or more projects from the **PROJECT(S)** drop-down list.




6. Select the number of days from the **LAST N DAYS** drop-down list.

7. The default **Display Type** is **EventsDual**.
8. Select report visibility: Public or Private.
9. Click **Preview**.
10. Click **Create**.  
The report is created and the **View Report** page appears.

#### Print or download charts

You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.




11. Click **Back to Reports List** to go back to the Reports dashboard.

#### Agile reports

Here's a list of TeamForge agile reports such as the release burn up and burn down reports.

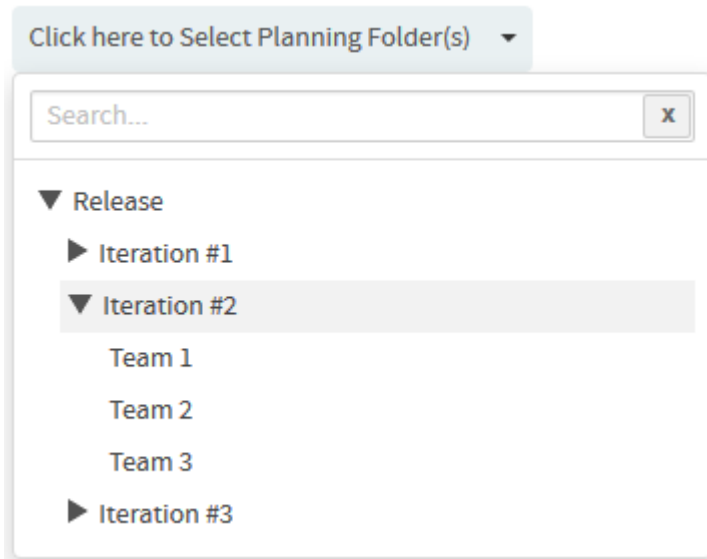
#### Cumulative Flow Chart

The cumulative flow chart shows the progress of backlog items by status for a sprint or a release.




- Generate this report to see the rolled up status chart of scheduled work items in a release or a sprint.
- By viewing the rolled up status of backlog items by date, you can forecast whether you are on track or not, adjust the scope if required and identify bottlenecks in your release or sprint.
- You have 'Date' in the X axis and 'Plan Estimate' in terms of number of points in the Y axis.

1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Cumulative Flow Chart** from **Agile Reports**.
4. Type a report title and description.
5. Select a planning folder from the **SELECT PLANNING FOLDER(S)** drop-down list.



6. Select one or more trackers from the **SELECT TRACKER(S)** drop-down list.




7. Leave the **CHART DISPLAY TYPE** as Area, which is the only available chart type for this report.
8. Select report visibility: Public or Private.
9. Click **Preview**.
10. Click **Create**.

The report is created and the **View Report** page appears.

#### **Print or download charts**


You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.




11. Click **Back to Reports List** to go back to the Reports dashboard.

### Release Burn Up Chart


The Release Burn Up Chart shows the work progress to date against the total planned work. This chart shows the total planned work, total work completed to date and the rate of progress (velocity).



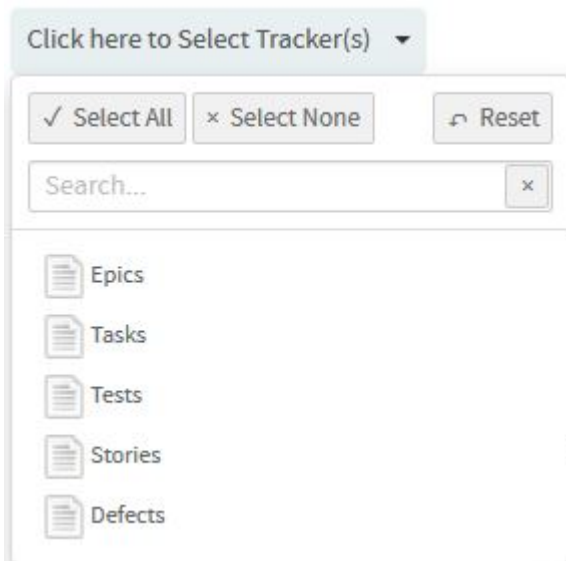
 **Important:** You must have the planning folder's start and end dates defined to generate this chart.

This chart includes a trend line to show the planned work scope change between a release's start and end dates (in other words, the start and end dates of the release's planning folder). Optionally, you can also include a trend line in the chart that forecasts when planned work might be completed depending on the average rate of progress (velocity).

1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Release Burn Up Chart** from **Agile Reports**.
4. Type a report title and description.
5. Select a planning folder from the **SELECT PLANNING FOLDER(S)** drop-down list.



6. Select one or more trackers from the **SELECT TRACKER(S)** drop-down list.




7. Optionally, select the **Include Forecast** check box.
8. Select either Points or Hours.
9. Leave the **CHART DISPLAY TYPE** as Trendlines, which is the only available chart type for this report.
10. Select report visibility: Public or Private.
11. Click **Preview**.
12. Click **Create**.

The report is created and the **View Report** page appears.

#### **Print or download charts**

You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.



13. Click **Back to Reports List** to go back to the Reports dashboard.


### Burn Down Chart

The Burn Down Chart shows your project's progress on a daily basis.

As work gets completed sprint-by-sprint for a release, your backlog tends to decrease. The burn down chart tells the story of how much work is left to be done versus how much time is left. With time along the X axis, the amount of work (backlog) measured in story points is on the Y axis. You can also see the average velocity per sprint in the release burn down chart. In addition to creating burn down charts by release, you can configure this report for a selected sprint planning folder to see its progress on a day-to-day basis.

- The story told in your burn down chart is only as reliable as the underlying data.
- You must have the planning folder's start and end dates defined to generate this chart.
- You can have the **POINTS** field enabled or disabled for a tracker. If a release planning folder consists of both "points-enabled" and "points-disabled" trackers, the release burn down chart shows grey-colored bars for all sprints, meaning, there are work items without points estimation data. This is because, the "points-disabled" tracker work items are considered as unestimated backlog items as they don't have points data. As a workaround, while configuring the release burn down chart, you can include "points-enabled" trackers alone to have the release burn down chart behave as expected.

### Things to consider if you are generating burn down chart for a release planning folder




- It is recommended to have all your sprint planning folders organized hierarchically within your release planning folder. Release burn down chart can be generated only if that planning folder has child (sprint) planning folders.

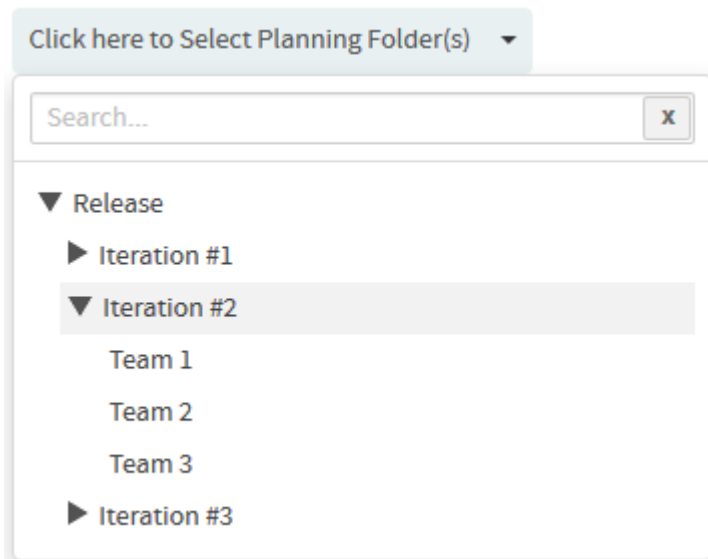


- Within a selected release planning folder, all immediate child planning folders are considered as sprint planning folders.
- The bar chart shows the work left to be done (estimated as points) on the first day of every sprint.
- The release burn down calculation is based on (story) points and no effort data is considered. Meaning, the release burn down will not behave as expected if you have effort data alone and no points data.


#### Things to consider if you are generating burn down chart for a sprint planning folder




- The sprint burn down calculation is based on remaining effort data and no (story) points data is considered. Meaning, the sprint burn down will not behave as expected if you have points data alone and no remaining effort data.
  - You can exclude weekends or specific days from your report if you are creating a burn down chart for a sprint planning folder.
  - The sprint burn down chart shows the number of tasks and hours remaining for a selected sprint planning folder.
1. Click **REPORTS** from the **Project Home** menu.
  2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
  3. Select **Burn Down Chart** from **Agile Reports**.
  4. Type a report title and description.
  5. Select a planning folder from the **SELECT PLANNING FOLDER(S)** drop-down list.



6. Select one or more trackers from the **SELECT TRACKER(S)** drop-down list.



7. Select one of the options, **Release** or **Sprint**, to create burn down chart for either a release planning folder or a sprint planning folder.


 **Tip:** Select the Exclude **Weekends** check box to exclude weekends from the report.

8. Leave the **CHART DISPLAY TYPE** as Burndown, which is the only available chart type for this report.
9. Select report visibility: Public or Private.
10. Click **Preview**.
11. Click **Create**.

The report is created and the **View Report** page appears.

#### **Print or download charts**


You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.



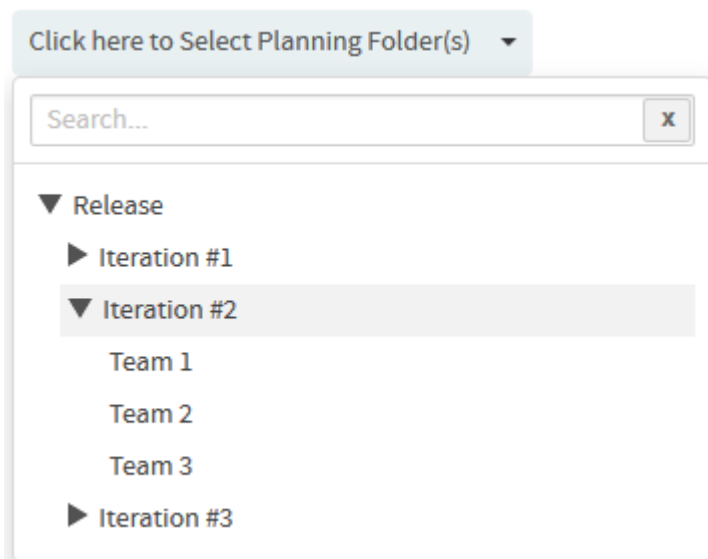
12. Click **Back to Reports List** to go back to the Reports dashboard.

### Committed vs Done vs Missed


The Committed vs Done vs Missed chart shows a comparison of the amount of work (in terms of story points) committed, completed and missed for a sprint.



1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Committed vs Done vs Missed** from **Agile Reports**.
4. Type a report title and description.
5. Select a planning folder from the **SELECT PLANNING FOLDER(S)** drop-down list.



6. Select one or more trackers from the **SELECT TRACKER(S)** drop-down list.




7. Leave the **CHART DISPLAY TYPE** as Columndone, which is the only available chart type for this report.
8. Select report visibility: Public or Private.
9. Click **Preview**.
10. Click **Create**.

The report is created and the **View Report** page appears.

#### **Print or download charts**

You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.



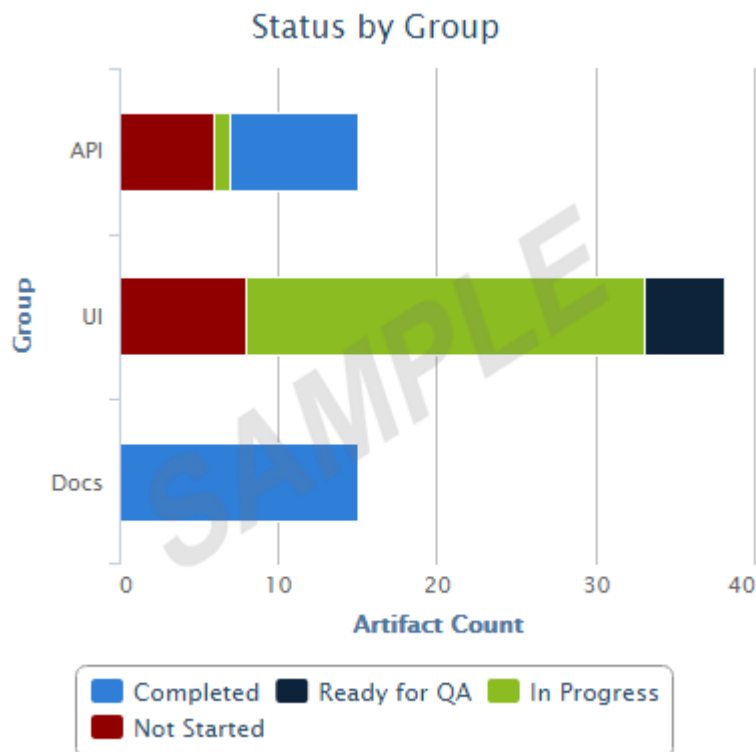
11. Click **Back to Reports List** to go back to the Reports dashboard.


### *Distribution reports*

Here's a list of reports to see distribution of work items such as tracker artifacts by parameters such as status, size, effort and so on.


#### Status Distribution by Area/Group

You can generate this chart for a specific sprint or release planning folder (optionally you can include child planning folders, if any) and take stock of the number of artifacts in different statuses and have them grouped by Group, Category, Customer, Assigned To, Priority or Teams.




 **Note:** Unlike other reports (that use datamart), this report runs on your operational database.

1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Status Distribution by Area/Group** from **Distribution Reports**.
4. Type a report title and description.
5. Select a planning folder from the **SELECT PLANNING FOLDER(S)** drop-down list.



6. Optionally, select the **Include: Child planning folders** check box to have the child planning folders, if any, included.
7. Select one or more trackers from the **SELECT TRACKER(S)** drop-down list.




8. Select an option from the **GROUP BY** drop-down list to have the chart grouped by Group, Category, Customer, Assigned To, Priority or Teams.
9. Leave the **CHART DISPLAY TYPE** as Bar, which is the only available chart type for this report.
10. Select report visibility: Public or Private.
11. Click **Preview**.
12. Click **Create**.

The report is created and the **View Report** page appears.

#### **Print or download charts**


You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.




13. Click **Back to Reports List** to go back to the Reports dashboard.


#### Total Size by Area/Group

You can generate this chart for a specific sprint or release planning folder (optionally you can include child planning folders, if any) and take stock of the total size (as in total of story points or estimated, actual or remaining effort) of artifacts and have them grouped by Group, Category, Customer, Assigned To, Priority, Tracker or Teams. You can generate this report for all artifacts or just for open or closed artifacts.




 **Note:** Unlike other reports (that use datamart), this report runs on your operational database.

1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Total Size by Area/Group** from **Distribution Reports**.
4. Type a report title and description.
5. Select a planning folder from the **SELECT PLANNING FOLDER(S)** drop-down list.



6. Select one or more trackers from the **SELECT TRACKER(S)** drop-down list.




7. Optionally, select the **Include: Child planning folders** check box to have the child planning folders, if any, included.
8. Select what you want to total from the **TOTAL BY** drop-down list. You can total by Story Points or Estimated Effort or Actual Effort or Remaining Effort.
9. Select an option from the **GROUP BY** drop-down list to have the chart grouped by Group, Category, Customer, Assigned To, Priority, Tracker or Teams.
10. Select one of the options, Include artifacts: **All** (default) or **Open** or **Closed**, to include all artifacts or just open or closed artifacts respectively.
11. Leave the **CHART DISPLAY TYPE** as Bar, which is the only available chart type for this report.
12. Select report visibility: Public or Private.
13. Click **Preview**.
14. Click **Create**.

The report is created and the **View Report** page appears.

#### **Print or download charts**




You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.




15. Click **Back to Reports List** to go back to the Reports dashboard.

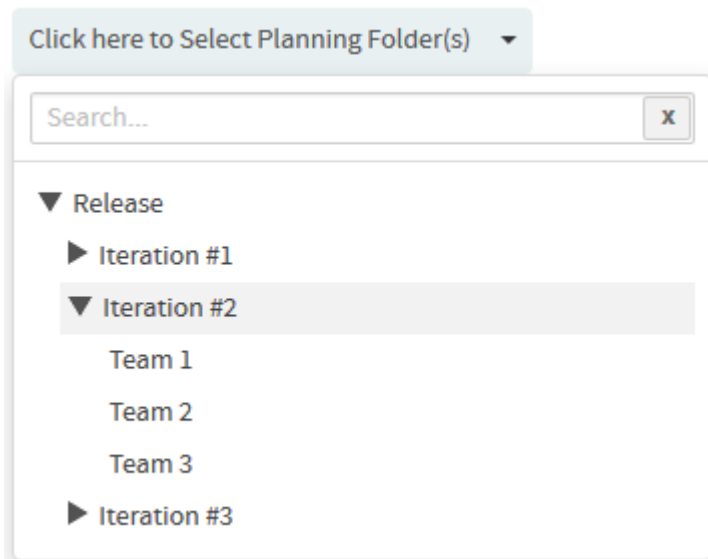
#### Average Size by Area/Group

You can generate this chart for a specific sprint or release planning folder (optionally you can include child planning folders, if any) and take stock of the average size (as in average of story points or estimated, actual or remaining effort) of artifacts and have them grouped by Group, Category, Customer, Assigned To, Priority, Tracker or Teams. You can generate this report for all artifacts or just for open or closed artifacts.




 **Note:** Unlike other reports (that use datamart), this report runs on your operational database.

1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Average Size by Area/Group** from **Distribution Reports**.
4. Type a report title and description.
5. Select a planning folder from the **SELECT PLANNING FOLDER(S)** drop-down list.



6. Optionally, select the **Include: Child planning folders** check box to have the child planning folders, if any, included.
7. Select one or more trackers from the **SELECT TRACKER(S)** drop-down list.




8. Select what you want to average from the **AVERAGE BY** drop-down list. You can average by Story Points or Estimated Effort or Actual Effort or Remaining Effort.
9. Select an option from the **GROUP BY** drop-down list to have the chart grouped by Group, Category, Customer, Assigned To, Priority, Tracker or Teams.
10. Select one of the options, Include artifacts: **All** (default) or **Open** or **Closed**, to include all artifacts or just open or closed artifacts respectively.
11. Leave the **CHART DISPLAY TYPE** as Bar, which is the only available chart type for this report.
12. Select report visibility: Public or Private.
13. Click **Preview**.
14. Click **Create**.

The report is created and the **View Report** page appears.

#### **Print or download charts**


You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.




15. Click **Back to Reports List** to go back to the Reports dashboard.


#### Total Size by Tracker Type

You can generate this chart for a specific planning folder (optionally you can include child planning folders, if any) and take stock of the total size of artifacts belonging to one or more tracker types and have them aggregated by Story Points or Actual Effort or Estimated Effort or Remaining Effort. This report is more useful if you want to know the total size by tracker types for a release and so is typically run against a specific release planning folder. You can generate this report for all artifacts or just for open or closed artifacts.



 **Note:** Unlike other reports (that use datamart), this report runs on your operational database.

1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Total Size by Tracker Type** from **Distribution Reports**.
4. Type a report title and description.
5. Select a planning folder from the **SELECT PLANNING FOLDER(S)** drop-down list.




6. Optionally, select the **Include: Child planning folders** check box to have the child planning folders, if any, included.
7. Select what you want to aggregate from the **AGGREGATE BY** drop-down list. You can aggregate by Story Points or Estimated Effort or Actual Effort or Remaining Effort.
8. Select one of the options, Include artifacts: **All** (default) or **Open** or **Closed**, to include all artifacts or just open or closed artifacts respectively.
9. Leave the **CHART DISPLAY TYPE** as Bar, which is the only available chart type for this report.
10. Select report visibility: Public or Private.
11. Click **Preview**.
12. Click **Create**.

The report is created and the **View Report** page appears.

#### Print or download charts

You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.




13. Click **Back to Reports List** to go back to the Reports dashboard.


#### Artifact Distribution Chart (Multiple Trackers)

Using this chart, you can report the status of artifacts belonging to more than one tracker or planning folder. You can also create charts based on a combination of more than one tracker and planning folder.


## Open 7.x Artifacts by Tracker Name



1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Artifact Distribution Chart (Multiple Trackers)** from **Distribution Reports**.
4. Type a report title and description.
5. Select one or more trackers from the **SELECT TRACKER(S)** drop-down list.



6. Select a planning folder from the **SELECT PLANNING FOLDER(S)** drop-down list.




7. Select one of the options from the **DISTRIBUTE BY** drop-down list.
8. Select one of the options from the **OPEN VS CLOSE** drop-down list.
9. Leave the **CHART DISPLAY TYPE** as Pie, which is the only available chart type for this report.
10. Select report visibility: Public or Private.
11. Click **Preview**.
12. Click **Create**.

The report is created and the **View Report** page appears.

#### Print or download charts


You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.



13. Click **Back to Reports List** to go back to the Reports dashboard.

#### Artifact Distribution Chart (Single Tracker)

This chart comes in handy if you want to report the status of artifacts belonging to a specific tracker.




1. Click **REPORTS** from the **Project Home** menu.

2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Artifact Distribution Chart (Single Tracker)** from **Distribution Reports**.
4. Type a report title and description.
5. Select a tracker from the **SELECT TRACKER(S)** drop-down list.
6. Select one or more planning folders (select the check boxes) from the **SELECT PLANNING FOLDER(S)** drop-down list.
7. Select one of the options from the **DISTRIBUTE BY** drop-down list.
8. Select one of the options from the **OPEN VS CLOSE** drop-down list.
9. Leave the **CHART DISPLAY TYPE** as Pie, which is the only available chart type for this report.
10. Select report visibility: Public or Private.
11. Click **Preview**.
12. Click **Create**.  
The report is created and the **View Report** page appears.

### Print or download charts

You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.



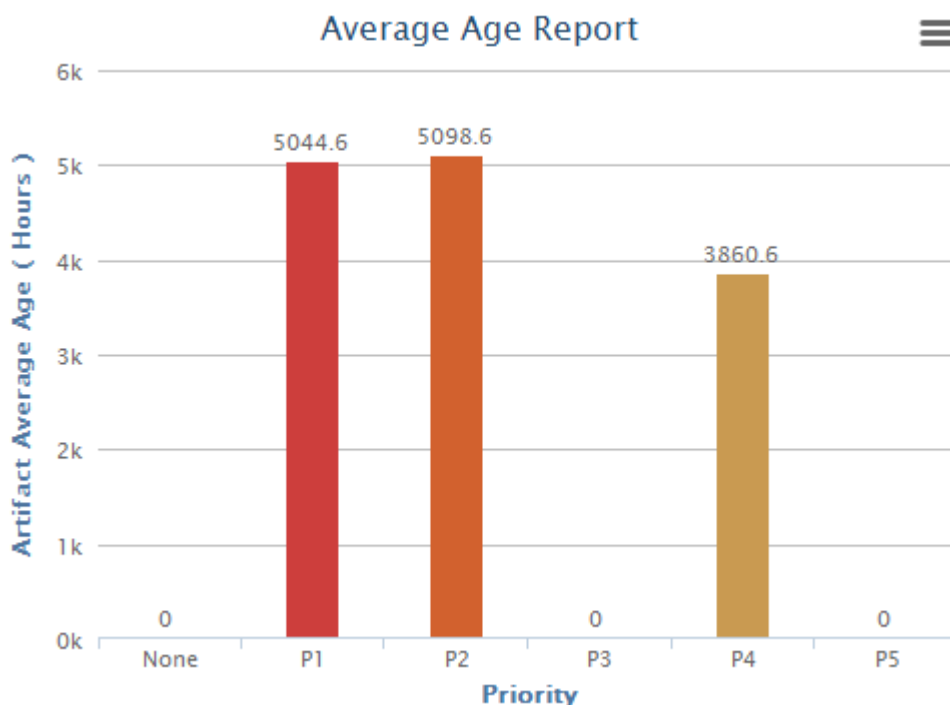
13. Click **Back to Reports List** to go back to the Reports dashboard.

### *Trend reports*

Here's a list of reports to know some artifact trend data.

#### Average Age Report

The average age chart lets you know the average age of artifacts in one or more trackers you select.



You can:

- Generate this report to know the average age of artifacts in various units of time such as number of hours, days or weeks.
- Group this report either by artifact priority such as P1, P2 and so on or by both artifact priority and category.
- Generate this report for either open or closed artifacts.
- Exclude weekends from being included in the average age calculation.

1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Average Age Report** from **Trend Reports**.
4. Type a report title and description.
5. Select one or more trackers from the **SELECT TRACKER(S)** drop-down list.

Click here to Select Tracker(s) ▾

✓ Select All
✕ Select None
↺ Reset

Search... ✕

Epics

Tasks


Tests

Stories

Defects



6. Select a planning folder from the **SELECT PLANNING FOLDER(S)** drop-down list.




7. Select one of the time units such as hours, days or weeks from the **AVG. TIME IN** drop-down list.
8. Select either **Priority** or **Priority & Category** from the **GROUP BY** drop-down list.
9. If required, select **EXCLUDE Weekends** check box to exclude weekends from being included in the average age calculation.
10. Select either Open Only or Closed from the **OPEN VS. CLOSE** drop-down list.
11. Leave the **CHART DISPLAY TYPE** as Stackedcolumns, which is the only available chart type for this report.
12. Select report visibility: Public or Private.
13. Click **Preview**.
14. Click **Create**.

The report is created and the **View Report** page appears.

#### Print or download charts


You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.



15. Click **Back to Reports List** to go back to the Reports dashboard.

#### Artifact Open/Close Chart (Multiple Trackers)

Use this chart to know the number of opened and closed artifacts over a period of time across one or more trackers or planning folders. This chart also helps in knowing the number of associated commits in specific repositories over a period of time. You can also know the total number of open artifacts through this chart.



The number of opened and closed artifacts, and the number of associated commits are shown as color-coded bars (scale on left Y coordinate axis) and the total number of open artifacts is shown as a color-coded line (scale on right Y coordinate axis). The X coordinate axis shows the time scale of the chart.


At least one of the reporting parameters such as the **Tracker ID**, **Planning Folder ID** or the **Repository ID** is required to generate this chart. The following table lists the various reporting parameters for Artifact - Open/Close Chart (Multiple Trackers).

1. Click **REPORTS** from the **Project Home** menu.
2. Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
3. Select **Artifact Open/Close Chart (Multiple Trackers)** from **Trend Reports**.
4. Type a report title and description.
5. Select one or more trackers from the **SELECT TRACKER(S)** drop-down list.

The dialog box titled "Click here to Select Tracker(s)" contains the following elements:

- Buttons:  Select All,  Select None, and a Reset button.
- Search bar: "Search..." with a clear (x) button.
- List of planning folders:
  - Epics
  - Tasks
  - Tests
  - Stories
  - Defects

6. Select one or more planning folders (select the check boxes) from the **SELECT PLANNING FOLDER(S)** drop-down list.
7. Select one or more repositories from the **SELECT REPOSITORY(S)** drop-down list.




8. Select one or more artifact priorities.  
For example, select P1 and P2 to include P1 and P2 artifacts in your chart. Select the **None** to include artifacts that have no priority assigned to them.
9. Select one of the report durations: Last Week/Last Month/Last Quarter/Last Year.
10. Leave the **CHART DISPLAY TYPE** as Dualaxeslinecolumn, which is the only available chart type for this report.
11. Select report visibility: Public or Private.
12. Click **Preview**.
13. Click **Create**.

The report is created and the **View Report** page appears.

#### Print or download charts

You can print charts or download them as .PNG, .JPG, .SVG or .PDF files using the print/download quick function icon.



14. Click **Back to Reports List** to go back to the Reports dashboard.

#### *Table reports: Task and Tracker reports*

These are reports on trackers and tasks in a tabular format.

- **Task reports:** Task reports display selected summary data about project tasks. You can generate reports on the tasks in a selected project or across multiple projects.
  - a) Click **REPORTS** from the **Project Home** menu.
  - b) Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
  - c) Select **Task** from **Table Reports**.
  - d) **Title** and **Description:** Type a title and description for your report.
  - e) **Project(s):** Select the project or projects from which you want task data to be reported.
  - f) **Status(s):** Select the status of the tasks to be included in your report.
  - g) **Priority(s):** Select one or more priorities of tasks to be included in your report.
  - h) Select one or more **Assigned To** or **Submitted By** project members.  
Click the '+' icon to add members.

ASSIGNED TO:

Select Assigned To +

SUBMITTED BY:

Select Submitted By +

- i) **Start Date** and **End Date**: Select the start and end date ranges of tasks to be included in your report.
- j) **Report Field(s)**: Select the fields you want displayed in your report.
- k) Select one of the report visibility options: Public (default) or Private.
- l) Click **Create**.

The **View Report** page appears. The task report is created.

- **Tracker reports**: Tracker reports give you a summary of the history of tracker artifacts. A tracker report can cover the tracker artifacts in a selected tracker or across all trackers in a project.

- a) Click **REPORTS** in the project navigation bar.
- b) Click **Create** in the **List Reports** page.  
The **Select Report Type** page appears.
- c) Select **Tracker** from **Table Reports**.

To create a tracker report across multiple projects, click **Show All Projects** and select the trackers from the required projects.

- d) **Title** and **Description**: Type a title and description for your report.
- e) **Select Tracker(s)**: Select the trackers from which you want tracker artifact data to be reported.
- f) **Group By**: Have the report grouped by one of the following: Assigned To, Category, Customer, Group, Priority, Status or Team.
- g) **Summary Statistics**: Select one of the following statistics to summarize the report: Count of Artifacts, Sum of Points or Sum of Effort.
- h) **Priority(s)**: Select one or more priorities of artifacts to be included in your report
- i) **Select Artifact Maturity(s)**: Select one or all of the following values to have all or new or modified artifacts included in the report respectively: Any or New or Edited
- j) Select date ranges for create date (**Submitted On**), last edited date (**Last Modified**), or closed date (**Closed**).
- k) Select one or more **Assigned To** or **Submitted By** project members.

Click the '+' icon to add members.

ASSIGNED TO:

SUBMITTED BY:

l) Select a planning folder from the **Select Planning Folder(s)** drop-down list.

m) Select one of the report visibility options: Public (default) or Private.

n) Click **Next**.

A list of filters (such as category, status, reported in release, fixed in release and so on) for the selected trackers are shown. Select the filter values from the drop-down lists and click **Create**. Based on the criteria being selected for filtering, it returns the configured rows.

The **View Report** page appears. The default value of the token `MAX_REPORT_ROWS` is set to '300' in `site-options.conf` file. For trackers with more than 300 artifacts, the tracker report is generated based on the criteria:

- If you have selected a single tracker having more than 300 artifacts, the tracker report shows only the first 300 artifacts. To see all the artifacts in that tracker, click the **Export** button.
- If you have selected multiple trackers with more than 300 artifacts in all, then the tracker report shows the artifacts based on the following calculation:

**No. of artifacts in report** =  $MAX\_REPORT\_ROWS \div \text{No. of selected trackers}$

For example, if you have selected 3 trackers with more than 300 artifacts in all, the first 100 artifacts is displayed in the tracker report for each tracker. You can export the report for each tracker to view all the artifacts in it.

#### **Important:**


In a tracker report, when you select more than one value from a multi-select, user-defined field as filter and if an artifact is associated with all of the selected values, then that artifact's record is duplicated for each of the selected values.

For example, assume that you have a 'Select User' multi-select, user-defined field with values 'User 1', 'User 2' and 'User 3' in a tracker report. All these three values are associated with 'artifact 1001'. Select all three values as filter and generate the tracker report. You will see 'artifact 1001' record being duplicated, that is, you will see three individual 'artifact 1001' records created for each of the three users.

### **Run, edit, export or delete a report**

Run a report to get fresh data on the status of your project.

The report data is regenerated each time the report is run. You can also refresh the data by clicking the 'Refresh' icon in the **List Reports** page.



- 👉 **Tip:** You can modify the report criteria before generating the report data. For example, you might want to run the same report each week, using all of the same report criteria except the start and end dates.
- 👉 **Note:** Report results contain only those items that you have permission to view. Other project members with different permissions might get different results when running the same report.

1. Click **REPORTS** from the **Project Home** menu.
2. Click the title of the desired report. The report is generated and displayed.
3. To edit a report, click **REPORTS**, select the report you want to edit (select the check box), click **Edit**, make the changes, and click **Save**.
4. To export report data, click **REPORTS**, select the table report you want to export (click the report title).

👉 **Note:** You can export only Table reports.

👉 **Important:** For multi-select user-defined select user field, you will export individual record for each user.

5. On the **Export Data** window, select an export format (.csv, .xml, .xlsx or tab-delimited file) and select the columns you want on the report and click **Export**.
6. To delete reports, click **REPORTS**, select one or more reports you want to delete (select the check boxes) and click **Delete**.

A confirmation message is displayed. Click **OK** to delete the selected reports.

### Customize reports

You can customize your reports by modifying the parameters in `config.ini` file. For example, some reports require a particular data source such as datamart and some require a particular license type such as ALM.

You may want to hide or display reports depending on the availability of their data source or the type of license so that users do not see reports which they cannot run.

For example:

- There are reports which require a data store (datamart or EventQ's events) to be enabled.
- There are many reports that do not apply to SCM users.

Further, you can also change a report's title, description or timeToLive duration of cached data and so on.

You can customize reports by modifying the parameters in `config.ini` file. For example, if you want to customize the "Average Size by Area/Group" report, modify the parameters in `/opt/collabnet/teamforge/var/cliserver/app/reports/pkg/averageSizeByArea/config.ini` file. The following illustration shows a sample `config.ini` file.

```

[main]
title = Average Size by Area/Group
description = Shows the average story points or effort per artifact by group, customer, category, assigned to, priority
outputType = bar
fields = planId,trackerIds,recursive,averageBy,groupBy,includeArtifacts
timeToLive = 2
almRequired = true
category = DistributionReports

[planId]
label = Planning Folder ID
type = wizard
object = plan
max = 1
required = true

[trackerIds]
label = Tracker ID
type = wizard
object = tracker
max = 10
required = false

[recursive]
label = Include
type = checkbox
values = Child planning folders

[averageBy]
label = Average by
type = select
values = Story Points,Estimated Effort,Remaining Effort,Actual Effort
default = Story Points

[groupBy]
label = Group By
type = select
values = Group,Category,Customer,Assigned To,Priority,Tracker,Team
default = Group

[includeArtifacts]
label = Include artifacts
type = radio
values = All,Open,Closed
default = All


[bar]
title = Average Artifact [:var:box.field.averageBy:] by [:var:box.field.groupBy:] ([:var:box.field.includeArtifacts:])
width = 500
height = 400
xAxis = Average [:var:box.field.averageBy:]
yAxis = [:var:box.field.groupBy:]

```

Here's a sample list of parameters you can modify:

- *almRequired*
  - If *almRequired*=true: The site should be in the ALM mode and the users should have an ALM license to operate the ALM reports.
  - If *almRequired*=false: The site can be in either of the modes; both ALM and SCM licensed users can operate these reports.
- *devopsrequired*
  - If *devopsrequired*=true: Users should have DevOps license to operate DevOps reports.
- *scmrequired*
  - If *scmrequired*=true: Users should have Version Control license to operate SCM reports.
- *eventsRequired*
  - If *eventsRequired*=true: EventQ's events data store should be installed. Users can operate these reports irrespective of the site mode and license type.


- If `eventsRequired=false`: Users cannot operate these reports irrespective of the site mode and license type.
  - `datamartRequired`
    - If `datamartRequired=true`: Datamart should be installed. Users can operate these reports irrespective of the site mode and license type.
    - If `datamartRequired=false`: Users cannot operate these reports irrespective of the site mode and license type.
  - `category`
    - You can change the category of a report by specifying it in `config.ini`. For example, to change the category of a specific report from 'Agile' to 'Distribution', specify as shown below:
 

```
category=DistributionReports
```
  - `timeToLive`
    - If `timeToLive=2`: The cached report data invalidated and refreshed every 2 hours.
    - The `timeToLive` parameter sets the duration (in hours) after which the cached report data is invalidated and refreshed from the data source. The default `timeToLive` for reports that use operational database is 2 hours. It is 24 hours for reports that use datamart. You can change the `timeToLive` duration depending on your site's requirements. However, shorter `timeToLive` duration can impact performance, though.
-  **Important:** Once you are done with your changes to `config.ini` file, run the `post-install.py` script to go live with the changes.

### Advanced reporting and datamart access

Using external reporting and OLAP tools, query the datamart directly and generate reports. The database schema diagrams provide the means to create advanced query scripts to extract required information from the datamart.

Accessing the datamart for reporting provides more analytical data than is provided by the TeamForge user interface options. The external tool must connect directly to the datamart and is then granted read-only permission to all data in the datamart.

 **Note:** It is recommended that you limit the access to the datamart to a limited set of trusted users.

#### Datamart schemas

The `Users`, `SCM`, and `Tracker` schemas are documented here for use in queries to the datamart. The datamart uses a "Star Schema" design for tables.


#### User schema

Query the user schema to obtain useful information in the user database tables detailed here in a schema diagram.

User schema contains user login information; TeamForge captures one fact row for each user that is logged in during the day.



## User schema diagram



## Description of user schema

### etl\_job

Used to track the ETL run status. There is one record per ETL run for a job, for example, Tracker ETL or User ETL. etl\_job has a 1-to-many relationship with audit\_dimension since a job may update more than one fact table. All report generation queries must "join" the etl\_job table with the condition etl\_job.status=1, thereby discarding data from incomplete ETL runs.

### audit\_dimension

Holds metadata about fact table records. There is one record per fact table for an ETL run.

### date\_dimension

Conformed dimension used for all transaction times.

### user\_dimension

Used for string user attributes and is a "slowly changing dimension of type 2 (SCD-2)." is\_super\_user, status, and license\_type are the SCD-2 fields.

### activity\_dimension

Conformed dimension that stores the activity or transaction names for various activities being tracked.

### user\_transaction\_fact

A fact-less fact table with user data of "daily" granularity.

## Sample queries

You can obtain useful user information by querying the user database, and further refine the results by using filters on the "date", "user type" (admin or non), "status", and "license type" fields. For example:

- Number of users who are logged in, by day, over a period of time:

```
SELECT c.date_of_trans as Date, count(distinct(b.id)) as NumUsers
FROM user_transaction_fact a, user_dimension b, date_dimension c, etl_job
d
WHERE a.user_key=b.user_key and a.trans_date_key=c.date_key and
a.job_key=d.job_key
and d.status=1 and c.date_of_trans >= '2012-12-17' and c.date_of_trans <=
'2012-12-21'
GROUP BY c.date_of_trans
```


- List of users who have logged in:

```
SELECT c.date_of_trans as Date, b.username as UserName
FROM user_transaction_fact a, user_dimension b, date_dimension c,
etl_job d
WHERE a.user_key=b.user_key and a.trans_date_key=c.date_key and
a.job_key=d.job_key
and d.status=1 and c.date_of_trans >= '2012-12-17' and
c.date_of_trans <= '2012-12-21'
GROUP BY c.date_of_trans, b.username
```

## SCM schema

Query the SCM schema to obtain useful commit information in the SCM tables detailed here in a schema diagram.

## SCM schema diagram



## Description of SCM schema

### etl\_job

Used to track the ETL run status. There is one record per ETL run for a job, for example, Tracker ETL or User ETL. **etl\_job** has a 1-to-many relationship with **audit\_dimension** since a job may update more than one fact table. All report generation queries must "join" the **etl\_job** table with the condition **etl\_job.status=1**, thereby discarding data from incomplete ETL runs.

### audit\_dimension

Holds metadata about fact table records. There is one record per fact table for an ETL run.

### date\_dimension

Conformed dimension used for all transaction times.

<b>activity_dimension</b>	Conformed dimension that stores the activity or transaction names for various activities being tracked.
<b>user_dimension</b>	Used for string user attributes and is a "slowly changing dimension of type 2 (SCD-2)." <code>is_super_user</code> , <code>status</code> , and <code>license_type</code> are the SCD-2 fields.
<b>project_dimension</b>	Used for storing project attributes and is a "slowly changing dimension of type 2 (SCD-2)." <code>is_deleted</code> and <code>project_access_level</code> are the SCD-2 fields.
<b>repository_dimension</b>	Used for storing repository attributes and is a "slowly changing dimension of type 1."
<b>scm_transaction_fact</b>	<p>The fact table for SCM activities with "transaction" granularity. TeamForge inserts a row in this table for every SCM activity that it processes in a transaction.</p> <ul style="list-style-type: none"> <li>• TeamForge object id , if available.</li> <li>• Number of files added, deleted, modified, moved, copied, if applicable.</li> </ul>

### Sample queries

You can obtain useful SCM information by querying the SCM database. For example:

- Number of SCM commits, sorted by date:

```
select b.date_of_trans as Date,
       count(a.scm_transaction_fact_key) as NumCommits
from scm_transaction_fact a, date_dimension b
where a.trans_date_key=b.date_key
group by b.date_of_trans
```

- Number of SCM commits, with quarterly trend:

```
select 'Q' || b.quarter as Quarter,
       count(a.scm_transaction_fact_key) as NumCommits
from scm_transaction_fact a, date_dimension b
where a.trans_date_key=b.date_key
group by b.quarter
```

- List of users who made commits.

```
select b.username as UserName,
       count(a.scm_transaction_fact_key) as NumCommits
from scm_transaction_fact a, user_dimension b
where a.user_key=b.user_key
group by b.username
```

- Project-wise commit data:

```
select b.id as ProjectId, b.title as ProjectName,
       count(a.scm_transaction_fact_key) as NumCommits
from scm_transaction_fact a, project_dimension b
where a.project_key=b.project_key
group by b.id, b.title
```

- Commits by date, in a specific project:

```
select c.date_of_trans as Date, b.id as ProjectId, b.title as
ProjectName,
       count(a.scm_transaction_fact_key) as NumCommits
```

```

from scm_transaction_fact a, project_dimension b,
date_dimension c
  where a.project_key=b.project_key and
a.trans_date_key=c.date_key
      and b.id='proj1008'
group by c.date_of_trans, b.id, b.title

```

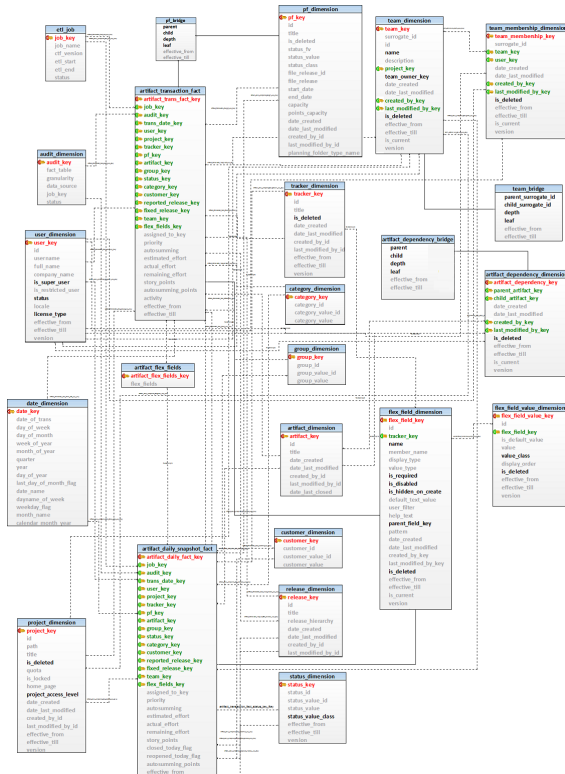
**Tracker schema**

Query the tracker schema to obtain useful information in the tracker database tables detailed here in a schema diagram.

TeamForge constructs the state or field values of the artifact during each update. The schema addresses both activity queries and total count queries.

TeamForge includes information for the activities in fixed fields, default artifact fields, and flex fields. This information includes artifact create, move, update for fixed-value changes, delete, open\_to\_close, and close\_to\_open.

**Tracker schema diagram**



**Description of tracker schema**

**etl\_job**

Used to track the ETL run status. There is one record per ETL run for a job, for example, Tracker ETL or User ETL. etl\_job has a 1-to-many relationship with audit\_dimension since a job may update more than one fact table. All report generation queries must "join" the etl\_job table with the condition etl\_job.status=1, thereby discarding data from incomplete ETL runs.

**audit\_dimension**

Holds metadata about fact table records. There is one record per fact table for an ETL run.

<b>date_dimension</b>	Conformed dimension used for all transaction times.
<b>user_dimension</b>	Used for string user attributes and is a "slowly changing dimension of type 2 (SCD-2)." <code>is_super_user</code> , <code>status</code> , and <code>license_type</code> are the SCD-2 fields.
<b>project_dimension</b>	Used for storing project attributes and is a "slowly changing dimension of type 2 (SCD-2)." <code>is_deleted</code> and <code>project_access_level</code> are the SCD-2 fields.
<b>pf_dimension</b>	The planning folder dimension for storing planning folder attributes. Use this table to generate reports without planning folder hierarchy.
<b>pf_bridge</b>	<p>A table for the planning folder hierarchy containing end-of-day status. <code>pf_bridge</code> is a "slowly changing dimension of type 2 (SCD-2)." You must limit queries to a given parent planning folder while joining with the <code>pf_bridge</code> table. Use <code>pf_bridge</code> to generate reports around planning folder hierarchy by joining with parent tables such as <code>artifact_daily_snapshot_fact</code> or <code>artifact_transaction_fact</code>.</p> <p>Joining <code>pf_bridge</code> with <code>artifact_transaction_fact</code> generates correct results only if end-of-day in queries is set to "12:00 a.m.". While formulating queries with <code>pf_bridge</code>, please note:</p> <ul style="list-style-type: none"> <li>• <code>parent</code> and <code>child</code> fields — Contain values from <code>pf_dimension.pf_key</code>; a depth of 0 indicates self.</li> <li>• <code>leaf</code> field — Is true if the child is a leaf node, otherwise false. Every planning folder will have an entry here with a depth of 0; every parent planning folder will have entries for all of its children, recursively, up to the leaf node of all branches.</li> <li>• <code>effective_from</code> and <code>effective_till</code> fields — Indicate the period when the parent-child relationship is correct, and can be used in queries to get the hierarchy for a specified time period.</li> </ul>
<b>tracker_dimension</b>	Used for storing tracker attributes and is a "slowly changing dimension of type 2 (SCD-2)." <code>is_deleted</code> is the changing field that act as a filter for reports .
<b>artifact_dimension</b>	Used for storing artifact data and is a "slowly changing dimension of type 1."
<b>group_dimension</b>	Used for holding values of the TeamForge tracker field "group". <code>group_dimension</code> has values for all artifacts from all Trackers.
<b>status_dimension</b>	Used for holding values of the TeamForge tracker field "status". The <code>status_value_class</code> field represents the meta-status of an artifact and has values "Open" and "Close". <code>status_dimension</code> has values for all artifacts from all Trackers.

<b>category_dimension</b>	Used for holding values of the TeamForge tracker field "category". <code>category_dimension</code> has values for all artifacts from all Trackers.
<b>customer_dimension</b>	Used for holding values of the TeamForge tracker field "customer". <code>customer_dimension</code> has values for all artifacts from all Trackers.
<b>release_dimension</b>	Used for holding values of the TeamForge tracker fields "Reported in Release" and "Fixed in Release". <code>release_dimension</code> has values for all artifacts from all Trackers.
<b>artifact_transaction_fact</b>	Every change in fixed or default artifact field values, or to project or tracker artifacts, results in a row inserted to <code>artifact_transaction_fact</code> . Changes to flex-field values, adding comments, attachments, and so on do not add a row to the table. The <code>artifact_dimension.date_last_modified</code> field has the time in the source tracker at the time of the ETL run. <code>artifact_transaction_fact</code> can be used to generate reports around activities such as create, update, delete and move, and for intra-day reports. <code>artifact_transaction_fact</code> has a "transaction" granularity.
<b>artifact_daily_snapshot_fact</b>	An aggregate table that holds the daily snapshot data or end-of-day status. <code>artifact_daily_snapshot_fact</code> can be used to generate reports for artifact close & re-open counts. It is recommended to use this table for end-of-day reports as it has fewer rows compared to <code>artifact_transaction_fact</code> . <code>artifact_daily_snapshot_fact</code> has a "daily" granularity.
<b>team_dimension</b>	A dimension for holding values of Teams. This is a "slowly changing dimension of type 2 (SCD-2)". Use this table to generate reports without team hierarchy.
<b>team_membership_dimension</b>	This is more or less a factless fact table that holds the changes in terms of Team memberships. This is a "slowly changing dimension of type 2 (SCD-2)".
<b>team_bridge</b>	A table for handling team hierarchy containing end-of-day status. <code>team_bridge</code> is a "slowly changing dimension of type 2 (SCD-2)." You must limit queries to a given parent team while joining with the <code>team_bridge</code> table. Use <code>team_bridge</code> to generate reports around team hierarchy (ex: a report on all artifacts assigned to Team 1 and its child teams) by joining with parent tables such as <code>artifact_daily_snapshot_fact</code> or <code>artifact_transaction_fact</code> .  Joining <code>team_bridge</code> with <code>artifact_transaction_fact</code> generates correct results only if end-of-day in queries is set to "12:00 a.m.". While formulating queries with <code>team_bridge</code> , please note:

- `parent_surrogate_id` and `child_surrogate_id` fields — Contain values from `team_dimension.surrogate_id`; a depth of 0 indicates self.
- `leaf` field — Is true if the child is a leaf node, otherwise false. Every team will have an entry here with a depth of 0; every parent team will have entries for all of its children, recursively, up to the leaf node of all branches .
- `effective_from` and `effective_till` fields — Indicate the period when the parent-child relationship is correct, and can be used in queries to get the hierarchy for a specified time period.

**Note:** `team_dimension` is a "slowly changing dimension of type 2". So, make sure that the JOIN clause in queries that join `team_bridge` and `team_dimension` includes the date and time as well.

**flex\_field\_dimension**

Used for storing information about flex fields and is a "slowly changing dimension of type 2 (SCD-2)".

**flex\_field\_value\_dimension**

Used for storing information about flex field values and is a "slowly changing dimension of type-2 (SCD-2)".

**artifact\_dependency\_dimension**

Used for storing relationship between artifacts and is a "slowly changing dimension of type-2 (SCD-2)".

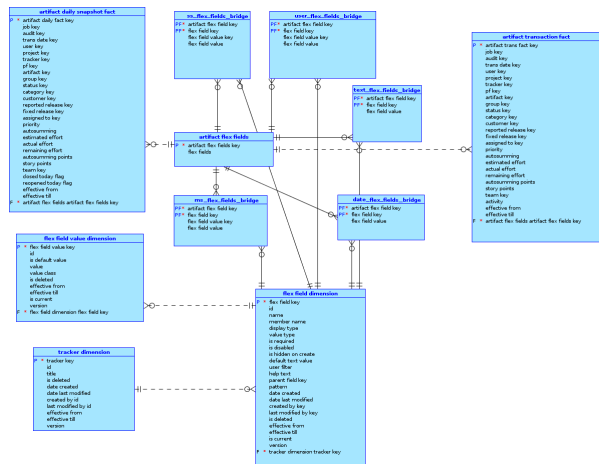
**artifact\_dependency\_bridge**

A table for the artifacts hierarchy containing end-of-day status. This is a "slowly changing dimension of type-2 (SCD-2). Use this table to generate reports around artifact hierarchy.

**artifact\_flex\_fields**

A table to bind the flex field updates on artifacts into a well formed XML, which stores the 'field key', 'value' and 'type' of the flex field.

**Schema diagram for XML to non-XML conversion**





## Description of schema used for XML to non-XML conversion

<b>flex field dimension</b>	Used for storing information about flex fields and is a "slowly changing dimension of type 2 (SCD-2)".
<b>flex field value dimension</b>	Used for storing information about flex field values and is a "slowly changing dimension of type-2 (SCD-2)".
<b>artifact flex fields</b>	A table to bind the flex field updates on artifacts into a well formed XML, which stores the 'field key', 'value' and 'type' of the flex field.
<b>artifact daily snapshot fact</b>	An aggregate table that holds the daily snapshot data or end-of-day status. artifact_daily_snapshot_fact can be used to generate reports for artifact close & re-open counts. It is recommended to use this table for end-of-day reports as it has fewer rows compared to artifact_transaction_fact. artifact_daily_snapshot_fact has a "daily" granularity.
<b>artifact transaction fact</b>	Every change in fixed or default artifact field values, or to project or tracker artifacts, results in a row inserted to artifact_transaction_fact. Changes to flex-field values, adding comments, attachments, and so on do not add a row to the table. The artifact_dimension.date_last_modified field has the time in the source tracker at the time of the ETL run. artifact_transaction_fact can be used to generate reports around activities such as create, update, delete and move, and for intra-day reports. artifact_transaction_fact has a "transaction" granularity.
<b>tracker dimension</b>	Used for storing tracker attributes and is a "slowly changing dimension of type 2 (SCD-2)." is_deleted is the changing field that act as a filter for reports .
<b>ss_flex_fields_bridge</b>	A table to hold details about single select flex field assigned to an artifact. To get the details of the assigned artifact, the "ss_flex_fields_bridge" table must be joined with the "artifact_transaction_fact" table and the "artifact_dimension" table.
<b>user_flex_fields_bridge</b>	A table to hold details about user flex field assigned to an artifact. To get the details of the assigned artifact, the "user_flex_fields_bridge" table must be joined with the "artifact_transaction_fact" table and the "artifact_dimension" table.
<b>text_flex_fields_bridge</b>	A table to hold details about text flex field assigned to an artifact. To get the details of the assigned artifact, the "text_flex_fields_bridge" table must be joined with the "artifact_transaction_fact" table and the "artifact_dimension" table.
<b>date_flex_fields_bridge</b>	A table to hold details about date flex field assigned to an artifact. To get the details of the assigned artifact, the "date_flex_fields_bridge" table must be joined

with the "artifact\_transaction\_fact" table and the "artifact\_dimension" table.

### ms\_flex\_fields\_bridge

A table to hold details about multi select flex field assigned to an artifact. To get the details of the assigned artifact, the "ms\_flex\_fields\_bridge" table must be joined with the "artifact\_transaction\_fact" table and the "artifact\_dimension" table.

### Sample queries

You can obtain useful tracker information by querying the tracker database. For example:

- Number of artifacts created in a tracker, sorted by date:

```
SELECT b.date_of_trans, count(a.artifact_key)
      FROM artifact_transaction_fact a, date_dimension b,
      tracker_dimension c
      WHERE a.trans_date_key=b.date_key and
      a.tracker_key=c.tracker_key
      and a.activity='Create' and c.title='Tracker-1'
      and b.date_of_trans >= '2011-10-31' and b.date_of_trans <=
      '2011-11-07'
      GROUP BY b.date_of_trans
      ORDER BY b.date_of_trans
```

- Number of artifacts created in a tracker, sorted by date and priority:

```
SELECT b.date_of_trans, d.title as Tracker, 'P'||a.priority as
      Priority,c.id ArtifactId
      FROM artifact_transaction_fact a, date_dimension b,
      artifact_dimension c, tracker_dimension d
      WHERE a.trans_date_key=b.date_key and
      a.artifact_key=c.artifact_key and a.tracker_key=d.tracker_key
      and a.activity='Create' and d.title='Tracker-1'
      and b.date_of_trans >= '2011-10-31' and b.date_of_trans <=
      '2011-11-07'
      ORDER BY b.date_of_trans, d.title, Priority
```

- Number of artifacts created on a particular day in a particular planning folder:

```
SELECT b.date_of_trans, c.id ArtifactId, c.title
      FROM artifact_transaction_fact a, date_dimension b,
      artifact_dimension c, pf_dimension d, pf_bridge e
      WHERE a.trans_date_key=b.date_key and
      a.artifact_key=c.artifact_key and a.pf_key=e.child
      and d.pf_key=e.parent and a.activity='Create' and
      d.title='Pf-1'
      and '2011-10-31' >= e.effective_from and '2011-10-31' <
      e.effective_till
      and b.date_of_trans = '2011-10-31'
      ORDER BY b.date_of_trans
```

- Number of closed tracker artifacts, sorted by day:

```
SELECT b.date_of_trans, count(a.artifact_key)
      FROM artifact_daily_snapshot_fact a, date_dimension b,
      tracker_dimension c
      WHERE a.trans_date_key=b.date_key and
      a.tracker_key=c.tracker_key and a.closed_today_flag='true'
      and c.title='Tracker-1' and b.date_of_trans >= '2011-10-31'
      and b.date_of_trans <= '2011-11-07'
      GROUP BY b.date_of_trans
```

```
ORDER BY b.date_of_trans
```

- List of artifacts in the "Open" state, sorted by day:

```
SELECT a.date_of_trans, c.id, c.title
       FROM date_dimension a inner join artifact_daily_snapshot_fact
       b
         on a.date_of_trans >= date(b.effective_from) and
       a.date_of_trans < date(b.effective_till)
         inner join artifact_dimension c on
       b.artifact_key=c.artifact_key
         inner join status_dimension d on b.status_key=d.status_key
       WHERE d.status_value_class='Open' and a.date_of_trans >=
       '2011-10-31'
         and a.date_of_trans <= '2011-11-02'
       ORDER BY a.date_of_trans
```

- List of artifacts assigned to a team (including child teams) on a given day:

```
select ad.id as artifact_id
       from artifact_transaction_fact fact
         inner join date_dimension dd on (dd.date_of_trans + 1) between
       fact.effective_from and fact.effective_till
         inner join team_dimension child_td on child_td.team_key =
       fact.team_key
         inner join team_bridge tb on tb.child_surrogate_id =
       child_td.surrogate_id and
         (dd.date_of_trans + 1) between tb.effective_from and
       tb.effective_till
         inner join team_dimension parent_td on tb.parent_surrogate_id =
       parent_td.surrogate_id and
         (dd.date_of_trans + 1) between parent_td.effective_from
       and parent_td.effective_till
         inner join artifact_dimension ad on ad.artifact_key =
       fact.artifact_key
         inner join etl_job ej on ej.job_key = fact.job_key and
       ej.status = 1
       where dd.date_of_trans = '2015-01-01' and parent_td.id =
       'team1002'
       order by artifact_id
```


### *Datamart access using external tools*

Use external tools to directly query the PostgreSQL or Oracle datamarts to access more TeamForge data than is available through options on the TeamForge user interface.

### **Accessing the datamart**

You can use OLAP or GUI tools to query your datamart directly to procure all the TeamForge data that is relevant to your analysis.

The external tool must be in the same network as the datamart to ensure fast access with a direct connection to the database. On the TeamForge host machine, you must enable `REPORTS_DATABASE_HOST` and `REPORTS_DATABASE_PORT` for remote access.

 **Note:** You only have read access; the datamart does not allow writes.

### **Enabling the datamart for access**

You must enable the datamart for direct queries and re-start the site. Once enabled, all data in the datamart can be queried using external tools.

Enable TeamForge (PostgreSQL Datamart) for queries from external tools

You must enable, then re-start TeamForge so that you can use an external tool to query the datamart directly.

You cannot use external reporting tools to connect directly to a datamart if you have a single-box installation using `localhost:SERVICES`.

1. Edit the `site-options.conf` file. Set `REPORTS_DB_ACCESS_HOSTS` to the IP address or IP address range (CIDR address) from which the tool will establish connection to the datamart. Specify multiple IPs as a comma-separated list.


For example: `REPORTS_DB_ACCESS_HOSTS = 10.0.0.2,10.2.1.0/24`

 **Note:**

- If the TeamForge site is not within your network and a Network Address Translation (NAT) is configured, then specify the NAT's external facing IP in the token.
- If you have used an "advanced" installation method to install TeamForge, you must also manually add the IP addresses to the `pg_hba.conf` file.

2. In `site-options.conf`, set `REPORTS_DATABASE_READ_ONLY_USER` and `REPORTS_DATABASE_READ_ONLY_PASSWORD` with your user and password.
3. On the TeamForge host machine, enable `REPORTS_DATABASE_HOST` and `REPORTS_DATABASE_PORT` for remote access.
4. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

You can now enable trusted users to establish connections by providing them with values of options

`REPORTS_DATABASE_HOST`, `REPORTS_DATABASE_PORT`, `REPORTS_DATABASE_READ_ONLY_USER`, and `REPORTS_DATABASE_READ_ONLY_PASSWORD`.

Enable TeamForge (Oracle Datamart) for queries from external tools

You must enable, then restart TeamForge so that you can use an external tool to query the datamart directly.

You can use external reporting tools to connect directly to a datamart only if you do not have a single-box installation using `localhost:SERVICES`.

1. In `site-options.conf`, set `REPORTS_DATABASE_READ_ONLY_USER` and `REPORTS_DATABASE_READ_ONLY_PASSWORD` with your user and password.
2. On the TeamForge host machine, enable `REPORTS_DATABASE_HOST` and `REPORTS_DATABASE_PORT` for remote access.
3. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

You can now enable trusted users to establish connections by providing them with values of options

`REPORTS_DATABASE_HOST`, `REPORTS_DATABASE_PORT`, `REPORTS_DATABASE_READ_ONLY_USER`, and `REPORTS_DATABASE_READ_ONLY_PASSWORD`.

## Common errors while connecting to PostgreSQL or Oracle datamarts

You may encounter these errors while configuring your datamart to allow queries from external tools.

### PostgreSQL database access error

Error: psql: FATAL: no pg\_hba.conf entry for host "111.111.111.111", user "test", database "datamart"

This error in a PostgreSQL datamart is because TeamForge is not granting access to your IP address.

Solution: Have the TeamForge administrator grant access to your IP address.

### PostgreSQL authentication error

Error: psql: FATAL: password authentication failed for user "test"

Solution: Specify the correct user name and password.

### PostgreSQL connection error

Error: psql: could not connect to server: Connection refused. Is the server running on host "<datamart-host>" and accepting TCP/IP connections on port 5632?

Solution:

- Check the host and port numbers; the database must be running on host "datamart-host" and accept TCP/IP connections on port 5632.
- Check if remote access is enabled on TeamForge; the Teamforge administrator can enable the required access.

### Oracle connection error

Error: ORA-12505, TNS:listener does not currently know of SID given in connect descriptor.

This error in an Oracle datamart occurs when some connection parameters are set incorrectly.

Solution: Check and set your connection parameters appropriately.

### Query PostgreSQL datamart based on flex fields

Create query scripts to query and extract the required information from the PostgreSQL datamart. Make sure these scripts are available to any user (including users with read-only permission) who wants to execute these scripts. This topic lists the functions and sample queries for a few specific use cases.

#### Important:

- Log into TeamForge as a Reporting user and run the queries. Note that this is a one-time process.
- Use the following flex field functions by joining the `flex_field_dimension`, `artifact_flex_field` and `artifact_transaction_fact` (or) `artifact_daily_snapshot_fact` tables as it depends on the *XML data* and *flex field key* generated by the platform.
- Common script used across all flex field functions:

```
CREATE OR REPLACE FUNCTION array_search(needle anyelement, haystack
  anyarray)
  RETURNS integer AS
  $BODY$
    SELECT i
      FROM generate_subscripts($2, 1) AS i
     WHERE $2[i] = $1
     ORDER BY i
  $BODY$
```

```
LANGUAGE sql STABLE
COST 100;
```

- **Function for Date flex fields**

**Input arguments**

- Flex field XML (derived automatically by joining the artifact\_transaction\_fact (or) artifact\_daily\_snapshot\_fact and artifact\_flex\_field tables)
- Flex field key (derived automatically from flex\_field\_dimension table and other tables)

**Output arguments**

Date flex field values

```
CREATE OR REPLACE FUNCTION get_artifact_date_value(artifact_xml xml,
  field_key integer)
  RETURNS character varying AS
$BODY$
DECLARE
field_value_key varchar(9055) default '';
BEGIN
  SELECT cast((xpath('/fields/field/@val', artifact_xml))
  [(array_search(field_key::text, (xpath('/fields/field/
@key', artifact_xml))::text[]))]
  as varchar) into field_value_key;
  Return field_value_key;
END;
$BODY$
LANGUAGE plpgsql VOLATILE
COST 100;
```

- **Function for Text flex fields**

**Input arguments**

- Flex field XML (derived automatically by joining the artifact\_transaction\_fact (or) artifact\_daily\_snapshot\_fact and artifact\_flex\_field tables)
- Flex field key (derived automatically from flex\_field\_dimension table and other tables)

**Output arguments**

Text flex field values

```
CREATE OR REPLACE FUNCTION replacen(artifact_xml1 xml)
  RETURNS text AS
$BODY$
select replace ((select
replace(artifact_xml1::varchar(3000), '<!', '&lt;!'), ']>', ']&gt;');
$BODY$
LANGUAGE sql STABLE
COST 100;

CREATE OR REPLACE FUNCTION get_artifact_text_value(artifact_xml xml,
  field_key integer)
  RETURNS character varying AS
$BODY$
DECLARE
  field_value_key varchar(9055) default '';

```

```

BEGIN
SELECT cast((xpath('/fields/field/@val',replacen(artifact_xml)::xml))
[(array_search(field_key::text,(xpath('/fields/field/
@key',replacen(artifact_xml)::xml))::text[]))]
as varchar) into field_value_key;
Return field_value_key;
END;
$BODY$
LANGUAGE plpgsql VOLATILE
COST 100;

```

- **Function for Single-select flex fields**

**Input arguments**

- Flex field XML (derived automatically by joining the artifact\_transaction\_fact (or) artifact\_daily\_snapshot\_fact and artifact\_flex\_field tables)
- Flex field key (derived automatically from flex\_field\_dimension table and other tables)

**Output arguments**

Values selected or stored in single-select flex field

```

CREATE OR REPLACE FUNCTION get_artifact_select_value(artifact_xml xml,
field_key integer)
RETURNS character varying AS
$BODY$
DECLARE
field_value_key varchar(9055) default '';
singleSelectValues varchar(9055) default '';
BEGIN
SELECT cast((xpath('/fields/field/@val',artifact_xml))
[(array_search(field_key::text,(xpath('/fields/field/
@key',artifact_xml))::text[]))]
as varchar) into field_value_key;
Select value into singleSelectValues from flex_field_value_dimension
where flex_field_value_key::text = (field_value_key::text);

Return singleSelectValues;
END;
$BODY$
LANGUAGE plpgsql VOLATILE
COST 100;

```

- **Function for Multi-select flex fields**

**Input arguments**

- Flex field XML (derived automatically by joining the artifact\_transaction\_fact (or) artifact\_daily\_snapshot\_fact and artifact\_flex\_field tables).
- Flex field key (derived automatically from flex\_field\_dimension table and other tables).
- Specific value(s) stored in the field or the keyword 'ALL' for retrieving all values stored in multi-select flex field.
- Logical conditional operator. Possible values are 'ALL' or 'ANY'. The argument value of 'ALL' performs a search equivalent to the 'AND' condition

and the value of 'ANY' performs a search equivalent to the 'OR' condition.

### Output arguments

Values selected or stored in multi-select flex field

```

-----First Execute Inner Function-----
-----Inner function start-----
CREATE OR REPLACE FUNCTION
  get_artifact_multiselect_values_any(artifact_xml xml, field_key integer,
    selectedfields text[])
RETURNS text AS
$BODY$
DECLARE
fieldValues text;
multiFieldValues text default '';
splitfield text[];
splitfield2 text[];
loop1 integer DEFAULT 1;
  loop2 integer DEFAULT 1;
arrayLength integer DEFAULT 1;
  arrayLength2 integer DEFAULT 1;
multifield text DEFAULT '';
filteredField text Default '';
BEGIN
  Select cast((xpath('/fields/field/@val',artifact_xml))
[array_search(field_key::text,(xpath('/fields/field/
@key',artifact_xml)::text[]))]
  as varchar(9055)) into fieldValues;
Select array_length(regexp_split_to_array(fieldValues, ','),1) into
  arrayLength;
Select regexp_split_to_array(fieldValues, ',') into splitfield;
IF arrayLength IS NOT NULL THEN
  FOR loop1 IN 1..arrayLength LOOP
    Select value into multiFieldValues from flex_field_value_dimension
  where flex_field_value_key = splitfield[loop1]::integer;
    multifield := concat(multifield,',',multiFieldValues);
    EXIT WHEN loop1 > arrayLength;
  END LOOP;

  multifield := rtrim(multifield,'" ');
  multifield := ltrim(multifield,',');

Select
  array_length(regexp_split_to_array(array_to_string(selectedFields,','),
  ','),1) into arrayLength2;
Select regexp_split_to_array(array_to_string(selectedFields,','), ',')
  into splitfield2;
IF selectedFields = Array['ALL'] THEN
filteredField:=multifield;
Return filteredField;
ELSE
FOR loop2 IN 1..arrayLength2 LOOP
IF  multifield ~ splitfield2[loop2] THEN
  filteredField := concat(splitfield2[loop2],',',filteredField);
END IF;
EXIT WHEN loop2 > arrayLength2;
END LOOP;
END IF;
END IF;
filteredField := rtrim(filteredField,'" ');
filteredField := rtrim(filteredField,',');
Return filteredField;

```





```

        filteredField := concat(v_array3[i],',',filteredField);

        count=count+1;
        EXIT;
        END IF;

    END LOOP;
    END LOOP;
ELSE
    results = false;

    END IF;
    IF (count=arrayLength3) THEN
        results=filteredField;
    ELSE
        results='';
    END IF;
    END IF;
ELSE
    select
get_artifact_multiselect_values_any(artifact_xml,field_key,selectedfields)
into returnvalues;
        IF returnvalues IS NOT NULL THEN
            results=returnvalues;
        ELSE
            results='';
        END IF;
    END IF;
    END IF;
    results := rtrim(results,'" ');
    results := rtrim(results,',');
    RETURN(results);
END;
$BODY$
LANGUAGE plpgsql VOLATILE
COST 100;

```

- **Function for Multi-user flex fields**

#### **Input arguments**

- Flex field XML (derived automatically by joining the artifact\_transaction\_fact (or) artifact\_daily\_snapshot\_fact and artifact\_flex\_field tables).
- Flex field key (derived automatically from flex\_field\_dimension table and other tables).
- Specific value(s) stored in the field or the keyword 'ALL' for retrieving all values stored in multi-select flex field.
- Logical conditional operator. Possible values are 'ALL' or 'ANY'. The argument value of 'ALL' performs a search equivalent to the 'AND' condition and the value of 'ANY' performs a search equivalent to the 'OR' condition.

#### **Output arguments**

Values selected or stored in multi-user flex field

```
-----First Execute Inner Function-----
```

```

-----Inner function start---
CREATE OR REPLACE FUNCTION get_artifact_user_values_any(artifact_xml xml,
  field_key integer, selectedfields text[])
  RETURNS text AS
$BODY$
DECLARE
fieldValues text;
multiFieldValues text default '';
splitfield text[];
splitfield2 text[];
loop1 integer DEFAULT 1;
  loop2 integer DEFAULT 1;
arrayLength integer DEFAULT 1;
  arrayLength2 integer DEFAULT 1;
multifield text DEFAULT '';
filteredField text Default '';
BEGIN
  Select cast((xpath('/fields/field/@val',artifact_xml))
[array_search(field_key::text,(xpath('/fields/field/
@key',artifact_xml)::text[]))]
  as varchar(9055)) into fieldValues;
Select array_length(regexp_split_to_array(fieldValues, ','),1) into
  arrayLength;
Select regexp_split_to_array(fieldValues, ',') into splitfield;
IF arrayLength IS NOT NULL THEN
  FOR loop1 IN 1..arrayLength LOOP
    Select full_name into multiFieldValues from user_dimension where
    user_key = splitfield[loop1]::integer;
    multifield := concat(multifield,',',multiFieldValues);
    EXIT WHEN loop1 > arrayLength;
  END LOOP;
  multifield := rtrim(multifield,' ');
  multifield := ltrim(multifield,',');

  Select
array_length(regexp_split_to_array(array_to_string(selectedFields,','),
','),1) into arrayLength2;
  Select regexp_split_to_array(array_to_string(selectedFields,','),
',') into splitfield2;
  IF selectedFields = Array['ALL'] THEN
    filteredField:=multifield;
    Return filteredField;
  ELSE
    FOR loop2 IN 1..arrayLength2 LOOP
      IF multifield ~ splitfield2[loop2] THEN
        filteredField :=concat(splitfield2[loop2],',',filteredField);
      END IF;
      EXIT WHEN loop2 > arrayLength2;
    END LOOP;
  END IF;
END IF;
filteredField := rtrim(filteredField,' ');
filteredField := rtrim(filteredField,',');
Return filteredField;
END;
$BODY$
LANGUAGE plpgsql VOLATILE
COST 100;

-----Inner Function End-----

create or replace FUNCTION get_artifact_user_value(artifact_xml xml,
  field_key integer,selectedfields text[],condition character varying)
  RETURNS text as

```

```

$BODY$
DECLARE
    fieldValues text;
    loop1 integer default 1;

    loop2 integer default 1;
    arrayLength1 integer default 1;
    arrayLength2 integer default 1;
    arrayLength3 integer default 1;
    multiFieldValues text default '';
    multifield text DEFAULT '';
    returnvalues text DEFAULT '';
    i integer default 1;
    j integer default 1;
    v_array1 text[];
    v_array2 text[];
    v_array3 text[];
    results text Default '';
    filteredField text Default '';
    count integer default 0;
BEGIN
    Select cast((xpath('/fields/field/@val',artifact_xml))
[array_search(field_key::text, (xpath('/fields/field/
@key',artifact_xml)::text[]))]
    as varchar(9055)) into fieldValues;
    Select array_length(regexp_split_to_array(fieldValues, ','),1) into
arrayLength1;
    Select regexp_split_to_array(fieldValues, ',') into v_array1;
    IF arrayLength1 IS NOT NULL THEN
        FOR loop1 in 1..arrayLength1 LOOP
            select full_name into multiFieldValues from user_dimension where
user_key=v_array1[loop1]::integer;
            multifield := concat(multifield,',',multiFieldValues);
            EXIT WHEN loop1 > arrayLength1;
            multifield := rtrim(multifield,'" ');
            multifield := ltrim(multifield,',');
            END LOOP;
            Select array_length(regexp_split_to_array(multifield, ','),1) into
arrayLength2;
            Select regexp_split_to_array(multifield, ',') into v_array2;
            Select
array_length(regexp_split_to_array(array_to_string(selectedfields,','),
','),1) into arrayLength3;
            Select regexp_split_to_array(array_to_string(selectedfields,','),
',') into v_array3;

            IF (UPPER(condition)=UPPER('ALL')) THEN
                IF selectedFields = Array['ALL'] THEN
                    results:=multifield;
                ELSE
            IF (v_array2 is not null AND v_array3 is not null) THEN
                FOR i in 1..arrayLength3 LOOP
                FOR j in 1..arrayLength2 LOOP
                    IF (v_array3[i]= v_array2[j]) THEN
                        filteredField := concat(v_array3[i],',',filteredField);

                        count=count+1;
                        EXIT;
                    END IF;

                END LOOP;
                END LOOP;
            ELSE
                results = false;

```

```

        END IF;
        IF (count=arrayLength3) THEN
            results=filteredField;
        ELSE
            results='';
        END IF;
        END IF;
    ELSE
        select
get_artifact_user_values_any(artifact_xml,field_key,selectedfields) into
returnvalues;
        IF returnvalues IS NOT NULL THEN
            results=returnvalues;
        ELSE
            results='';
        END IF;
    END IF;
    END IF;
    results := rtrim(results,'" ');
    results := rtrim(results,',');
    RETURN(results);
END;
$BODY$
LANGUAGE plpgsql VOLATILE
COST 100;

```

## Sample use cases and queries

### Note:

- Flex field Name, Value, Tracker title and so on that are used in the filtering conditions are case-sensitive.
- See [Case 7](#) below for building a case-insensitive search using upper() or lower() functions.

### Use case 1: Filter Date and Text flex fields

Suppose you want to retrieve data based on the following assumptions:

- Date flex field named 'CreatedDate' that has a value of '2015-09-02 00:00:00'
- Text flex field named 'Text102' with part of its value containing the pattern 'soft'
- Tracker title is 'Tracker101'

### Note:

- You can change the field name, tracker title and values based on the data in your system
- The date passed as input should be of the format “YYYY-MM-DD HH24:MI:SS”

### Sample query

```

select
t.Date as Date,
t.Project as Project,
t.Tracker as Tracker,
t.Planing_folder as Planing_folder,
'P'||t.Priority as Priority,
count(distinct t.artifact_key)
from

```

```

(
    select
    b.date_of_trans as Date,
    e.title as Project,c.title as
Tracker,
    d.title as Planing_folder,
    'P' || a.priority as Priority,

get_artifact_text_value(ff.flex_fields,ffd.flex
as TextFlexField,
    a.artifact_key
from
    artifact_transaction_fact a,
    date_dimension b,
    artifact_flex_fields ff,
    tracker_dimension c,
    pf_dimension d,
    project_dimension e ,
    flex_field_dimension ffd
where
    a.trans_date_key=b.date_key
    and ff.artifact_flex_fields_key
= a.flex_fields_key
    and a.tracker_key=c.tracker_key
    and a.pf_key=d.pf_key
    and a.project_key=e.project_key
    and
a.tracker_key=ffd.tracker_key
    and e.title='Project SCD-2 Test'
    and ffd.name in ('Text102')
    and c.title in ('Tracker101')
    and
get_artifact_text_value(ff.flex_fields,ffd.flex
like 'soft%'

union

    select    b.date_of_trans as Date,
    e.title as Project,c.title as
Tracker,
    d.title as Planing_folder,
    'P' || a.priority as Priority,

regexp_split_to_table(get_artifact_date_value(1
ffd.flex_field_key),E',') as
DateFlexField,
    a.artifact_key

from
    artifact_transaction_fact a,
    date_dimension b,
    artifact_flex_fields ff,
    tracker_dimension c,
    pf_dimension d,
    project_dimension e ,
    flex_field_dimension ffd
where
    a.trans_date_key=b.date_key
    and ff.artifact_flex_fields_key
= a.flex_fields_key
    and a.tracker_key=c.tracker_key
    and a.pf_key=d.pf_key
    and a.project_key=e.project_key
    and
a.tracker_key=ffd.tracker_key

```

```

and e.title='Project SCD-2 Test'
and ffd.name in ('CreatedDate')
and c.title in ('Tracker101')
and
get_artifact_date_value(ff.flex_fields,
ffd.flex_field_key)='2015-09-02
00:00:00') as t
group by 1,2,3,4,5
order by 1,2,3,4,5;

```

## Use case 2: Filter Single-select flex fields

Suppose you want to retrieve data based on the following assumptions:

- Flex field name: 'FSS1'
- Value selected or stored: 'S1'
- Tracker title: 'Tracker2'



### Note:

- You can change the field name, tracker title and values based on the data in your system

## Sample query

```

select
  b.date_of_trans as Date,
  pd.title as Project,
  td.title as Tracker,
  d.title as Planning_folder,
  'P'||a.priority as Priority ,
  get_artifact_select_value(aff.flex_fields,ffd.
as SingleSelectFlexField,
  count(distinct a.artifact_key)
TotalArtifacts
from
  artifact_transaction_fact a,
  date_dimension b,
  artifact_flex_fields aff,
  project_dimension pd ,
  tracker_dimension td,
  pf_dimension d,
  flex_field_dimension ffd
where
  a.trans_date_key=b.date_key
  and
  a.flex_fields_key=aff.artifact_flex_fields_key
  and a.tracker_key=td.tracker_key
  and a.project_key=pd.project_key
  and td.tracker_key=ffd.tracker_key
  and a.pf_key=d.pf_key
  and
  pd.title='ProjectTestFunctionSCD-2'
  and td.title in ('Tracker2')
  and ffd.name='FSS1'
  and
  get_artifact_select_value(aff.flex_fields,ffd.
  --and
  date(a.effective_from)<=date(now())
  and
  date(a.effective_till)>'2012-04-01'
group by 1,2,3,4,5,6
order by 1,2,3,4,5,6;

```

**Use case 3: Filter Text flex fields**

Suppose you want to retrieve data based on the following assumptions:

- Flex field name: 'Text101'
- Value contains: 'hello'
- Tracker name: 'Tracker101'

 **Note:**

- You can change the field name, tracker title and values based on the data in your system

**Sample query**


```
select
    b.date_of_trans as Date,
    e.title as Project,
    c.title as Tracker,
    d.title as Planing_folder,
    'P' || a.priority as Priority,

    get_artifact_text_value(ff.flex_fields,ffd.flex
as TextFlexField,
    count(distinct a.artifact_key)
Artifacts
from
    artifact_transaction_fact a,
    date_dimension b,
    artifact_flex_fields ff,
    tracker_dimension c,
    pf_dimension d,
    project_dimension e ,
    flex_field_dimension ffd
where
    a.trans_date_key=b.date_key
    and ff.artifact_flex_fields_key =
a.flex_fields_key
    and a.tracker_key=c.tracker_key
    and a.pf_key=d.pf_key
    and a.project_key=e.project_key
    and a.tracker_key=ffd.tracker_key
    and e.title='Project SCD-2 Test'
    and ffd.name = 'Text101'
    and c.title in ('Tracker101')
    and
    get_artifact_text_value(ff.flex_fields,ffd.flex
like 'hello%'
    --and
    date(a.effective_from)<=date(now())
    and
    date(a.effective_till)>'2012-04-01'
group by 1,2,3,4,5,6
order by 1,2,3,4,5,6;
```

**Use case 4: Filter Date flex fields**

Suppose you want to retrieve data based on the following assumptions:

- Flex field name: 'CreatedDate'
- Value: '2015-09-02 00:00:00'
- Tracker name: 'Tracker101'

 **Note:**



- You can change the field name, tracker title and values based on the data in your system

### Sample query

```
select
  b.date_of_trans as Date,
  e.title as Project,
  c.title as Tracker,
  d.title as Planing_folder,
  'P' || a.priority as Priority,

  regexp_split_to_table(get_artifact_date_value(
ffd.flex_field_key),E',') as
DateFlexField,
  count(distinct a.artifact_key)

from
  artifact_transaction_fact a,
  date_dimension b,
  artifact_flex_fields ff,
  tracker_dimension c,
  pf_dimension d,
  project_dimension e ,
  flex_field_dimension ffd
where
  a.trans_date_key=b.date_key
  and
  ff.artifact_flex_fields_key =
  a.flex_fields_key
  and
  a.tracker_key=c.tracker_key
  and a.pf_key=d.pf_key
  and
  a.project_key=e.project_key
  and
  a.tracker_key=ffd.tracker_key
  and e.title='Project SCD-2
Test'
  and ffd.name in
('CreateDate')
  and c.title in ('Tracker101')
  and
  get_artifact_date_value(ff.flex_fields,
ffd.flex_field_key)='2015-09-02
00:00:00'

group by 1,2,3,4,5,6
order by 1,2,3,4,5,6;
```

### Use case 5: Multi-select flex fields

Suppose you want to retrieve data based on the following assumptions:

- Multi-select flex field name: 'Country'
- Value: 'Russia,India'
- Tracker name: 'TrackerN'
- Conditional parameter: 'ALL'



#### Note:

- You can change the field name, tracker title and values based on the data in your system

- Flex field name, value and tracker title that are used in the SQL filter conditions are case sensitive
- If you want to select all the values in User flex field, then pass 'ALL' as the conditional parameter
- If you want to select any value, then pass 'ANY' as the conditional parameter

### Sample query

```

select b.date_of_trans as Date,
       e.title as Project,c.title as
       Tracker,
       d.title as Planing_folder,'P' ||
       a.priority as Priority,
       get_artifact_multiselect_value(ff.flex_fields,1)
       as MultiselectFlexField,
       count(distinct a.artifact_key)
       Artifacts
-- ,a.artifact_key
from
       artifact_transaction_fact a,
       date_dimension b,
       artifact_flex_fields ff,
       tracker_dimension c,
       pf_dimension d,
       project_dimension e ,
       flex_field_dimension ffd
where
       a.trans_date_key=b.date_key
       and ff.artifact_flex_fields_key
       = a.flex_fields_key
       and a.tracker_key=c.tracker_key
       and a.pf_key=d.pf_key
       and a.project_key=e.project_key
       and
       a.tracker_key=ffd.tracker_key
--and
       e.title='TestMultiFunction'
       and ffd.name = 'Country'
       and c.title in ('TrackerN')
       and
       get_artifact_multiselect_value(ff.flex_fields,1)
       =' '
--and
       date(a.effective_from)<=date(now())
       and
       date(a.effective_till)>'2012-04-01'
group by 1,2,3,4,5,6
order by 1,2,3,4,5,6;

```

### Use case 6: Filter User flex fields

Suppose you want to retrieve data based on the following assumptions:

- User flex field name: 'Select User'
- Value: 'user1,user2'
- Tracker name: 'TrackerN'
- Conditional parameter: 'ALL'

**Note:**

- You can change the field name, tracker title and values based on the data in your system
- Flex field name, value and tracker title that are used in the SQL filter conditions are case sensitive
- If you want to select all the values in User flex field, then pass 'ALL' as the conditional parameter
- If you want to select any value, then pass 'ANY' as the conditional parameter

**Sample query**

```

select b.date_of_trans as Date,
       e.title as Project,c.title as
       Tracker,
       d.title as Planing_folder,'P' ||
       a.priority as Priority,
       get_artifact_user_value(ff.flex_fields,ffd.flex
       as UserFlexField,
       count(distinct a.artifact_key)
       Artifacts
       --,a.artifact_key
from
       artifact_transaction_fact a,
       date_dimension b,
       artifact_flex_fields ff,
       tracker_dimension c,
       pf_dimension d,
       project_dimension e ,
       flex_field_dimension ffd
where
       a.trans_date_key=b.date_key
       and ff.artifact_flex_fields_key
       = a.flex_fields_key
       and a.tracker_key=c.tracker_key
       and a.pf_key=d.pf_key
       and a.project_key=e.project_key
       and
       a.tracker_key=ffd.tracker_key
       --and
       e.title='TestMultiFunction'
       and ffd.name = 'Select User'
       and c.title in ('TrackerN')
       and
       get_artifact_user_value(ff.flex_fields,ffd.flex
       ='
       --and
       date(a.effective_from)<=date(now())
       and
       date(a.effective_till)>'2012-04-01'
group by 1,2,3,4,5,6
order by 1,2,3,4,5,6;

```

**Use case 7: Filter Text flex fields, case-insensitive search**

Suppose you want to retrieve data based on the following assumptions:

- Flex field name: 'Text102'
- Value contains: 'soft'

- Tracker name: 'Tracker101'



**Note:**

- You can change the field name, tracker title and values based on the data in your system
- The `upper()` function can be replaced with the `lower()` function appropriately for case-insensitive search

**Sample query**

```
select
    b.date_of_trans as Date,
    e.title as Project,
    c.title as Tracker,
    d.title as Planing_folder,
    'P' || a.priority as Priority,

    get_artifact_text_value(ff.flex_fields,ffd.flex_fields)
as TextFlexField,
    count(distinct a.artifact_key)
as TotalCounts
from
    artifact_transaction_fact a,
    date_dimension b,
    artifact_flex_fields ff,
    tracker_dimension c,
    pf_dimension d,
    project_dimension e ,
    flex_field_dimension ffd
where
    a.trans_date_key=b.date_key
    and ff.artifact_flex_fields_key
= a.flex_fields_key
    and a.tracker_key=c.tracker_key
    and a.pf_key=d.pf_key
    and a.project_key=e.project_key
    and
a.tracker_key=ffd.tracker_key
    and
UPPER(e.title)=UPPER('Project SCD-2
Test')
    and
UPPER(ffd.name)=UPPER('Text102')
    and UPPER(c.title) =
UPPER('Tracker101')
    and
UPPER(get_artifact_text_value(ff.flex_fields,ffd.flex_fields))
like UPPER('soft%')

group by 1,2,3,4,5,6
order by 1,2,3,4,5,6;
```

*Query Oracle datamart based on flex fields*

Create query scripts to query and extract the required information from the Oracle datamart. Make sure these scripts are available to any user (including users with read-only permission) who wants to execute these scripts. This topic lists the functions and sample queries for a few specific use cases.



**Important:**

- Log into TeamForge as a Reporting user and run these queries. Note that this is a one-time process.

- Use the following flex field functions by joining the `flex_field_dimension`, `artifact_flex_field` and `artifact_transaction_fact` (or) `artifact_daily_snapshot_fact` tables as it depends on the *XML data* and *flex field key* generated by the platform.

- **Function for Date flex fields**

**Input arguments**

- Flex field XML (derived automatically by joining the `artifact_transaction_fact` (or) `artifact_daily_snapshot_fact` and `artifact_flex_field` tables)
- Flex field key (derived automatically from `flex_field_dimension` table and other tables)

**Output arguments**

Date flex field values

```
create or replace FUNCTION get_artifact_date_value( artifact_xml IN
XMLTYPE, field_key IN INTEGER )
RETURN varchar2
IS
flex varchar2(4000);
BEGIN
SELECT x.val INTO flex FROM dual ,
XMLTABLE ('/fields/field[@key=$keyalias]'
PASSING artifact_xml,field_key as "keyalias"
COLUMNS val VARCHAR2(4000) PATH '@val') x;
RETURN(flex);
END;
```

- **Function for Text flex fields**

**Input arguments**

- Flex field XML (derived automatically by joining the `artifact_transaction_fact` (or) `artifact_daily_snapshot_fact` and `artifact_flex_field` tables)
- Flex field key (derived automatically from `flex_field_dimension` table and other tables)

**Output arguments**

Text flex field values

```
create or replace FUNCTION replacen( artifact_xml1 IN XMLTYPE)
RETURN varchar2
IS
replaceval varchar2(3000);
BEGIN
SELECT REPLACE((REPLACE(cast(artifact_xml1 as
varchar2(3000)),' ','')),','') into replaceval from dual;
RETURN(replaceval);
END;
create or replace FUNCTION get_artifact_text_value( artifact_xml IN
XMLTYPE, field_key IN integer )
RETURN varchar2
IS flex varchar2(3200);
fields_xml XMLTYPE;
BEGIN
SELECT XMLTYPE(replacen(artifact_xml)) INTO fields_xml FROM dual;
SELECT x.val INTO flex FROM dual ,
XMLTABLE ('/fields/field[@key=$keyalias]'
```

```

        PASSING fields_xml,field_key as "keyalias"
        COLUMNS val VARCHAR2(400) PATH '@val' ) x;
RETURN (flex);
END;

```

- **Function for Single-select flex fields**

**Input arguments**

- Flex field XML (derived automatically by joining the artifact\_transaction\_fact (or) artifact\_daily\_snapshot\_fact and artifact\_flex\_field tables)
- Flex field key (derived automatically from flex\_field\_dimension table and other tables)

**Output arguments**

Values selected or stored in single-select flex field

```

create or replace FUNCTION get_artifact_select_value( artifact_xml IN
XMLTYPE, field_key IN INTEGER)
RETURN VARCHAR2
IS
    singleSelectValue VARCHAR2(100);
    flex INTEGER;
BEGIN
    SELECT x.val INTO flex FROM dual ,
    XMLTABLE ('/fields/field[@key=$keyalias]'
        PASSING artifact_xml,field_key as "keyalias"
        COLUMNS val VARCHAR2(30) PATH '@val') x;
    SELECT z.value INTO singleSelectValue from FLEX_FIELD_VALUE_DIMENSION
z where z.FLEX_FIELD_VALUE_KEY=flex;
RETURN (singleSelectValue);
END;

```

- **Function for Multi-select flex fields**

**Input arguments**

- Flex field XML (derived automatically by joining the artifact\_transaction\_fact (or) artifact\_daily\_snapshot\_fact and artifact\_flex\_field tables).
- Flex field key (derived automatically from flex\_field\_dimension table and other tables).
- Specific value(s) stored in the field or the keyword 'ALL' for retrieving all values stored in multi-select flex field.
- Logical conditional operator. Possible values are 'ALL' or 'ANY'. The argument value of 'ALL' performs a search equivalent to the 'AND' condition and the value of 'ANY' performs a search equivalent to the 'OR' condition.

**Output arguments**

Values selected or stored in multi-select flex field

```

-----Inner Function-----
create or replace FUNCTION get_artifact_multiselect_any(artifact_xml IN
XMLTYPE, field_key IN INTEGER,selectedfields IN varchar2)
RETURN varchar2
IS
    flexs varchar2(300);
    arraylength integer default 1;

```

```

loop1 integer default 1;
loop2 integer default 1;
i integer default 1;
j integer default 1;
v_array1 apex_application_global.vc_arr2;
v_array2 apex_application_global.vc_arr2;
v_array3 apex_application_global.vc_arr2;
results varchar2(200) ;

BEGIN
    SELECT x.val INTO flexs FROM dual ,
        XMLTABLE ('/fields/field[@key=$keyalias]'
            PASSING artifact_xml,field_key as "keyalias"
            COLUMNS val VARCHAR2(4000) PATH '@val') x;
    v_array1 := apex_util.string_to_table(flexs,',');

    FOR loop1 in 1..v_array1.count LOOP
        select value into v_array2(loop1) from flex_field_value_dimension
        where flex_field_value_key=v_array1(loop1);
    END LOOP;

    v_array3 := apex_util.string_to_table(selectedfields,',');

    FOR loop2 in 1..1 LOOP
        IF (v_array3(loop2)='ALL') THEN
            FOR i in 1..v_array2.count LOOP
                IF LENGTH(results)!=0 THEN
                    results:=results
                    ||','||v_array2(i);
                ELSE
                    results:=v_array2(i);
                END IF;
            END LOOP;
        ELSE
            FOR i in 1..v_array2.count LOOP
                FOR j in 1..v_array3.count LOOP
                    IF
                        (v_array3(j)=v_array2(i)) THEN
                            IF LENGTH(results)!=0
                                results:=results
                                ||','||v_array3(j);
                            ELSE
                                results:=v_array3(j);
                            END IF;
                        END IF;
                    END LOOP;
                END LOOP;
            END IF;
        END LOOP;

    RETURN(results);
    END;
/

```

```
-----End Inner function-----
```

```

    create or replace FUNCTION get_artifact_multiselect_value(artifact_xml
IN XMLTYPE, field_key IN INTEGER,selectedfields IN varchar2,condition IN
varchar2)
    RETURN varchar2
    IS
    flexs varchar2(300);
    loop1 integer default 1;
    loop2 integer default 1;
    i integer default 1;
    j integer default 1;
    v_array1 apex_application_global.vc_arr2;
    v_array2 apex_application_global.vc_arr2;
    v_array3 apex_application_global.vc_arr2;
    results varchar2(200) ;
    filteredField varchar2(200) ;
    valcount integer default 0;

    BEGIN
        SELECT x.val INTO flexs FROM dual ,
            XMLTABLE ('/fields/field[@key=$keyalias]'
                PASSING artifact_xml,field_key as "keyalias"
                COLUMNS val VARCHAR2(4000) PATH '@val') x;
        v_array1 := apex_util.string_to_table(flexs,',');

        FOR loop1 in 1..v_array1.count LOOP
            select value into v_array2(loop1) from flex_field_value_dimension
            where flex_field_value_key=v_array1(loop1);
        END LOOP;

        v_array3 := apex_util.string_to_table(selectedfields,',');

        IF (UPPER(condition)=UPPER('ALL')) THEN
            FOR loop2 in 1..1 LOOP
                IF (v_array3(loop2)='ALL') THEN
                    FOR i in 1..v_array2.count LOOP
                        IF LENGTH(results)!=0 THEN
                            results:=results ||','||v_array2(i);
                        ELSE
                            results:=v_array2(i);
                        END IF;
                    END LOOP;
                    filteredField:=results;
                    EXIT;
                ELSE

                    FOR i in 1..v_array3.count LOOP
                        FOR j in 1..v_array2.count LOOP
                            IF (v_array3(i)=v_array2(j)) THEN
                                IF LENGTH(results)!=0 THEN
                                    results:=results ||','||v_array3(j);
                                    valcount:=valcount+1;
                                ELSE
                                    results:=v_array3(j);
                                    valcount:=valcount+1;
                                END IF;
                            END IF;
                        END IF;
                    END IF;
                END IF;
            END IF;
        END IF;
    END IF;

```



```

                END LOOP;
            END LOOP;
            IF (valcount=v_array3.count) THEN
                filteredField:=results;
            ELSE
                filteredField:='';
            END IF;
        END IF;
    END LOOP;

ELSE
    select
    get_artifact_multiselect_any(artifact_xml,field_key,selectedfields) into
    results from dual;
    IF results IS NOT NULL THEN
        filteredField:=results;
    ELSE
        filteredField:='';
    END IF;
END IF;
filteredField:=TRIM(TRAILING ',' FROM filteredField);

RETURN(filteredField);
END;
```

- **Function for Multi-user flex fields**

#### Input arguments

- Flex field XML (derived automatically by joining the `artifact_transaction_fact` (or) `artifact_daily_snapshot_fact` and `artifact_flex_field` tables).
- Flex field key (derived automatically from `flex_field_dimension` table and other tables).
- Specific value(s) stored in the field or the keyword 'ALL' for retrieving all values stored in multi-select flex field.
- Logical conditional operator. Possible values are 'ALL' or 'ANY'. The argument value of 'ALL' performs a search equivalent to the 'AND' condition and the value of 'ANY' performs a search equivalent to the 'OR' condition.

#### Output arguments

Values selected or stored in multi-user flex field

```

-----Inner Function-----
create or replace FUNCTION get_artifact_user_any(artifact_xml IN
XMLTYPE, field_key IN integer,selectedfields IN varchar2)
RETURN varchar2
IS
flexs varchar2(300);
arraylength integer default 1;
loop1 integer default 1;
loop2 integer default 1;
i integer default 1;
j integer default 1;
v_array1 apex_application_global.vc_arr2;
v_array2 apex_application_global.vc_arr2;
```

```

v_array3 apex_application_global.vc_arr2;
results varchar2(200) ;

BEGIN
  SELECT x.val INTO flexs FROM dual ,
    XMLTABLE ('/fields/field[@key=$keyalias]'
      PASSING artifact_xml,field_key as "keyalias"
      COLUMNS val VARCHAR2(4000) PATH '@val') x;
  v_array1 := apex_util.string_to_table(flexs,',');

  FOR loop1 in 1..v_array1.count LOOP
    select full_name into v_array2(loop1) from user_dimension where
user_key=v_array1(loop1);
  END LOOP;

v_array3 := apex_util.string_to_table(selectedfields,',');

FOR loop2 in 1..v_array3.count LOOP
  IF (v_array3(loop2)='ALL') THEN
    FOR i in 1..v_array2.count LOOP
      IF LENGTH(results)!=0 THEN
        results:=results
||','||v_array2(i);
      ELSE
        results:=v_array2(i);
      END IF;
    END LOOP;
    EXIT;
  ELSE
    FOR i in 1..v_array2.count LOOP
      FOR j in 1..v_array3.count LOOP
        IF
(v_array3(j)=v_array2(i)) THEN
          IF LENGTH(results)!=0
THEN
            results:=results
||','||v_array3(j);
          ELSE
            results:=v_array3(j);
          END IF;
        END IF;
      END LOOP;
    END LOOP;
  END IF;
END LOOP;

RETURN(results);
END;

-----End Inner Function-----

create or replace FUNCTION get_artifact_user_value(artifact_xml IN
XMLTYPE, field_key IN INTEGER,selectedfields IN varchar2,condition IN
varchar2)
RETURN varchar2

```

```

IS
flexs varchar2(300);
loop1 integer default 1;
loop2 integer default 1;
i integer default 1;
j integer default 1;
v_array1 apex_application_global.vc_arr2;
v_array2 apex_application_global.vc_arr2;
v_array3 apex_application_global.vc_arr2;
results varchar2(200) ;
filteredField varchar2(200) ;
valcount integer default 0;

BEGIN
    SELECT x.val INTO flexs FROM dual ,
        XMLTABLE ('/fields/field[@key=$keyalias]'
            PASSING artifact_xml,field_key as "keyalias"
            COLUMNS val VARCHAR2(4000) PATH '@val') x;
    v_array1 := apex_util.string_to_table(flexs,',');

    FOR loop1 in 1..v_array1.count LOOP
        select full_name into v_array2(loop1) from user_dimension where
user_key=v_array1(loop1);
    END LOOP;

v_array3 := apex_util.string_to_table(selectedfields,',');

IF (UPPER(condition)=UPPER('ALL')) THEN
    FOR loop2 in 1..1 LOOP
        IF (v_array3(loop2)='ALL') THEN
            FOR i in 1..v_array2.count LOOP
                IF LENGTH(results)!=0 THEN
                    results:=results ||','||v_array2(i);
                ELSE
                    results:=v_array2(i);
                END IF;
            END LOOP;
            filteredField:=results;
            EXIT;
        ELSE

            FOR i in 1..v_array3.count LOOP
                FOR j in 1..v_array2.count LOOP
                    IF (v_array3(i)=v_array2(j)) THEN
                        IF LENGTH(results)!=0 THEN
                            results:=results ||','||v_array3(j);
                            valcount:=valcount+1;
                        ELSE
                            results:=v_array3(j);
                            valcount:=valcount+1;
                        END IF;
                    END IF;
                END LOOP;
            END LOOP;
            IF (valcount=v_array3.count) THEN
                filteredField:=results;
            ELSE
                filteredField:='';
            END IF;
        END IF;
    END IF;

```

```

        END IF;
    END LOOP;

ELSE
    select get_artifact_user_any(artifact_xml,field_key,selectedfields)
into results from dual;
    IF results IS NOT NULL THEN
        filteredField:=results;
    ELSE
        filteredField:='';
    END IF;
END IF;
filteredField:=TRIM(TRAILING ',' FROM filteredField);

RETURN(filteredField);
END;
```

## Sample use cases and queries

### Note:

- Flex field Name, Value, Tracker title and so on that are used in the filtering conditions are case-sensitive.
- See [Case 7](#) below for building a case-insensitive search using upper() or lower() functions.

### Use case 1: Filter Date and Text flex fields

Suppose you want to retrieve data based on the following assumptions:

- Date flex field named 'CreateDate' that has a value of '2015-07-07 00:00:00'
- Text flex field named 'TEXT' with part of its value containing the pattern 'cre'
- Tracker title is 'TestDateTracker'

### Note:

- You can change the field name, tracker title and values based on the data in your system
- The date passed as input should be of the format “YYYY-MM-DD HH24:MI:SS”

### Sample query

```

SELECT t.date_of_trans    "Date",
       t.Project          ,
       t.Tracker          ,
       t.PlanningFolder   ,
       t.Priority,
       count(distinct
t.artifact_key) "TotalArtifacts"
FROM (SELECT
b.date_of_trans,
e.title AS
Project,
c.title AS Tracker,
d.title AS
PlanningFolder,
'a.priority AS Priority,
a.artifact_key
```

```

FROM artifact_transaction_fact
a inner join date_dimension b on
(a.trans_date_key=b.date_key)
      inner join
artifact_flex_fields ff on
(ff.artifact_flex_fields_key =
a.flex_fields_key)
      inner
join tracker_dimension c on
(a.tracker_key=c.tracker_key)
      inner join pf_dimension
d on (a.pf_key=d.pf_key)
      inner join
project_dimension e on
(a.project_key=e.project_key)
      inner join
flex_field_dimension ffd on
(a.tracker_key=ffd.tracker_key)
and
ffd.name='Date'
and
c.title='TestDateTracker'
and
get_artifact_date_value(ff.flex_fields,ffd.flex
00:00:00'

UNION

SELECT          b.date_of_trans,
                e.title
AS Project,
                c.title AS Tracker,
                d.title
AS PlaningFolder,
                'P' ||
a.priority,
a.artifact_key
FROM artifact_transaction_fact
a inner join date_dimension b on
(a.trans_date_key=b.date_key)
      inner join
artifact_flex_fields ff on
(ff.artifact_flex_fields_key =
a.flex_fields_key)
      inner
join tracker_dimension c on
(a.tracker_key=c.tracker_key)
      inner join pf_dimension
d on (a.pf_key=d.pf_key)
      inner join
project_dimension e on
(a.project_key=e.project_key)
      inner join
flex_field_dimension ffd on
(a.tracker_key=ffd.tracker_key)
and
ffd.name='TEXT'
and c.title='TrackerTest'
and
get_artifact_text_value(ff.flex_fields,ffd.flex
LIKE 'cre%'
and d.id='plan1004' ) t

```

```
GROUP BY
  t.date_of_trans,t.Project,t.Tracker,t.PlaningFo
ORDER BY
  t.date_of_trans,t.Project,t.Tracker,t.PlaningFo
```

## Use case 2: Filter Single-select flex fields

Suppose you want to retrieve data based on the following assumptions:

- Flex field name: 'FSS1'
- Value selected or stored: 'S1'
- Tracker title: 'FTracker'



### Note:

- You can change the field name, tracker title and values based on the data in your system

## Sample query


```
SELECT      b.date_of_trans  "Date",
            e.title
            "Project",
            c.title
            "Tracker",
            d.title
            "Planingfolder",
            'P' || a.priority
            "Priority",
            count( distinct
a.artifact key) "TotalArtifacts"
FROM  artifact_transaction_fact
a inner join date_dimension b on
(a.trans_date_key=b.date_key)
            inner
join  artifact_flex_fields ff on
(ff.artifact_flex_fields_key =
a.flex_fields_key)
            inner
join  tracker_dimension c on
(a.tracker_key=c.tracker_key)
            inner join pf_dimension d on
(a.pf_key=d.pf_key)
            inner join project_dimension
e on  (a.project_key=e.project_key)

            inner join
flex_field_dimension ffd on
(a.tracker_key=ffd.tracker_key)
            and
ffd.name='FSS1'
            and
c.title='FTracker' and
get_artifact_select_value(ff.flex_fields,ffd.f
--where a.effective_from <=
sysdate and a.effective_till >
to_date('2012-04-01', 'YYYY-MM-DD')
GROUP BY
  b.date_of_trans,e.title,c.title,d.title,a.prior
ORDER BY
  b.date_of_trans,e.title,c.title,d.title,a.prior
```

**Use case 3: Filter Text flex fields**

Suppose you want to retrieve data based on the following assumptions:

- Flex field name: 'TEXT'
- Value contains: 'cre'
- Tracker name: 'TrackerTest'

 **Note:**

- You can change the field name, tracker title and values based on the data in your system

**Sample query**

```
SELECT          b.date_of_trans
               "Date",
               e.title
               "Project",
               c.title "Tracker",
               d.title
               "Planingfolder",
               'P' ||
a.priority    "Priority",

count( distinct a.artifact_key)
"TotalArtifacts"
FROM artifact_transaction_fact
a inner join date_dimension b on
(a.trans_date_key=b.date_key)
           inner join
artifact_flex_fields ff on
(ff.artifact_flex_fields_key =
a.flex_fields_key)
           inner
join tracker_dimension c on
(a.tracker_key=c.tracker_key)
           inner join pf_dimension
d on (a.pf_key=d.pf_key)
           inner join
project_dimension e on
(a.project_key=e.project_key)
           inner join
flex_field_dimension ffd on
(a.tracker_key=ffd.tracker_key)
and
ffd.name='TEXT'
and c.title='TrackerTest'
and
get_artifact_text_value(ff.flex_fields,ffd.flex
LIKE 'cre%'
and d.id='plan1004'
--where a.effective_from <=
sysdate and a.effective_till >
to_date('2012-04-01', 'YYYY-MM-DD')
GROUP BY
b.date_of_trans,e.title,c.title,d.title,a.prior
ORDER BY
b.date_of_trans,e.title,c.title,d.title,a.prior
```

**Use case 4: Filter Date flex fields**

Suppose you want to retrieve data based on the following assumptions:

- Flex field name: 'Date'

- Value: '2015-08-12 00:00:00'
- Tracker name: 'TestDateTracker'


**Note:**

- You can change the field name, tracker title and values based on the data in your system

**Sample query**

```

SELECT          b.date_of_trans
               "Date",
               e.title
               "Project",
               c.title "Tracker",
               d.title
               "Planingfolder",
               'P' ||
a.priority    "Priority",

count( distinct a.artifact_key)
               "TotalArtifacts"
FROM artifact_transaction_fact
a inner join date_dimension b on
(a.trans_date_key=b.date_key)
           inner join
artifact_flex_fields ff on
(ff.artifact_flex_fields_key =
a.flex_fields_key)
           inner
join tracker_dimension c on
(a.tracker_key=c.tracker_key)
           inner join pf_dimension
d on (a.pf_key=d.pf_key)
           inner join
project_dimension e on
(a.project_key=e.project_key)
           inner join
flex_field_dimension ffd on
(a.tracker_key=ffd.tracker_key)
           and
ffd.name='Date'
           and
c.title='TestDateTracker'
           and
get_artifact_date_value(ff.flex_fields,ffd.flex
00:00:00'
--where a.effective_from <=
sysdate and a.effective_till >
to_date('2012-04-01', 'YYYY-MM-DD')
GROUP BY
b.date_of_trans,e.title,c.title,d.title,a.prior
ORDER BY
b.date_of_trans,e.title,c.title,d.title,a.prior

```

**Use case 5: Multi-select flex fields**

Suppose you want to retrieve data based on the following assumptions:

- Multi-select flex field name: 'Multiselect'
- Value: 'M12,M11'
- Tracker name: 'TrackerTest'
- Conditional parameter: 'ALL'



**Note:**

- You can change the field name, tracker title and values based on the data in your system
- Flex field name, value and tracker title that are used in the SQL filter conditions are case sensitive
- If you want to select all the values in User flex field, then pass 'ALL' as the conditional parameter
- If you want to select any value, then pass 'ANY' as the conditional parameter

**Sample query**

```

SELECT          b.date_of_trans
  "Date",

e.title "Project",
                c.title
  "Tracker",

d.title  "Planingfolder",
        'P' ||
a.priority "Priority",
        count(distinct
a.artifact_key)
FROM
artifact_transaction_fact a
inner join date_dimension b on
(a.trans_date_key=b.date_key)
        inner join
artifact_flex_fields ff on
(ff.artifact_flex_fields_key =
a.flex_fields_key)
        inner join
tracker_dimension c on
(a.tracker_key=c.tracker_key)
        inner join pf_dimension d
on (a.pf_key=d.pf_key)
        inner join
project_dimension e on
(a.project_key=e.project_key)
        inner join
flex_field_dimension ffd on
(a.tracker_key=ffd.tracker_key)
        and
ffd.name='Multiselect'
        and
c.title='TrackerTest'
        and
get_artifact_multiselect_value(ff.flex_fields, fi
IS NOT NULL
GROUP BY
b.date_of_trans,e.title,c.title,d.title,a.prior
ORDER BY
b.date_of_trans,e.title,c.title,d.title,a.prior

```

**Use case 6: Filter User flex fields**

Suppose you want to retrieve data based on the following assumptions:

- User flex field name: 'Select User'

- Value: 'user1,user2'
- Tracker name: 'TrackerTestUser'
- Conditional parameter: 'ALL'


**Note:**

- You can change the field name, tracker title and values based on the data in your system
- Flex field name, value and tracker title that are used in the SQL filter conditions are case sensitive
- If you want to select all the values in User flex field, then pass 'ALL' as the conditional parameter
- If you want to select any value, then pass 'ANY' as the conditional parameter

**Sample query**

```

SELECT
    b.date_of_trans "Date",
    e.title "Project",
    c.title "Tracker",
    d.title "Planingfolder",
    'P' || a.priority
"Priority",
    count( distinct
a.artifact_key) "TotalArtifacts"
FROM artifact_transaction_fact
a inner join date_dimension b on
(a.trans_date_key=b.date_key)
    inner join
artifact flex_fields ff on
(ff.artifact_flex_fields_key =
a.flex_fields_key)
    inner join
tracker_dimension c on
(a.tracker_key=c.tracker_key)
    inner join pf_dimension d on
(a.pf_key=d.pf_key)
    inner join project_dimension e on
(a.project_key=e.project_key)
    inner join
flex_field_dimension ffd on
(a.tracker_key=ffd.tracker_key)
    and ffd.name='Select User'
    and
c.title='TrackerTestUser'
    and
get_artifact_user_value(ff.flex_fields,ffd.flex
IS NOT NULL
--where a.effective_from <=
sysdate and a.effective_till >
to_date('2012-04-01', 'YYYY-MM-DD')
GROUP BY
b.date_of_trans,e.title,c.title,d.title,a.prior
ORDER BY
b.date_of_trans,e.title,c.title,d.title,a.prior

```

**Use case 7: Filter Text flex fields, case-insensitive search**

Suppose you want to retrieve data based on the following assumptions:

- Flex field name: 'TEXT'
- Value contains: 'cre'
- Tracker name: 'TestDateTracker'


**Note:**

- You can change the field name, tracker title and values based on the data in your system
- The `upper()` function can be replaced with the `lower()` function appropriately for case-insensitive search

**Sample query**

```

SELECT
    b.date_of_trans "Date",
    e.title "Project",
    c.title "Tracker",
    d.title
"Planingfolder",
    'P' || a.priority
"Priority",
    count( distinct
a.artifact_key) "TotalArtifacts"
FROM artifact_transaction_fact
a inner join date_dimension b on
(a.trans_date_key=b.date_key)
    inner join
artifact_flex_fields ff on
(ff.artifact_flex_fields_key =
a.flex_fields_key)
    inner join
tracker_dimension c on
(a.tracker_key=c.tracker_key)
    inner join pf_dimension d on
(a.pf_key=d.pf_key)
    inner join project_dimension
e on (a.project_key=e.project_key)

    inner join
flex_field_dimension ffd on
(a.tracker_key=ffd.tracker_key)
    and
UPPER(ffd.name)=UPPER('TEXT')
    and
upper(c.title)=upper('TestDateTracker')
    and
upper(get_artifact_text_value(ff.flex_fields, f
like upper('cre%'))
--where a.effective_from <=
sysdate and a.effective_till >
to_date('2012-04-01', 'YYYY-MM-DD')
GROUP BY
    b.date_of_trans,e.title,c.title,d.title,a.prior
ORDER BY
    b.date_of_trans,e.title,c.title,d.title,a.prior

```

***TrackerInitialJob - parallel processing***


A parallel processing feature has been introduced in TeamForge 8.1 Patch 1 to improve the performance of the tracker initial load job.

A performance observation was conducted on the tracker initial load job when fetching data from TeamForge and loading them to datamart. For more information on the performance observation, see [TrackerInitialJob performance observation](#).

Based on the result, a parallel processing feature has been introduced to reduce the data processing time. For the observation results on parallel processing, see [here](#).

To enable parallel processing:


1. Change the JVM heap size by setting the Xmx value to `-Xmx2048m` in the `ETL_JAVA_OPTS` site options token (`/opt/collabnet/teamforge/etc/site-options.conf`).

 **Important:** To make this a permanent change, you must modify the above setting before installing the product. If you do not want to make this a permanent change due to hardware resource constraint or for any other reason, then change the value in the `set-env.sh` file located in `runtime/conf` directory as below.

- a) Open `set-env.sh` from `<teamforge-installer-base-dir>/runtime/conf` directory.
- b) In `ETL_JVM_OPTS` site options token, change the Xmx setting to `-Xmx2048m`.

2. Restart the CollabNet ETL service.

```
setenforce 1
/opt/collabnet/teamforge/bin/teamforge restart
```

 **Warning:** If the heap size (Xmx) is not set to 2048 MB for the ETL service, the tracker initial load job is more likely to throw an exception due to insufficient memory.

### Upgrade CLI reports to the latest version on Red Hat or CentOS

Use these instructions to upgrade the CLI reports to the latest version.

Though life cycle metric reports are released as part of regular TeamForge product releases (including patch releases), new reports are, at times, made available to customers as and when they are ready. When a new report is available, you can upgrade your CLI reports to the latest version.

1. Log on to TeamForge application server as root.
2. Install the latest `ctf_clireports`.

```
yum install ctf_clireports
```

3. Run the `post-install.py` script.

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

4. Run the following command and verify the CLI reports installation. Observe the version and release information.

```
rpm -qi ctf_clireports
```


### Control project access

To control who can access your project, consider the purpose of your project and the appropriate type of user.

#### Control access by user role

Project administrators use existing global project roles or create and assign roles to project members to define what those project members can do in the project.


If you choose to use existing global project roles, you can quickly assign relevant roles to your project members. It is a good practice to create a new role in your project only when a suitable global project role is not available.

 **Note:** If your project is a subproject of any other project, your project may have inherited some roles from the parent project. To work with the properties and permissions associated with inherited roles, you must go to the parent project where those roles are specified.

#### Create a role


A role defines the applications that project members with that role can use, and the specific things project members can do in each application.

Any project administrator can create and assign a role.

 **Tip:** It is a good idea to check the existing global project roles via **Project Admin > Permissions > Roles > View: global project roles** before creating any new role in a project.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin Menu**, click **Permissions**.
3. On the **ROLES** tab, select **View: Roles Created For a Project**.
4. Click **Create**.
5. On the **Create Role** page, write a name and description for the role.
6. To let this role be automatically added to any private subprojects of this project, clear the **PREVENT INHERITANCE** option.

By default, roles are inherited by public and gated subprojects, but not private subprojects.

 **Tip:** If this project is a subproject of any other project, you may already have inherited some roles.

7. To allow project members to be able to request this role, select **PROJECT MEMBERS CAN REQUEST THIS ROLE**.


Project members can submit requests for requestable roles, which the project administrator can approve or reject.


 **Tip:** Select **GRANT AUTOMATICALLY ON REQUEST** to skip the need to approve or reject a request.

8. Click **Create**.

The role is created. The **Edit Role** page appears.


9. For each application listed on the **Role Permissions** page, select the permissions and resources you want to make available to users with this role.

 **Tip:** You can specify the permissions for Binaries and integrated applications here too.

 **Note:** You can specify access to individual top-level folders, but not to specific subfolders. However, in the case of documents, you can specify access to individual subfolders as well.

10. Click **Save**.

The role is created. You can assign it to project members at any time.


 **Tip:** Based on the earlier settings, a project member may be able to submit a request for the role.

### *Create a project administrator role*


The project administrator is responsible for managing the project's users and roles.

The project creator is assigned the Founder Project Admin role, a special role granting all project and application administration permissions for the project. You can transfer the Founder Project Admin role to another user.

If the project creator is a CollabNet TeamForge administrator, no Founder Project Admin is created.

 **Note:** By default, project administrators do not have application administration permissions, such as Tracker Admin or Task Admin. Application administration permissions can be included in a project administrator role, but must be assigned separately.

1. On the **Role Permissions** page, select **Project Admin Permissions**.

 **Note:** If this is an inherited role, you can not edit the permissions associated with it. You can edit the project members and user groups to whom this inherited role is assigned.

- If you want the role to contain only the project administrator permissions to manage users and roles, **Project Admin Permissions** is all you need to select.
- If you want the role to contain additional application administration or other permissions, check the additional permissions.


2. Click **Save**.

All project members assigned this role have project administrator permissions to manage the project's users and roles.


### Change a role

If users need to do things that are not allowed by a role you have assigned to them, you may need to change the permissions associated with that role.

When you edit a role, all project members with that role get the updated permissions automatically.

 **Note:** If your project is a subproject of any other project, your project may have inherited some roles from the parent project. To work with the properties and permissions associated with inherited roles, you must go to the parent project where those roles are specified.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin** menu, click **Permissions**.
3. From the list of roles created for this project, click the role you want to edit.

 **Note:** You can assign a role to direct project members and user groups of the project regardless of whether the role belongs directly to this project or is inherited from a parent project.


4. On the **Role** page, make the changes you need.
  - To edit the title or other role details, click **Edit**. Make the required changes and click **Update**.
  - To edit the role's permissions, choose an application from the left side of the page and select or deselect permissions and resources.
  - To edit the project members to whom the role is assigned, click **Assigned Project Members**.
  - To edit the user groups to whom the role is assigned, click **Assigned Groups**.
5. Click **Save**.

The role is modified.

### Give roles to a project member


A project member can have any number of roles. As project administrator, you must assign each project member's roles.

Permissions are cumulative. The project member has all of the access permissions allowed by all of the assigned roles, plus any permissions that may have been assigned globally using application permissions.

 **Note:** If your project is a subproject of any other project, you may have inherited some roles, project members or user groups from the parent project. You can assign any role to any project member, regardless of whether it's a global project role, or the role belongs directly to the project, or is inherited from a parent project.


1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin** menu, click **Permissions**, then click the **User-Role Matrix** tab.

Observe the users listed on the left and all the available roles (global, direct and inherited) on the right.

 **Note:** Users can be assigned global, direct and inherited roles. Inherited users have all the permissions that they have in the parent project.

- If you assign another inherited role to an inherited user, that user gets both the roles.
- If you assign a role to a user in a project, that user becomes a current member in that project.

3. Select roles for each project member.

 **Remember:** A user's license type also influences what the user can see and do on your site. A user's license type supersedes any role assignments. Ask your site administrator how many licenses of each kind are available for your users. For more information, see [How do TeamForge licenses work?](#) on page 948.

The USER\_ROLE MATRIX page is also equipped with a smart search (filter) function that makes it easy to filter a user by name and assign roles. Use it to quickly search for a user to which you want to assign one or more roles.

The screenshot shows the 'Project Admin Menu' on the left with 'Permissions' selected. The main area displays 'Project Members (3)' with a table. A search box with 'mark' is active over the table.

NAME ▲	MARK X	GLOBAL PROJECT ROLES			
Mark George (mgeorge)	<input type="checkbox"/>				
Mark Lovell (mlovell)	<input type="checkbox"/>				
Mark Phippard (markphip)	<input type="checkbox"/>				

4. Click **Save**.

The roles are now assigned to each project member.


*Give a role to multiple project members*


A role can be assigned to many users at once.

The user-role matrix provides a convenient way to add project members to a role, but it can become unwieldy if the project has a large number of users or roles. When that is the case, try assigning roles to multiple users at one time.


1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin** menu, click **Permissions**.
3. Click the name of the role that you want to assign to project members.
4. On the **Role Permissions** page, click the **Assigned Project Members** tab.

The **Assigned Project Members** page shows all users who currently have the role.

 **Tip:** You can assign the role only to the direct project members of the project.

 **Remember:** A user's license type also influences what the user can see and do on your site. A user's license type supersedes any role assignments. Ask your site administrator how many licenses of each kind are available for your users. For more information, see [How do TeamForge licenses work?](#) on page 948.

5. Click **Add** and use the **Find a User** window to move your desired users into the **Selected Users** list.

 **Tip:** You can search by full or partial user name or full name to find the right project members.

6. Click **OK**.

The project members are now assigned the role.

*Handle a role request from a project member*


A project member in CollabNet TeamForge can ask to be granted a role. As the project administrator, it's up to you to approve or reject such requests.


When a CollabNet TeamForge project member submits a request for a role in a project, the request is placed in the **User Membership** section of the **Project Administration** page, pending approval by a project administrator. The request is also displayed in the **Items Pending My Approval** section of each project administrator's **My Page**.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin** menu, click **User Membership**, then click the **Pending Requests** tab.

3. Under **Role Requests**, select the user whose role request you want to approve or reject.

- Click **Approve** to approve the request and assign the role to the user.
- Click **Reject** to deny the request.


 **Tip:** To view the permissions granted via the role, click the role name, if required.

 **Note:** Before approving or rejecting a role request, you can check the roles already assigned to the user by clicking the user name and selecting the **Roles** tab.

The user receives an email notification when the request is approved or rejected.

#### *Assign a global project role on request*

As a project administrator you can edit a requestable global project role for your project. You can update the settings to immediately assign the role to a project member on request.

 **Note:** Only site administrators or restricted site administrators with "Role-Edit" permission can edit global project role details. Project administrators can only edit requestable global project role's grant settings for their projects.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin** menu, click **Permissions**.
3. From the list of global project roles, click the requestable global project role you want to edit.
4. On the **Edit Role** page, click **Edit**.
5. Select the **Grant Automatically on Request** option and click **Save**.  
Selecting this option enables the project member requesting the role to be assigned the same immediately, without waiting for an approval.

The role is modified for your project.

#### *Assign roles to a user group*


Enable multiple users to do something all at once by giving their group a role.

A user group can have any number of roles. Each member of the user group has all the access permissions allowed by all of the assigned roles, plus any permissions that may have been assigned by other methods, such as application permissions or individually assigned roles.

Roles and the associated permissions can be inherited. If your project is a subproject of any other project, you may have inherited some roles or user groups. You can assign any inherited role to any user group.


1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin** menu, click **Permissions**, then click the **Group-Role Matrix** tab.

Observe the user groups listed on the left and all the available roles (global, direct and inherited) on the right.

 **Note:** User groups can be assigned global, direct and inherited roles. The inherited user groups have all the permissions that they have in the parent project.

- If you assign another inherited role to an inherited group, the members of that group get both the roles.
- If you assign a role to a user group in a project, that user group becomes a direct user group in that project.

3. Select the roles you want for each user group and click **Save**.

 **Restriction:** When you give a group access to a CVS or Wandisco Subversion repository, members of the group can view the repository but cannot do repository actions, such as commit and update. You must assign those permissions to users individually.

The roles are now assigned to each user group.


#### *Assign user groups to a role*

To manage permissions for a lot of groups or roles at once, try assigning user groups to roles.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin** menu, click **Permissions**.




- To display existing global, direct or inherited roles, click the **Global project roles**, **Roles Created For a Project** or **Roles Inherited From Parent Project** option of **View**.


 **Note:** If your project is a subproject of another project, you may have inherited some roles or user groups. The inherited role details like role name, description and the source project name are listed.

You can assign direct user groups to a global/direct/inherited role.

- Click the name of the role that you want to assign to user groups.
- On the **Role Permissions** page, click the **Assigned Groups** tab. The **Assigned Groups** page shows all user groups that are currently assigned the role.

 **Tip:** You can assign the role only to the direct user groups of the project.

- Click **Add**.
- Type some of the group's name in the **Name (search)** box and click **Find**.
- In the **Find a user group** window, select the user groups that you want to add, and click **Finish**.

 **Restriction:** When you give a group access to a CVS or Wandisco Subversion repository, members of the group can view the repository but cannot do repository actions, such as commit and update. You must assign those permissions to users individually.

The user groups are now assigned the role.

#### *View users and user groups assigned to a role*

Before adding a role to a project member or user groups, see all the users and user groups who are assigned that role through inheritance.

A user or user group can have any number of roles. Roles and the associated permissions can be inherited via project hierarchy or project groups.

- Click **Project Admin** from the **Project Home** menu.
- On the **Project Admin** menu, click **Permissions**.
- In the **View** drop-down, select **Global project roles** and from the results, a specific role.
- In the **Assigned Project Members** tab, click the **View** drop-down and make a selection.
  - Direct Members** displays the users who are directly assigned the role in the project.
  - Inherited Members** displays the users who inherit the role from parent projects as well as project groups.
- In the **Assigned Groups** tab, click the **View** drop-down and make a selection.
  - Direct User Groups** displays the user groups that are directly assigned the role in the project.
  - Inherited User Groups** displays the user groups that inherit the role from parent projects as well as project groups.

The roles are now assigned to each user group.

#### **Control access by user class**


To avoid having to create and assign a lot of similar roles for individual users, give access to applications to whole classes of users whenever possible.


For each application (tasks, documents, file releases, trackers, and discussion forums), you can assign permissions globally based on user type.

For example, if you know that you want all project members to be able to view and submit to all project trackers, set the application's permissions to reflect this. You need to configure these settings only once.

- Click **PROJECT ADMIN** from the **Project Home** menu.
- On the **Project Admin** page, configure your project access settings and click **Next**.
- On the **Edit Default Access Application Permissions** page, click the + symbol to expand the section for which you want to assign permissions.
- For each application and resource, choose a user type from the drop-down menus.
  - All users of the selected type will have the specified permissions: view, submit/view, or post/view.

- For discussion forums and trackers, you can also specify submit or post permissions.

 **Note:** You can specify access to top-level folders, but not to specific subfolders.


 **Note:** If you want to control access to an application or resource that is not displayed on the **Edit Default Access Application Permissions** page, you can do so using role-based access control.

## 5. Click **Finish**.


### Control access by project type

Projects can be open only to project members, open to everyone in the world, or something in between.

By default, all new projects are created as private projects, accessible only to project members. Your system administrator can change the default access level for new projects.

 **Important:** Users who do not have access to a project cannot see it on the **Home** page, in the **All Projects** list, or in search or reporting results.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin Menu**, click **Permissions**, then click the **DEFAULT ACCESS PERMISSIONS** tab to see the project's current access setting.
3. Click **Edit**.
4. On the **Edit Default Access Permissions** page, select the kind of access you want to allow to your project.
  - **Private** - Project members only.
  - **Gated community** - Project members and unrestricted users.
  - **Public** - All users.

 **Note:** If the option is not available, your system administrator has prohibited changing the access levels of projects on the site.

### Let users see other users' roles


You can enable some project members to view the roles assigned to other project members.

For example, if your project includes both core team members and consultants, you may want to restrict full visibility of user details to the core team members.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin Menu**, click **Permissions**.
3. On the **ROLES** tab, click the role you want to edit.
 

For example, if you have divided project members into a "Core Team" role and a "Consultant" role, click the "Core Team" role.
4. Under **Project Admin Permissions**, select **View User Membership**.


Now only project members who have the role you edited can see the roles held by other project members.

 **Note:** They still can't see the actual permissions included in those roles, unless they have Project Admin status.


### Lock or unlock a project


To ensure that no changes occur in a project while you are collating or migrating project data, lock the project. You must have project administration permissions or be a site administrator to lock or unlock a project.

To lock or unlock a project in TeamForge, go to **Project Settings** and lock/unlock the project.

 **Note:** A locked project does not allow any member (including project administrators and site administrators) to make any changes to the project. Besides that, a locked project can not be set as the parent project for any other project and tasks like adding, editing or deleting integrated applications are also not allowed.

Click **PROJECT ADMIN** from the **Project Home** menu.

The project is locked or unlocked as desired. The lock icon  **Locked** appears on all the project pages while the project is locked.


-  **Note:** If a locked project has an integrated application, for example, project tracker, all the project tracker pages are also non-editable while the project is locked. The user who has access permissions for the integrated application can only view the pages.

### Control access to source code

It's a good idea to make sure your source code can only be used by people who have business with it.

You can control which users can view or commit source code. You can make these distinctions at the repository level or at the path level within a repository.


1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Settings** page, click **Permissions**.
3. Click the role whose access to source code you want to control.

-  **Tip:** If the appropriate role does not exist, you must create it first. See [Create a role](#) on page 292.

4. On the **Edit Role** page, click **Source Code**.
5. Under **Permissions for Specific Repositories**, select a repository in the project and specify this role's access to the repository as a whole.
  - To block users with this role from seeing this repository at all, select **No Access**.
  - To allow users with this role to see everything in the repository but not commit to it, select **View Only**.
  - To allow users with this role to commit to any path in the repository, select **View and Commit**.
6. If you want to control access to a specific path within the repository for users with this role, select **Path-based Permissions**.

See [Who can access source code?](#) and [Wildcard-based access control and path-based permissions in TeamForge](#) for more information.

- a) Click **Add**.
- b) Specify a path in the repository.

-  **Tip:** The path does not have to exist when you specify it. You can set a permission for a path that may be created later.


- c) Select **No Access**, **View Only**, or **View and Commit** for this path.

#### **Note:**

- If none of the available permissions (**View Only**, **View and Commit**, or **Path-based Permissions**) is selected for any repository, and none of the options under **Source Code Permissions** is selected, users with this role do not see the **Source Code** toolbar button.
- If two paths have different permissions, the permissions on the lower-level path take effect. For example, consider a role that has "No Access" set for the path `/branches/version3/users`, but has "View and Commit" access to `/branches/version3/users/vijay`. Users with this role can:
  - Check code in and out of the `vijay` directory.
  - Click down through all the directories in that path, including `users`.

Users with this role cannot check files in and out of the `users` directory or monitor commits to `users`.

7. Click **Save**.

-  **Note:** You can also restrict the information that goes out with commit notification emails. See [Who can access source code?](#) on page 978.

### Control HTML headers in hand-coded project pages

When a HTML page that you created in Microsoft Word or Frontpage looks strange, you may be able to fix it by suppressing HTML head content.

1. Click **PROJECT ADMIN** from the **Project Home** menu.


2. On the **Project Settings** section, select **Show custom web page** and then select **Preserve HTML head content**.
3. Click **Save**.

Pages created in Microsoft Word or Frontpage should render correctly.

### **Allow anonymous Subversion checkouts**

To grant anonymous checkout access while restricting write access to a Subversion repository, set the project's default access permissions to public, provide Source Code View permission to all users, and limit other permissions to specific user classes.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin Menu**, click **Permissions**, then the **DEFAULT ACCESS PERMISSIONS** tab.
3. Click **Edit**.
4. Choose the **Public** project access option on the **Edit Default Access Permissions** page and click **Next**.
5. Under **Application Permissions** on the **Edit Default Access Application Permissions** page, choose **Allow all Site Users and Guests** from the drop-down for Source Code View permission.

 **Note:** You can give all users checkout access to all repositories or a specific repository.

6. For other application permissions, choose a user class based on your access requirement.

Users can now check out from a repository without entering a password.

### **Show or hide an application**

To help users focus on the relevant parts of your project, choose which applications they can see in the **Project Home** menu.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin Menu** page, click **Tools**.
3. Select the applications that you want to be visible to users in your project, and click **Save**.

Only the application buttons you have set to **VISIBLE** appear in the Project Home menu of the project page.

- Removing buttons from this menu does not remove the application from the site. Users with the appropriate permissions can still get to the hidden applications by other means, such as clicking a link in an email. The hidden application button will not appear.
- When an application button is set to **Visible**, but a given user does not have permission to use that application, that user still cannot see the application button. For example, if a user has a role that does not permit access to any source code repositories, that user does not see the **Source Code** toolbar button, even if the button is enabled for the project.
- If the Tasks tool is hidden for a project, the option to run reports on the Task tool is also hidden.
- If you create a project template from a project that has hidden applications, any project you create from that template will have the same applications hidden.

### **Who can post to discussions by email?**

To help reduce the risk of spam or other mischief, you may need to limit the users who can post to your project's discussion forums by email.

To leverage the advantages of community collaboration, you should keep your forums as open as you can. However, some projects require tighter control over who can participate in discussions. TeamForge enables you to balance openness against privacy along a spectrum of choices.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Discussion Settings**.
3. Set the value of **EMAIL POSTING** to one of these choices (listed from most restrictive to least restrictive).

**This option...**

**Allow only forum admins**

**has this effect:**

Only users with the "discussion admin" permission can post by email.

**This option...****Users with roles & permissions****All logged in users****Allow known email addresses only****All site users and guests****has this effect:**

Only registered users with the "topic post" permission can post by email. This is the default. (For discussions marked as private, this is the least restrictive setting possible.)

All users who are registered on the site can post by email.


Users can post by email only if they have explicitly been added to the monitoring list. (This can include people who are not registered site users.)

Anyone can post by email.

If you select a value that's more liberal than the value your site administrator has set for the site as a whole, the site administrator's setting rules. For example, suppose the site administrator has provided that only users with the appropriate role ("Users with roles & permissions") can post by email. If your project requires extra security, you can choose to accept email only from forum administrators. However, you cannot accept email posts from a less restrictive category of users, such as "All logged in users."

**Manage project membership**


You can add or remove project members on their request or on your own initiative.

 **Note:** If your project is a subproject of any other project, your project may have inherited some roles from the parent project. To work with the properties and permissions associated with inherited roles, you must go to the parent project where those roles are specified.


**Add users to a project**

Before a person can work on a project, you have to make them a member of the project.


You can make any registered user on your TeamForge site a project member. You can assign roles to the user at the same time.

 **Note:** Project members, roles and the associated permissions can be inherited and reused in subprojects. If your project is a subproject of any other project, you may have inherited some roles, project members, or user groups.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin Menu**, click **User Membership**.
3. Click **Add**.
4. On the **Add Users** page, find the users you want by one of these methods:
  - Under **FIND USERS**, filter the list of site users eligible to join this project. You can filter by full or partial name or user name.
 



 **Note:** Search text is not case-sensitive.
  - Browse the list of registered users on the site. Sort them by name, user name, email address or membership status.

The inherited project members continue to hold the inherited roles with corresponding permissions, as specified in the source projects.

 **Note:** If a site has a great many users, you must filter them first to narrow down the list. This helps avoid slowing down the system.

5. Select the users you want to add.
6. Under **ASSIGN ROLES**, select the roles you want the users to have.

You can select any global project role, role created just for this project, or inherited role that is available.


-  **Note:** If your project is a child project of another project, the members of the parent project become inherited members of your project. The user roles specified in the parent project are available in your project provided the role inheritance is not prevented. If you assign a role in your project to a user, that user becomes a direct member of your project.
-  **Tip:** If you prefer, you can skip this step and assign roles later on the **Project Admin > Permissions** page. Note that using the **Assigned Project Members** page, you can assign roles only to the direct project members.

7. Save your changes.

- Click **Save** to return to the **User Membership** page.
- Click **Save and Add More** to keep adding users.

### Remove a user from a project

When you remove a user from a project, all items such as tasks and tracker artifacts that were assigned to the user are re-assigned to None.

-  **Note:** If your project is a subproject of any other project, your project may have inherited some users from the parent project. To remove an inherited user, you must go to the parent project where that user is a direct member.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin Menu**, click **User Membership**.
3. From the list of current project members, select the user you want to remove and click **Remove**.

The user is now removed from the project.

### Handle a request for project membership

A registered CollabNet TeamForge user can ask to be a member of a project. As the project administrator, it's up to you to approve or reject such requests.

When a CollabNet TeamForge user submits a request for project membership, the request is placed in the **User Membership** section of the **Project Administration** page, pending approval by a project administrator. The request is also displayed in the **Items Pending My Approval** section of each project administrator's **My Page**.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin Menu**, click **User Membership**, then click the **PENDING REQUESTS** tab.
3. Under **Users Requesting to Join Project**, select the user you want to approve or reject.
  - Click **Approve** to approve the request and add the user to the project.
  - Click **Reject** to deny the request.
4. To view details about the user or add a comment before approving or rejecting the request, click the user name.

This is optional.

The user receives an email notification when the request is approved or rejected.

### Handle a request to leave a project

A TeamForge user who wants to leave a project must submit a request. The project administrator can approve or reject the request.

A request to leave is placed in the **User Membership** section of the **Project Administration** page, pending approval by a project administrator. The request is also displayed in the **Items Pending My Approval** section of each project administrator's **My Page**.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin Menu**, click **User Membership**, then click the **PENDING REQUESTS** tab.
3. Under **Users Requesting to Leave Project**, select the user whose request you want to approve or reject.
  - Click **Approve** to approve the request and remove the user from the project.
  - Click **Reject** to deny the request.
4. To view the user details or add a comment before approving or rejecting the request, click the user name.

This is optional.


The user receives an email notification when the request is approved or rejected.

### Create a project template

To make it easier to start projects, provide project templates based on existing projects.


A project template is used to populate new projects with the structure and configuration of the source project. It can also include the actual contents of the project it is based on, such as tasks, tracker items and documents.

When you create a project template, it is available only to projects on your own site. To make it available more widely, you can export the project template and share it with others. Ask your CollabNet representative for more information about doing this.

 **Note:** The access settings for a project template are the same as the access settings for the project on which the template is based. For example:

- If you create a template from a project that has hidden applications, any project you create from that template will have the same applications hidden.
- If you create a template from a project that is private, any project created from that template will also be private.


1. Set up the project that will serve as the basis for the new project template.

 **Tip:** If you need a clean project template, you may want to create a new project specifically for this purpose.

2. Click **PROJECT ADMIN** from the **Project Home** menu.


3. On the **Project Settings** page, click **Create Project Template**.

4. Write a name and description for the template.

 **Tip:** Consider the uses that you or other project managers might have for this project template, and include keywords in your description that are related to those uses.

5. Choose the items that you want to make available in new projects that are created from this template.

- Every new project includes a **Trackers** component into which you can add new trackers; a **Documents** folder into which you can add more document folders and documents; and so on.

 **Note:** When you select "Artifacts and their dependencies" from optional content, it copies the artifacts along with their associated tags.

- If you wish, you can include the actual trackers from the current project's **Trackers** component; the actual documents from the current project's **Documents** component, and so on.

6. Click **OK**.

The project template is created. The template is available to all users who have the required permissions to create new projects. When the templates are used, the number of projects associated to individual template is displayed in the **projects created** field .

The template name and description appear on the **Template** tab of the **Projects** list, accessible from your personal navigation bar.

### Categorize a project

Organizing projects by categories can help users find what they need on a site quickly and easily.

Project categories express the relationships among projects. If your project is in a category, it is visible to users browsing projects from the main **Project Categories** page.

Your project can be in one or more project categories.

1. Click **PROJECT ADMIN** from the **Project Home** menu.

2. On the **Project Admin Menu**, click the **PROJECT CATEGORIZATION** tab.

A list of all the categories that the project belongs to is displayed.

3. Click **Add**.

The **Select Category** window shows all the available project categories.

4. If you find a project category that your project should be in, select that category and click **Add**.

Your project is now a member of the selected project category. You can repeat this process to add your project to any number of project categories.


### Make other applications available to project members

You can make it easy for project members to use a wide variety of applications and sites from inside TeamForge.

#### Link an external application

To make an application or site outside of CollabNet TeamForge available in your project, create a linked application.

You can create as many linked applications per project as you need.

 **Tip:** Some sites employ page code that disables this feature. For example, some Google apps automatically generate user login data and append it to their URLs, which prevents them from opening in an iFrame.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin Menu**, click **Tools**.  
A list of all the applications in the project is displayed on the **Project Tools** section.
3. Click **Add Tool**.
4. On the **Add Tool** page, enter a name for the linked application.  
This name will appear, with an icon you choose, in the **Project Home** menu.
5. Enter the server location or URL for the linked application.
6. If you are a TeamForge administrator, select whether you want to use single sign on for the linked application.
  - By default, only site administrators can edit (turn on/off) single sign on (SSO) for linked applications. However, you can set the site options token `ONLY_SITE_ADMIN_CAN_EDIT_SINGLE_SIGN_ON` to false to have both site and project administrators turn SSO on and off. For more information, see [ONLY\\_SITE\\_ADMIN\\_CAN\\_EDIT\\_SINGLE\\_SIGN\\_ON](#) on page 760.
  - If you use single sign on, TeamForge users can automatically log into the linked application.
  - If you do not use single sign on, users must log into the linked application using its native authentication system.
7. Click **Choose File** and select a gif, .jpg, or .png file to serve as the new icon for the linked application.  
Make the image 25 pixels wide and 20 pixels high.  
This icon appears with the application name in the **Project Home** menu.
8. Click **Save**.

A button for the linked application is added to your **Project Home** menu. Clicking it launches the application in the main TeamForge project window.

#### Change a linked application

When the use patterns of a linked application change, you may need to change the way the application integrates with CollabNet TeamForge .

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. From the list of applications, click the application you want to edit.  
The **Edit Tool** page is displayed.
3. Edit the name of the linked application, if required.
4. Edit the linked application URL, if required.
5. Enable or disable single sign on for the linked application.  
By default, only site administrators can edit (turn on/off) single sign on (SSO) for linked applications. However, you can set the site options token `ONLY_SITE_ADMIN_CAN_EDIT_SINGLE_SIGN_ON` to false to have both site and project administrators turn SSO on and off. For more information, see [ONLY\\_SITE\\_ADMIN\\_CAN\\_EDIT\\_SINGLE\\_SIGN\\_ON](#) on page 760.
6. Click **Choose File** and select a gif, .jpg, or .png file to serve as the new icon for the linked application.  
Make the image 25 pixels wide and 20 pixels high.



This icon appears with the application name in the **Project Home** menu.

#### 7. Click **Update**.


### **Integrate an external application into your project**

To make an application or site outside of TeamForge available to your users seamlessly from inside your TeamForge project, bring it in as an integrated application.

If the application you want to use is not yet available for your project, ask your site administrator to set it up.

You can use as many integrated applications as you wish, after your site administrator has made them available.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Admin Menu**, click **Tools**.  
The **Project Tools** section displays the list of all the current applications in the project.
3. Click **Add Tool**.
4. On the **Add New Tool** page, select the desired application. and all the relevant paramters to configure the tool are displayed.
  - a) If your site administrator has made it possible, specify a prefix for the resources created by the application you are setting up.  
A prefix enables TeamForge to provide handy links between objects managed by this application and other TeamForge objects. For example, if you are bringing in a blogging application:
    - You can connect a blog post with a TeamForge artifact using a special link called an "association."
    - For each blog post, you can give readers a simplified address, known as a "go URL."

 **Note:** For project-level associations and go URLs the application's prefix is permanent after you save it.
  - b) Set any other configuration parameters you need and save your changes.
    - Click **Save** to return to the **Project Tools** page.

An icon for the integrated application is added to your **Project Home** menu. Clicking it launches the application in the main TeamForge project window.

### **Export an integrated application**


TeamForge site administrators can export an integrated application's variables in XML format and have them available for editing.

1. Go to **My Workspace > Admin**.
2. Click **INTEGRATED APPS** from the **Projects** menu.
3. To export the integrated application, select the integrated application name.
4. From the **View Integrated Application** page, click **Export**.

The integrated application will be exported.

### **Hide / Remove a tool (linked or integrated application)**

To stop making an application or site outside of TeamForge available to your users from inside your TeamForge projects seamlessly, hide or disintegrate an application.

-  **Note:** You may have integrated several external applications per project to maximize the integrated applications feature's utility. However, because each integrated application adds an icon to the project navigation bar, a large number of integrated applications can cause horizontal scrolling. Consider removing the linked or integrated applications that are not in use.
1. Click **PROJECT ADMIN** from the **Project Home** menu.
  2. On the **Project Admin Menu**, click **Tools**.
  3. To hide an integrated application from the list of tools displayed, clear the **Visible** check box and click **Save**. The selected integrated application is hidden for the project. The icon of the hidden linked or integrated application disappears from your **Project Home** menu.



**Warning:** Hiding a linked or integrated application, hides all associations to the integrated application too. It also hides the component type in the Project Home's **Create Component** page.

- You can also click the **Delete** icon to delete a linked or integrated application permanently from your project.



**Warning:** Deleting an integrated application (such as Review Board, EventQ or Binaries) means a permanent removal of the application including all related integrated application data from your project. Exercise caution before deleting an integrated application.

### Add a parent project to your project

You can coordinate work among multiple projects; enable the user, user group and role inheritance by adding a common parent project to several projects.



**Note:** A subproject can have only one parent. You can change or remove that parent at any time.



**Note:** You must be a project administrator or site administrator to add/edit or remove the parent projects.

- Click **PROJECT ADMIN** from the **Project Home** menu.
- On the **Project Settings** page, click **Add Parent**.
- Choose a parent project that it makes sense for your project to belong to, and click **Update**.



**Note:** You cannot add a parent project for the 'Look' project as it is a special project in itself.

### Separate a subproject from its parent project

When a subproject grows beyond its original scope, you may want to make it stand-alone project or move it to a different project hierarchy.

By removing the association with a parent project, you can manage the subproject as a separate project.



**Note:** Only one parent project can be selected for a subproject. However, the parent project can be changed or removed, as required.

- Click **PROJECT ADMIN** from the **Project Home** menu.
- On the **Project Settings** page, click **Edit Parent**.
- Change the parent project as required.

You can be a project administrator or a site administrator to change or remove a parent project. As a project administrator, you can remove or change a parent project only if you have administrator permissions for both the projects that are being linked.



**Note:** A parent project can be removed or changed only when its members, user groups and roles are not in use in any other project. In other words, you can not remove/change a parent project while its members, user groups or roles are in use in any other project.

- On the **Choose a Parent Project** page, select the desired parent project and click **Update**.



**Note:** If the project hierarchy exists, the project and its subprojects are moved only under the project from which members, user groups or roles are inherited. If project hierarchy does not exist and no inheritance is in use, the project is made a 'Root' project.

- Click **Remove Parent** in the **Project Settings** page to remove the association with parent project.
- On the pop-up message box, click **OK**, if you wish to make the project a 'Root' project.

The project hierarchy is changed or removed in accordance with role based access control and inheritance rules.

### Require a wiki version comment

To help keep track of changes to a Wiki page, require users to record a comment about their changes when they save the page.

- On any project page, click **PROJECT ADMIN** from the **My Page** menu.
- On the **Project Settings** page, click **Wiki Settings**.

### 3. Select **REQUIRES VERSION COMMENT ON EDIT**.

Now users who edit a page in your project's Wiki must leave a note describing their changes.

#### **Provision hardware for project members**

You can use TeamForge Lab Management to provision physical or virtual hardware on which project members can remotely build and test their code.



**Tip:** For more information about the new features in TeamForge Lab Management 2.2, see [this introductory video](#). (You'll need a free open.collab.net membership to watch it.)

#### **Set up hardware for your team to use**

When you set up Lab Management, your team members can use TeamForge to access their own virtual machines for developing and testing.

1. Go to your Lab Management project.
2. Select **Project Admin** from the **Project Home** menu.
3. On the **Project Admin** menu, click **Tools**.
4. Click **Add Tool**.
5. Select **Other** from **Select Tool Type** drop-down.
6. Enter the display name as **Lab Mgmt**.
7. Select the **Show in Toolbar** checkbox.
8. Enter the server location or URL of your Lab Management server.
9. If you are a TeamForge administrator, select whether you want to use single sign-on for the linked application.
  - If you use single sign-on, TeamForge manages authentication for Lab Management, and users don't have to log into Lab Management after they have logged into TeamForge.
  - If you do not use single sign-on, users must log into Lab Management using its native authentication system.
10. Click **Browse** and select this icon for the linked application:



This icon will appear with the name **Lab Mgmt** in the project navigation bar.

11. Click **Save**.

A **Lab Mgmt** button is added to your **Project Home** menu.. Clicking it launches the application in the main TeamForge project window.

#### **Reallocate a system**

To enable a project member to use a system, you must reallocate the system to that user.

You must be a Domain Administrator or Project Administrator to free a host that does not already belong to you.

1. Notify the user you're taking the system from through email or some other means, so that they don't lose any important work which they may not have saved on that system.
 

The user whose system was deallocated will get an email after the deallocation, but it is a courtesy to notify the user beforehand.
2. On your project's **Project Home** page, click the host you want to reallocate.
3. Click **Free Host**.
 

The host is now free and can be allocated by anyone in the project. TeamForge Lab Management emails the user that their system has been deallocated by an administrator.
4. To discourage other users from allocating it, put a note in the host's **Description** stating something like "This host reserved for Jane."
5. Notify the user who is going to allocate the new system that their system is now ready.

#### **Edit a project**

As a project administrator in TeamForge Lab Management, you can edit certain properties for your project.

1. On the project home page for the project that you wish to edit, click **Edit Project**.

2. Edit the following parameters:

**Project Summary**

A brief, one-line summary of the project.

**Project Description**

A more detailed explanation of the project.

**Project MOTD (Message of the Day)**

The Project MOTD is displayed to all users in your projects. This message also may display to users when they log in to client nodes in your project.

**Project-specific host allocation time limits**

In this section, project administrators can control how long users in their project can allocate hosts for. If you set this value to 0 (zero), there is no limit on the time a host can be allocated.

If you reduce this value, users' hosts in your project may be deallocated. If deallocations will occur as a result of your lowering of the maximum allocation time, you will be warned which systems will be affected, and given a chance to change your mind.

**Delete this Project/Undelete this Project**

If your project is not deleted and has no hosts, you will be given the option to delete this project from TeamForge Lab Management. If your project is deleted, you will be given the option to undelete it.


**Note: only Domain Admins will be able to delete and undelete projects**

**Dictate which profiles can be used in your project**

To control which operating system profiles the users in your project can build hosts with, adjust the settings in your profile library.

1. On the **Profile Library** page, click **Add/Remove Profiles**.
2. In the list of all potential profiles that you can add to your project, select the profiles you wish to be allowed in your project and click **Save Changes**.

Your changes take effect immediately.

 **Note:** To revert any changes you have made, click **Cancel** or **Reset**.

**Deliver a build**


To get your software into users' hands, upload your build to the Project Build Library.

*Get the PBL upload client*

You must download the pbl.py script to transfer files into the Project Build Library.

The PBL upload client is free, open-source, and freely modifiable and distributable.


Download the PBL upload client from <http://cubit.open.collab.net/pbl/>.

 **Note:** The API operations are fully documented, for users who might want to develop their own PBL upload client.


*Upload a file*

To upload a file, run the `pbl.py upload` command.

In this example, we upload a file named `Release.zip` from our local machine into the public area of the project `myproject`, in the directory `/foo/bar/baz/`.

 **Note:** If any part of the requested path does not already exist, `pbl.py` creates the intermediate directories.

Run the `pbl.py upload` command like this, substituting the correct values for your situation.



```
pbl.py upload --api-user=username --api-key=713cdf90-2549-1350-80c3-2d0bcf9a1697 --api-url http://$external_host/cubit_api/1 --project=myproject -t pub -r /foo/bar/baz -d "This is the description." /home/Release.zip
```



**Tip:** Wildcards are accepted in the filename argument. If the file argument is a directory, or a wildcard which includes one or more directories, `pbl.py` recursively uploads all the subdirectories underneath the parent. All the files in the recursive upload get the same description.

Once this operation has completed, you can download this file from `https://$external_host/myproject/pub/foo/bar/baz/Release.zip`.

For other options to `pbl.py`, run `pbl.py help upload`.

#### *Change the description of a file*

You can change the description associated with a file or directory without changing the file itself or the md5 checksum of the file.

In this example, you have already authenticated and saved your user name and key credentials in your home directory.

Run the `pbl.py changedesc` command like this:

```
pbl.py changedesc -l http://$external_host/cubit_api/1 --project=myproject -t pub -r /foo/bar/baz/Release.zip -d "This is the new description"
```

#### *Move a file*


With the `pbl.py move` command, you can move files or directories within a project, or even between projects.

The syntax for this command is a bit different than the rest of the commands, because the other commands only operate on one project or file or directory at a time, and that is not the case the the move operation.

To move a file, run the `pbl.py move` command with these options.


In the simplest case, we move a file, or a directory and all its contents, from one name to another.

<b>--srcproj projname</b>	The name of the project the source file is located in.
<b>--destproj projname</b>	The name of the project to move the file to. If left blank, defaults to value of <code>--srcproj</code> .
<b>--srcpath path</b>	The path to the file or directory to move.
<b>--destpath path</b>	The destination path for the file or directory specified in <code>--srcpath</code> . Two important things to note about this option: <ul style="list-style-type: none"> <li>• If you specify a path which does not exist, that path will be automatically created for you as part of the move.</li> <li>• If the <code>--destpath</code> parameter ends with a slash ("/"), the destination will be assumed to be a directory. If it does not end with a slash, the destination will be assumed to be a file. An example of this behavior is below. This is approximately how the UNIX "mv" command behaves.</li> </ul>
<b>--srctype {pub priv}</b>	The visibility type of the source file, either <code>pub</code> or <code>priv</code> .
<b>--desttype {pub priv}</b>	The visibility type of the destination file, either <code>pub</code> or <code>priv</code> .
<b>--force</b>	If the destination file exists, the <code>--force</code> option must be used to replace it.

 **Note:** Because `destpath` does not end with a slash -- `/foo/bar/baz/Release_old.zip` -- the last component of the path is interpreted as "a file named 'Release\_old.zip'".

```
pbl.py move -l http://$external_host/TeamForge Lab Management_api/1 --
srcprj=myproject --srctype=pub --srcpath=/foo/bar/baz/Release.zip --
destpath=/foo/bar/baz/Release_old.zip
```

To move a file from one project to another, and also change it from public to private, run the command like this.

 **Note:** Because `destpath` ends with a slash -- `/foo/bar/baz/archive/` -- the last component of the path is interpreted as "a directory named 'archive'".


```
pbl.py move -l http://$external_host/TeamForge Lab Management_api/1 --
srcprj=myproject --destproj=myproject_archive --srctype=pub --desttype=priv
--srcpath=/foo/bar/baz/Release.zip --destpath=/foo/bar/baz/archive/
```

### Allow clouds in a project

To enable project members to allocate hosts from a public cloud in your Lab Management site, you must turn on a setting to allow the inclusion of public clouds in your project.

1. In **Administration > Projects**, click **Edit** for your project.
2. Turn on the **Allow Public Clouds** setting.

If you want to be able to control the hosts your project members can select from, turn this setting OFF. In this case, your project members can only select systems from a cloud your project explicitly owns. Your project's **Allowed Clouds** page lists all the available clouds from which project members can allocate hosts.

 **Note:** If you think that your project needs more systems than those available from a cloud, you need to ask the cloud administrator for your site to increase this number.

### Provision a Continuous Integration server from a Lab Management cloud

The Lab Management Cloud plugin enables Hudson and Jenkins to automatically create slaves from Lab Management clouds.

*Install the Lab Management Cloud plugin for Hudson and Jenkins*

Get the plugin from openCollabNet and upload it using the Plugin Manager.

1. Download the `labmanagement.hpi` file from openCollabNet. (Please check back later for availability.)
2. In the Hudson or Jenkins Plugin Manager page, click the **Advanced** tab.
3. In the Upload Plugin section, browse to the location where you saved the `.hpi` file and click **Upload**.

In the Plugin Manager's **Installed** tab, you should see the Lab Management Cloud plugin enabled.

*Configure the Lab Management Cloud plugin for Hudson and Jenkins*

When you add a new cloud, you set up templates specifying details such as the host type and profile for the nodes that will be provisioned from this cloud.

1. In the Cloud section of the Hudson or Jenkins configuration page, click **Add a new cloud** and select **CollabNet Lab Management**.

Lab Management options are displayed. Here's an example:

### Cloud

CollabNet LabManagement

LabManagement URL  ?

Username  ?

Password  ?


API Key  ?

Project  ?

Cloud  ?

Templates

Name	<input type="text" value="ctf6.1"/>	?
Host Type	<input type="text" value="basic_i386"/>	?
Size	<input type="text" value="1cpu-1gb-60gb"/>	?
Profile	<input type="text" value="ctf6.1-dev-r5.4-i386"/>	?
Labels	<input type="text" value="ctf 6.1"/>	?
# of Executors	<input type="text" value="1"/>	?

2. Enter the URL of the Lab Management Manager node. For example, <https://mgr.cloud.sp.collab.net/>.
  3. Provide the user name and password.  
This user must have access to the Lab Management project where new hosts will be allocated.
  4. Enter the user's API key for the Lab Management web service.
-  **Tip:** Copy and paste the key from the user's Lab Management home page.
5. Click **Test Connection** to make sure that the options you provided are valid.
  6. Specify the Lab Management project where new hosts will be added.
  7. Select a **Cloud** from which new hosts will be allocated.

You can find the list of valid cloud names in the Lab Management Manager interface. For example:

Name	Summary
DevCloud ▼	VMware-server clouds
PublicCloud ▼	Public cloud for all projects
VaporCloud ▼	Short term allocation cloud

8. Click **Add Template** to create a template that Hudson will start.

a) Specify a name that characterizes the template.

The **name** is used to identify the template and is displayed in various parts of the Hudson interface.


b) For **Host Type**, select the type of hardware you want for the node.

c) For **Size**, select the amount of resources -- memory, CPU, and disk -- assigned to this template.

You can find the exact amount of resources for each size in the Lab Management Manager's **Clouds > Allocate From Cloud** page. For example:

d) For **Profile**, select the Lab Management profile to be used for the template.

e) Provide a list of **labels** separated by white spaces.

 **Note:** The names and labels you provide for the templates are used to specify which jobs will use this cloud.

f) Enter a value for **# of Executors**.

This controls the number of concurrent builds that Hudson or Jenkins can perform. So the value affects the overall system load that might be incurred. A good value to start with is the number of processors on your system.

When using Hudson or Jenkins in the master/slave mode, setting this value to 0 would prevent the master from doing any building on its own. Slaves may not have zero executors, but may be temporarily disabled using the button on the slave's status page.

9. Click **Save**.

In the **Manage Nodes** page, you'll see an option to provision a node from Lab Management.



Jenkins » nodes

- [Back to Dashboard](#)
- [Manage Jenkins](#)
- [New Node](#)
- [Configure](#)

S	Name ↓	Architecture	Clock Difference
	<a href="#">master</a>	Linux (amd64)	In sync

Provision via LabManagement ▼

### Run a Hudson or Jenkins job on a host provisioned from Lab Management cloud

When you run a job on a host provisioned from a Lab Management cloud, use a template you defined earlier.

1. In the Hudson or Jenkins configuration page for the job, provide a name and description.
2. Select the **Restrict where this project can be run** option to tie the job to the Lab Management cloud.
3. Enter a label expression.

To always run this job on a specific node, just specify its name. However, when several nodes could be available and you don't want to tie the job to a specific one, enter an expression based on the name and label values of a Lab Management cloud template you configured earlier.

Project name: Elastic Provisioning test

Description: Runs a build on a dynamically provisioned node from a Lab Management cloud.

Discard Old Builds

This build is parameterized

Disable Build (No new builds will be executed until the project is re-enabled.)

Execute concurrent builds if necessary (beta)

Restrict where this project can be run

Label Expression: ctf

4. Specify any other options you want and click **Save**.

When a node is no longer used, it will be brought down and released.

### Set up Tags

Creating tags and tagging items such as documents, artifacts and so on can aid in classification, marking ownership of work items, marking items as milestones, releases and requirements, and so on. Project Administrators can set up tags to be used by project members in a project. However, project members with CREATE/EDIT permissions can create tags, if required.


Tags, once set up, can be used for tagging items such as Documents, Tracker Artifacts, and so on. To start with, TeamForge 17.1 supports tagging for Documents. Tagging will be extended for other objects in due course.

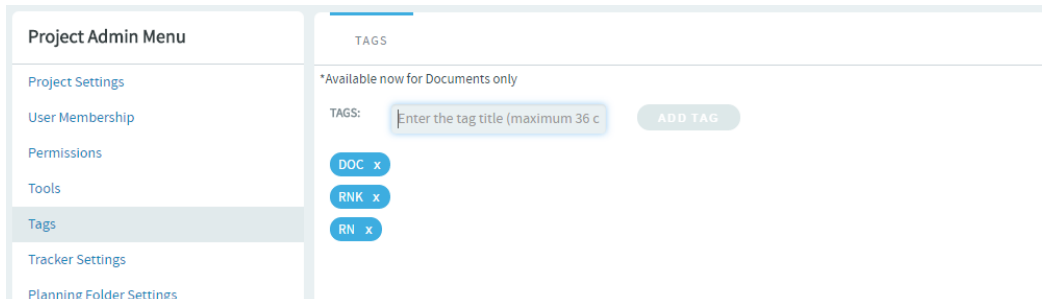
**Note:** While project members with CREATE/EDIT permissions can create new tags in a project, only project administrators can delete tags.

### Set up tags

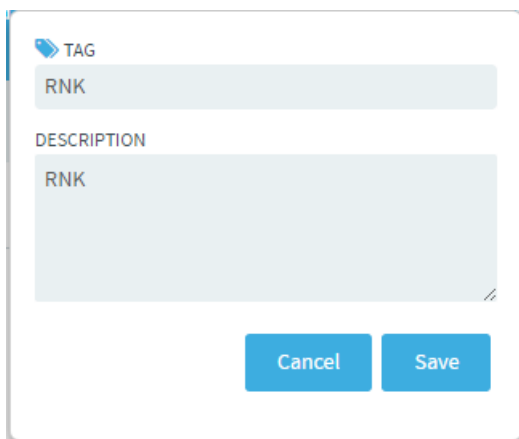
1. Select a project from **My Workspace** menu.

2. Select **Project Home > Tags**.
3. Type a tag title and click **Add Tag**.

 **Note:** A tag can be of 64 characters in length and can contain both alphanumeric and special characters.



4. Click a tag to edit a tag's title and description, edit the tag and click **Save**.



5. To delete a tag, click the "X" mark of the tag. A confirmation message appears. Click **OK** to delete the tag.



**Caution:** When you delete a tag, all associations between the specific tag and TeamForge objects are removed. Exercise caution before deleting tags.

## Support project teams with CollabNet TeamForge

As a site administrator, you can help project managers and members of the site work together effectively throughout the application life cycle.

### Identity Management: Set up TeamForge OAuth 2.0 Authorization Framework

With the new TeamForge Identity Management built on OpenID Connect (OIDC) and OAuth 2.0 authorization frameworks, TeamForge can now act as an ID Provider (IdP). As an IdP, TeamForge can authorize a third-party client application to obtain limited access to its services either on behalf of a Resource Owner (user) or on behalf of the client application itself. This topic discusses TeamForge OAuth architecture, scopes, supported authorization grant types, Federated Identity Management in TeamForge, and so on. This topic also discusses step by step how to add a client application that wants to access TeamForge's services.

### TeamForge OAuth nomenclature

It is assumed that you are familiar with the OAuth and OpenID Connect frameworks, terms and concepts. Here's a few links and references for further reading.

Term	Description
OAUTH	<a href="#">The OAuth 2.0 Authorization Framework RFC 6749</a>
OIDC	<a href="#">OpenID Connect Core 1.0</a>

Term	Description
BEARER TOKEN	<a href="#">The OAuth 2.0 Authorization Framework: Bearer Token Usage RFC 6750</a>
SAMLBEARER	<a href="#">SAML 2.0 Bearer Assertion Profiles for OAuth 2.0</a>
JWE	<a href="#">JSON Web Encryption (JWE)</a>
JWS	<a href="#">JSON Web Signature (JWS)</a>
JWT	<a href="#">JSON Web Token (JWT)</a>
JWA	<a href="#">JSON Web Algorithms (JWA)</a>
JWK	<a href="#">JSON Web Key (JWK)</a>

### TeamForge OAuth overview

In the traditional client-server authentication model, client applications must have access to the Resource Owner's credentials in order to request and access protected resources on the Resource Server. Sharing Resource Owner's credentials such as usernames and passwords to third party applications poses several problems, security being one of the major ones.

Frameworks such as the OAuth 2.0 and OpenID Connect help in mitigating such security risks and hence TeamForge's authentication model has been completely rebuilt on top of OpenID Connect and OAuth 2.0 Authorization frameworks.

### TeamForge OAuth roles

TeamForge OAuth authorization process involves interaction between four entities (roles) such as the user, IdP, SP and client application. The following table lists the roles that an entity can assume in the TeamForge OAuth setup.

Entity/Role	OAuth Term	OIDC Term	SAML Term	Description
User	Resource Owner	End User	User	Human or machine that intends to access a resource of an SP.
Service Provider (SP)	Resource Server	Relying Party	Service Provider, Relying Party	Entity that demands authentication to grant access to some resource.  The SP trusts the IdP to perform the authentication on its behalf.
Client	Client	-	Browser	Application interacting with SP and IdP on the User's behalf.
Identity Provider (IdP)	Authorization Server	Provider	Identity Provider, Asserting Party	Entity that proves identity and authenticates an user.

### OAuth 2.0 protocol flow

The following illustration shows the OAuth 2.0 protocol flow.



### Abstract Protocol Flow

- (A) The client application requests authorization from the Resource Owner.
- (B) The Resource Owner responds to the client with a credential called 'Authorization Grant'.
- (C) The client then redeems the 'Authorization Grant' with the Authorization Server for a valid 'Access Token' (D).
- (D) The client then reaches out to the Resource Server with the Access Token and gains access to the protected resource.

For more information about the OAuth 2.0 protocol flow, see [OAuth 2.0 Protocol Flow](#).

It's recommended that you read [The OAuth 2.0 Authorization Framework RFC 6749](#) and get familiar with **OAuth Roles, Protocol Flow, Access Tokens, Grant Types, Client Types** and so on before you proceed.

TeamForge can act both as an IdP and Service Provider and at times as a client too. The TeamForge Web Application is one of the system defined clients that use TeamForge IdP's OAuth services. For more information about the system defined clients, see [System defined clients](#) on page 323.

### TeamForge API authentication and access tokens

Both TeamForge REST and SOAP APIs use OAuth 2.0 access tokens for authentication. Clients can obtain access tokens from the token endpoint which is located at `/oauth/auth/token`.

#### Access Tokens

The tokens used by the TeamForge API are Bearer Tokens as specified in [RFC 6750](#). This means that it is possible and allowed to share tokens among multiple clients or to have clients pass tokens to intermediate services, which then delegate tokens to TeamForge. TeamForge tokens use a standardized format, [JSON Web Token \(JWT\)](#). However, clients should consider access tokens to be opaque in order to guarantee compatibility with future TeamForge versions. It is the client's responsibility to protect the access token against theft. This means that access tokens should only be transmitted over SSL-secured connections and should not be persisted.

#### TeamForge OAuth scopes

Scopes can be used to restrict which services a token can be used for. Limiting the number of scopes decreases the potential damage that could occur in case an access token is stolen, so it is advisable to restrict the number of scopes to a minimum.

TeamForge currently supports the following list of scopes.

Scope	Description
<b>urn:ctf:services:ctf</b>	CollabNet TeamForge application services. Use this scope to call the TeamForge REST API.
<b>urn:ctf:services:svn</b>	Subversion services. Use this to call the Subversion REST API.
<b>urn:ctf:services:gerrit</b>	Git/Gerrit services. Use this to call the Git REST API.
<b>urn:ctf:services:soap60</b>	SOAP60 services. Use this needed to call the TeamForge SOAP API or the EventQ Reporting API.

### Client ID and Client Secret

In order to use TeamForge's OAuth services, a client application must be registered with TeamForge. Upon registration, the client gets its client credentials: Client ID and Client Secret. For more information about how to get the Client ID and Client Secret, see [Add clients](#) on page 323.

<b>Client ID</b>	The Client ID is a unique identifier used by TeamForge to identify the client application and is used in building the Authorization URLs.
<b>Client Secret</b>	The Client Secret is a unique identifier used to authenticate the client application's identity.

### TeamForge OAuth grant types

When a client requests access token from an Authorization Server, it uses the Authorization Grant, which is a credential that represents the Resource Owner's authorization to access protected resources. Of the four standard grant types defined in the *The OAuth 2.0 Authorization Framework RFC 6749*, TeamForge supports the following three grant types:

- [Authorization Code grant type \(authorization\\_code\)](#) on page 317
- [Resource Owner Password Credentials Grant Type \(password\)](#) on page 319
- [Client Credentials grant type \(client\\_credentials\)](#) on page 320

In addition, TeamForge supports the following custom grant type:

- [Anonymous grant type \(urn:ctf:grant\\_type:anonymous\)](#) on page 320 (rn:ctf:grant\_type:anonymous)


### TeamForge system defined grant types

TeamForge also supports the following system defined grant types. However, these grant types are used by TeamForge's system defined clients to obtain access tokens and are therefore not supposed to be used for other custom clients.

- JSession (urn:ctf:grant\_type:jsession)
- SCM Admin (urn:ctf:grant\_type:scmrequestkey:scmadmin)
- SCM Viewer (urn:ctf:grant\_type:scmrequestkey:scmviewer)
- Token (urn:ctf:grant\_type:token)

### Authorization Code grant type (authorization\_code)

The Authorization Code grant type can be used if the client intends to use TeamForge as its IdP. This grant type, if used by clients, provides an access token for a valid authorization code. This is a redirection based grant type and hence you must add one or more Redirect URIs while adding clients that use this grant type.

 **Note:** Include a comma separated list of Redirect URIs in case you want to set up multiple URIs.

### How to get an authorization code?

Clients (Relying Party) must request the authorization code from the TeamForge IdP URL (/oauth/oidc/authorize) with the following parameters:

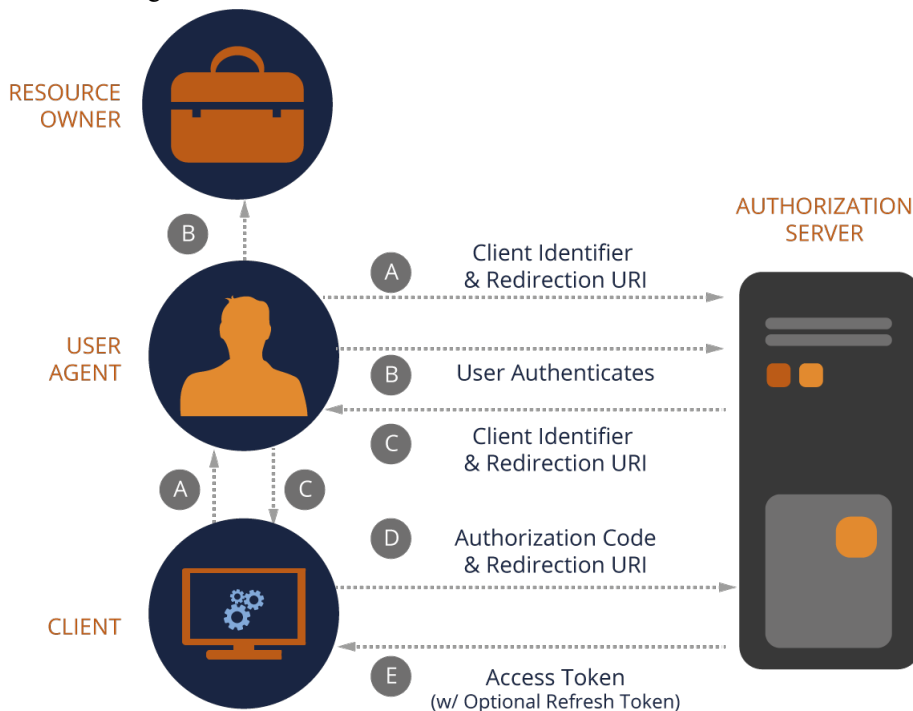
```
scope = open_id (mandatory for OIDC clients)
response_type = code (mandatory for OIDC clients)
client_id = <Client ID generated while registering the client with
TeamForge>
redirect_uri = <URL encoded Redirect URI>
```

After successful authorization, TeamForge IdP provides the authorization code to the redirect URI configured while adding the client. Here's a sample OIDC response after a successful authorization:

```
https://<your_redirect_uri>?code=<authorization_code>
```

Up on getting the authorization code, the client can redeem it with the Authorization Server to obtain an access token.

The following illustration shows the authorization code flow.



### Authorization Code Flow

#### Usage example to obtain an access token

```
# Base URL of TeamForge site.
site_url="https://teamforge.example.com"

# Requested scope (all)
scope="urn:ctf:services:ctf urn:ctf:services:svn urn:ctf:services:gerrit
urn:ctf:services:soap60"

# Client's authorization code
code=<your authorization code>

curl -d "grant_type=authorization_code&client_id=<your Client
ID>&client_secret=<your Client Secret>
&scope=$scope&code=$code" $site_url/oauth/auth/token
```

A successful response will return the HTTP 200 status code and the response body will contain something like this:

```
{
  "access_token" :
  "eyJraWQyOiIwIiwiaWwiYWNxIjoiaUlMyNTYifQ.eyJzdWIiOiJhZG1pbiIsImFlZCI...\" ,
  "token_type" : "Bearer"
}
```

In compliance with the Bearer Tokens specification (RFC 6750), TeamForge expects access tokens to be passed in the Authorization header:

```
GET /resource HTTP/1.1
Host: teamforge.example.com
Authorization: Bearer SAKsa921hjsi...
```

The only exception to this is the SOAP API which expects the token to be passed as part of the SOAP request payload in accordance with the API documentaton.

For more information, see [Authorization Code grant type \(authorization\\_code\)](#) on page 317.

### Resource Owner Password Credentials Grant Type (password)

Clients can use the Resource Owner Password Credentials grant type to get an access token for a given username and password.



### Resource Owner Password Credentials Flow

#### Usage example to obtain an access token

```
# Base URL of TeamForge site.
site_url="https://teamforge.example.com"

# TeamForge authentication credentials.
username="foo"
```

```
password="bar"

# Requested scope (all)
scope="urn:ctf:services:ctf urn:ctf:services:svn urn:ctf:services:gerrit
urn:ctf:services:soap60"

curl -d "grant_type=password&client_id=<your Client ID>&scope=
$scope&username=$username&password=$password"
$site_url/oauth/auth/token
```

You can see this from this example that it is sending an HTTP POST request to the `/oauth/auth/token` endpoint, and the content of the request body contains the following data:

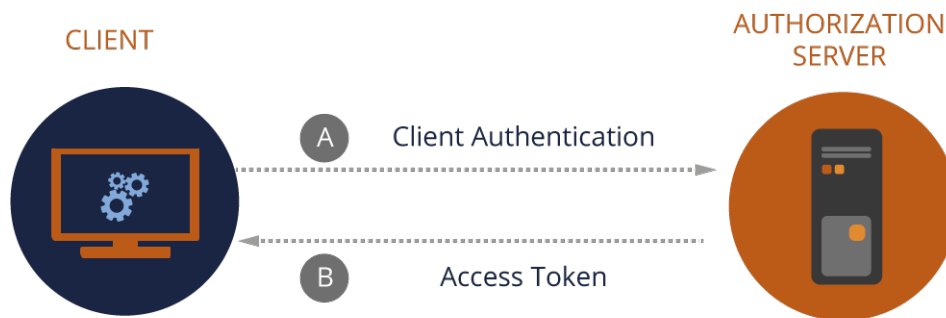
- **grant\_type=password:** Indicates you are providing username and password.
- **client\_id=<your Client ID>:** Client ID generated while adding the client.
- **scope= ...:** Space-separated list of requested scopes.
- **username=value&password=value:** The valid TeamForge user credentials.

For more information, see [Resource Owner Password Credentials Grant Type \(password\)](#) on page 319.

### Client Credentials grant type (client\_credentials)

Clients can use the Client Credentials grant type to get an access token for a given Client ID and Client Secret. Use this grant type only for trusted and confidential clients.

- 👉 **Important:** Clients that use the Client Credentials grant type must be assigned with a valid TeamForge license in order for the clients to be able to access TeamForge resources as OAuth clients.



### Client Credentials Flow

#### Usage example to obtain an access token

```
# Base URL of TeamForge site.
site_url="https://teamforge.example.com"

# Requested scope (all)
scope="urn:ctf:services:ctf urn:ctf:services:svn urn:ctf:services:gerrit
urn:ctf:services:soap60"

curl -d "grant_type=client_credentials&client_id=<your Client
ID>&client_secret=<your Client Secret>
&scope=$scope" $site_url/oauth/auth/token
```

For more information, see [Client Credentials grant type \(client\\_credentials\)](#) on page 320.

### Anonymous grant type (urn:ctf:grant\_type:anonymous)

Use this grant type for clients that need to access TeamForge's public resources.



### Usage example to obtain an access token

```
# Base URL of TeamForge site.
site_url="https://teamforge.example.com"

# Requested scope (all)
scope="urn:ctf:services:ctf urn:ctf:services:svn urn:ctf:services:gerrit
urn:ctf:services:soap60"

curl -d "grant_type=urn:ctf:grant_type:anonymous&client_id=<your Client
ID>&client_secret=<your Client Secret>
&scope=$scope" $site_url/oauth/auth/token
```

### Unsupported grant types

The following grant types are not supported by TeamForge.

- refresh\_token
- urn:ietf:params:oauth-grant-type:saml2-bearer

### Federated Identity Management

Federated Identity Management refers to the process of storing user credentials with an IdP (such as TeamForge, SAML and so on) and having the Service Provider (SP) rely on the IdP to validate user credentials when he/she tries to access its resources.

TeamForge supports federated identity. By default, federated login is disabled in TeamForge. Site Administrators must enable federated login and select an IdP.

1. Log on to TeamForge as a Site Administrator.
2. Select **My Workspace > Admin**.
3. Select **My Page > Identity**.
4. Select the **Federation** tab.
5. Select the **Use Federated Login** check box and select an IdP from the drop-down list.

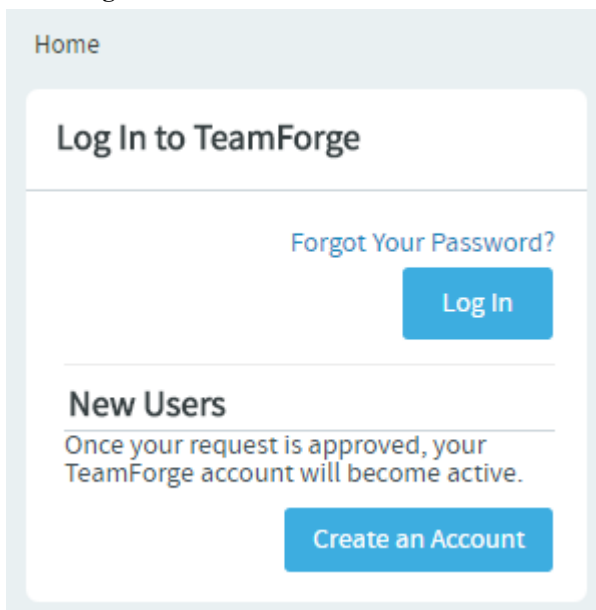
The screenshot shows the 'Identity' section of the TeamForge administration interface. The 'FEDERATION' tab is active. Under the heading 'USE FEDERATED LOGIN:', there is a checked checkbox. Below that, the 'SELECT IDP:' dropdown menu is set to 'teamforge'. At the bottom right of the form, there are 'Cancel' and 'Save' buttons.

6. Click **Save**.

### Logging into TeamForge in a Federated Login Setup

Once enabled, TeamForge Web Application, by virtue of being a system defined client application (see [System defined clients](#) on page 323), follows the OAuth authorization/authentication process to access TeamForge services on behalf of the Resource Owner (user).

1. Click **Login**.



Home

## Log In to TeamForge

[Forgot Your Password?](#)

**Log In**

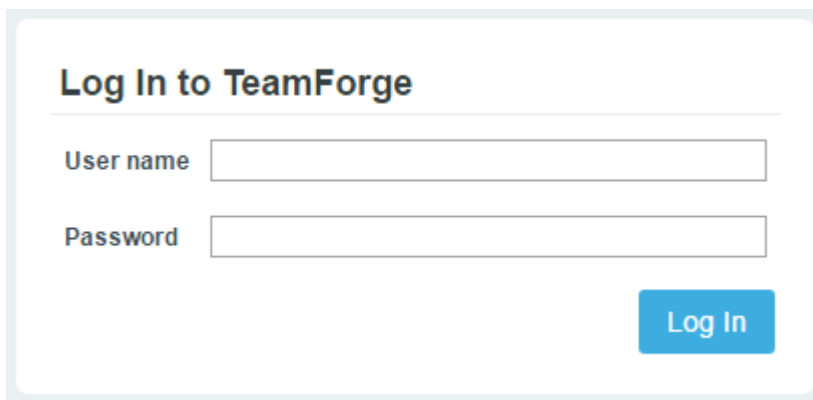
---

### New Users

Once your request is approved, your TeamForge account will become active.

**Create an Account**

You are redirected to an authorization page.



## Log In to TeamForge

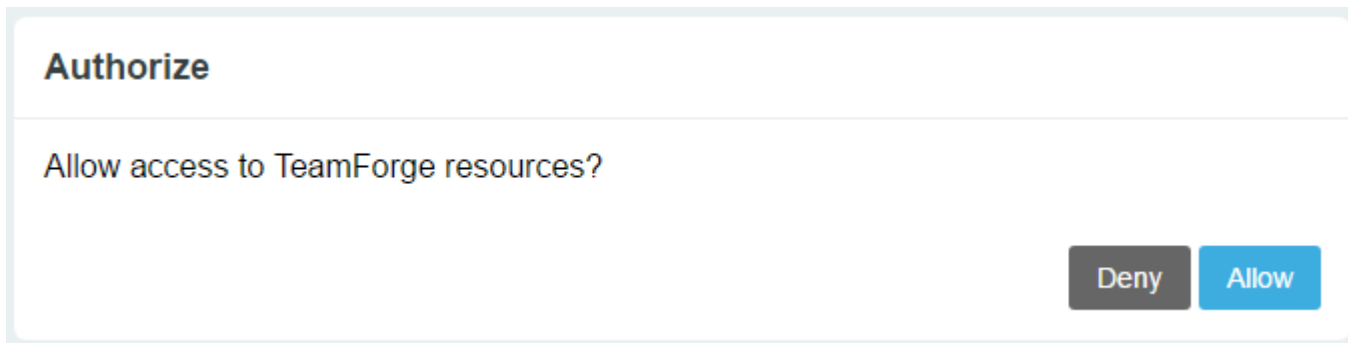
User name

Password

**Log In**

2. Type the user name and password and click **Log In**.

You, as a Resource Owner, are then prompted to either allow or deny the client application (TeamForge Web Application) to access the resources on the Resource Server (TeamForge Application Server).



## Authorize

Allow access to TeamForge resources?

**Deny** **Allow**

3. Click **Allow** or **Deny** to allow or deny access respectively.

## System defined clients

TeamForge OAuth consists of the following system defined clients:

- TeamForge Web Application Client (ctfweb)
- CodeBrowser Client (codebrowser)
- SOAP60 Client (soap60-client0)
- SCM Client (scm-client)
- API Client (api-client)
- App Client (app-client)

These clients have been pre-defined as illustrated in the following table.

System Defined Client	Grant Type	Services	Token Life Time (in seconds)
ctfweb	authorization_code urn:ctf:grant_type:jsession	soap60,ctf,gerrit,svn	3600
codebrowser	urn:ctf:grant_type:anonymous	soap60	3600
soap60-client	password urn:ctf:grant_type:anonymous	soap60	3600
scm-client	urn:ctf:grant_type:scmrequestkey:scmviewer urn:ctf:grant_type:scmrequestkey:scmadmin	soap60,ctf	3600
api-client	password	soap60,ctf,gerrit,svn	3600
app-client	urn:ctf:grant_type:token	soap60,ctf,gerrit,svn	3600

You cannot edit these system defined clients. The following error message shows up if you try to edit these clients.

```
Error editing oauth client: System defined clients are not allowed to be modified.
```

## Add clients

In order to use TeamForge's OAuth services, a client application must be registered with TeamForge. Upon registration, the client gets its client credentials: Client ID and Client Secret.

### Client ID

The Client ID is a unique identifier used by TeamForge to identify the client application and is used in building the Authorization URLs.

### Client Secret

The Client Secret is a unique identifier used to authenticate the client application's identity.

To add a TeamForge OAuth client:

1. Log on to TeamForge as a Site Administrator.
2. Select **Admin** from the **My Workspace** menu.
3. Select **Identity** from the **Projects** menu.
4. Click **Add Client** from the **OAuth** tab.

The **Add Client** dialog box appears.

## Add Client ✕

**CLIENT DESCRIPTION:**

**REDIRECT URIS:**

**GRANT TYPE: \***

**SERVICES: \***

**TOKEN LIFE TIME: \***

5. Enter client information.

The following parameters are required to register a client application with TeamForge:

- **Client Application Name:** Type a name for your client.
- **Redirect URIs:** Include a comma separated list of Redirect URIs in case you want to set up multiple URIs. Redirect URIs are required only for clients that use the Authorization Code grant type.
- **Grant Type:** Select the client's grant type.
- **Services (scopes):** Select one or more TeamForge OAuth scopes.
- **Token Life Time (in seconds):** Type the duration (in seconds) for which the token is valid.
- **License Types:** Select a TeamForge license for the client. This is required only for clients that use the Client Credentials grant type.

6. Click **Save**. Once the client is added successfully, a message containing the client's credentials is shown.

Client added successfully. Client Id: u9llqWroUkTCimd3ql5Foca8SGuFfJi

Client Secret: CoBF6jFx4EntK5KG8dsAeyu30s-cYTPzbqLeAljOVm6TqGRHBEcojii4-Y2GR0n1eti17O2UnU


Please copy and keep your client secret safe as you can't fetch your secret once you close this message.



**Warning:** Once the client is added successfully, the Client ID and Client Secret are generated and shown in the message. You must copy and keep your Client Secret safe before closing the message as you cannot fetch your client secret later from anywhere else.

### Move a site from CollabNet Enterprise Edition to CollabNet TeamForge 17.4

As a domain administrator, you will find some aspects of TeamForge similar to CollabNet Enterprise Edition and others quite different. Here are some steps you can take to ease the transition for yourself and your site's users.

 **Tip:** Training is available to help you manage the transition. See the [CollabNet training portal](#) for details.

#### 1. Get a new password for TeamForge 17.4 site


Your login name is migrated from CollabNet Enterprise Edition to TeamForge 17.4, but your password is not.

To create your new password, click **Forgot Your Password?** on the TeamForge 17.4 home page.

Your TeamForge site emails you a link you can use to set a new password.

#### 2. Be aware of password policy differences.

TeamForge 17.4 uses a slightly different password policy from your CollabNet Enterprise Edition site. To review the site's password policy settings, click **View Password Policy** on the **Reset Password** page.

 **Note:** Users' password histories are not migrated. This means that if "Size of user's password history" was set in CollabNet Enterprise Edition, the history will be cleared on migration.

#### 3. Site-wide broadcasts are a good way to keep your site's users up to date during the transition.

See [Communicate with all site users at once](#) on page 335.

#### 4. Each of the CollabNet Enterprise Edition tools has its counterpart in CTF. Project information and resources work a little differently in each one.

For a detailed feature-to-feature comparison, see [Migrating tools from Enterprise Edition to TeamForge 17.4](#) on page 411.

#### 5. For the end user, permissions mostly work the same way on the new site as on the CEE site. However, it may be useful to know about the different ways permissions are handled behind the scenes.

For a detailed comparison of permission handling between the two platforms, see [How do user permissions in CollabNet Enterprise Edition map to permissions in TeamForge 17.4?](#) on page 971

#### 6. Project templates are not migrated from CollabNet Enterprise Edition to CTF.

CollabNet will provide support to manually recreate templates after the site is upgraded.

For information about how project templates work in CTF, see [Create a project template](#) on page 303.

#### 7. TeamForge 17.4 URLs are different from CollabNet Enterprise Edition URLs. Backward compatibility is provided for certain important URLs.

For detailed information about how URLs are handled, see [How do URLs in CollabNet TeamForge 17.4 compare with URLs in CollabNet Enterprise Edition?](#) on page 1009

#### 8. Remote Publishing

Remote publishing gets its own repository now. It is listed as a publishing repository in the SCM listing page. Please check out the publishing repository and use it for all remote publishing activities.

#### 9. Customize Project Home Page

Post migration, please check the **Project Admin > Project Settings** Home Page settings to ensure that the desired content is displayed as the Home page.

If the **Show default project pages and components** option is selected, the project home page displays project pages. You will see the content of `index.html` (if you had one in CollabNet Enterprise Edition) in the project home page as an html component. Moving forward, you will have to edit this html component through the PCE only and you can not edit `index.html` directly in the repository to modify the project home page.

If the **Show custom web page** option is selected, the project home page displays content from the `www/index.html` of the 'Publishing' repository of the project. In this case, any changes to the `www/index.html` file will be visible in the project home page.

You can refer this help page for details: [Create a custom project home page](#) on page 141

#### 10. Change of CVS URL

CVS URLs used in CollabNet Enterprise Edition will work in TeamForge 17.4, but the repository path is different. See the **Source Code** tab of the TeamForge 17.4 project for the new checkout command for your CVS repository. The CVS working copies must be recreated (checked-out afresh).

## 11. Subversion

Existing Subversion working copies can be used as available and do not require a fresh checkout, except for the remote publishing content under the "www" folder.

## 12. Re-do project requests/posts

Unapproved project requests or discussion posts are not migrated to TeamForge 17.4. If you have project requests in CollabNet Enterprise Edition that were not approved until migration, it will have to be done afresh in TeamForge 17.4.

## 13. CollabNet Enterprise Edition's audit history

Audit history from CollabNet Enterprise Edition is not migrated into TeamForge 17.4, but the data is archived. Please get in touch with your CollabNet representative for accessing the data.


### Delegate site administration tasks with site-wide roles

To spread out the site administration burden, consider creating additional site administrators with permissions to do almost everything a full site administrator can do.


To allow some CollabNet TeamForge users to use one or more TeamForge tools across projects, give them a role that allows them to use project tools in any project on the site.

#### Create additional site administrators


To assist in the administration of the TeamForge site, a person must have a site administrator user account with a corresponding role on that site. TeamForge administrators can create suitable site-wide roles and delegate site administration responsibilities.

 **Note:** You can choose site administration permissions through site-wide roles.


1. Go to **My Workspace > Admin**.
2. Click **Roles** from the **Projects** menu.
3. Click **Create**.
4. On the **SITE-WIDE ROLE** tab, write a name and description for the role.  
The role name is case-sensitive.
5. To prevent inheritance of the role into private projects, select the **PREVENT ACCESS** option.


 **Note:** Selecting the option to prevent role inheritance does not affect access to public and gated projects. On selecting **Prevent access**, the user may not be allowed to do project-permissions related tasks in private projects.


6. Click **Create**.  
The restricted site administrator role is created. The **Edit Site-wide Role Permissions** page appears.

 **Note:** You can select the permissions for site administration tools as well as for applications available across all projects.

7. Select the apt site administration and/or project permissions listed on the **ROLE PERMISSIONS** tab, to match the responsibility assigned to a user with that role.

 **Tip:** You may not want to risk delegating the task of deleting projects, users, groups, roles or categories.

 **Note:** If you are creating a site-wide role that has Project Tracker's "Configure - Site" permission, you must also assign the "Role - View" permission.


 **Note:** To manage artifact types globally, users must have project administrator permissions in a site-wide role.

The role is created. You can assign it to site members at any time.

#### Modify the additional site administrators privileges

If restricted site administrators need to do things that are not allowed by a role you have assigned to them, you may need to change the permissions associated with that role.

When you edit a role, all restricted site administrators with that role get the updated permissions automatically.

 **Tip:** You may have prevented the access permission into private projects earlier. Now you can modify the restriction or change other relevant permissions.


1. Go to **My Workspace > Admin**.
2. Click **Roles** from the **Projects** menu.
3. From the list of roles, click the role you want to edit or select the check box and click **Edit**.
4. On the **Edit Site-wide Role Permissions** page, make the changes you need.
  - To edit the title or description of the role, click **Edit**.
  - To edit the site administration and/or project permissions, choose an application from the left side of the page and select or deselect permissions and resources.
  - To edit the site members to whom the role is assigned, click **Assigned Users** tab.
5. Click **Save**.

### **Make selected users additional site administrators**

You can empower site users to assist in site administration by giving them a suitable role.

Depending upon the permissions you grant via site-wide roles, you can select site users who could be granted the privilege.

1. Go to **My Workspace > Admin**.
2. Click **ROLES** from the **Projects** menu.  
The existing site-wide roles are listed.
3. Click the role that you want to assign to the site users.
4. On the **Edit Site-wide Role Permissions** page, click the **ASSIGNED USERS** tab.  
All users who currently have the role are displayed.
5. Click **Add**.
6. In the **Find a User** window, select the site users you want to add, and move them from the **Found Users** list to the **Selected Users** list. Click **Add**.


 **Note:** You can search by full or partial user name or full name to find the desired site members.

7. Click **OK**.


The additional site administrators are now ready to act! Their names are added to the **Assigned Users** list.

### **Give project-independent access for project tools**

To allow some CollabNet TeamForge users to use one or more CollabNet TeamForge tools across several projects, create site-wide roles with specific project permissions, minus site administrative permissions. Assign these site-wide roles to those who may need to access the project tools in any project.

 **Note:** For example, you may want a user to be able to use Tracker across several projects. You don't need to create and assign a role supporting the task individually across all projects. Just do it one time as a site-wide role and assign it to the user.


1. Go to **My Workspace > Admin**.
2. Click **Roles** from the **Projects** menu.
3. Click **Create**.
4. On the **Create Site-wide Role** page, write a name and description for the role.  
The role name is case-sensitive.
5. To prevent inheritance of the role into private projects, select the **PREVENT ACCESS** option.

 **Note:** Selecting the option to prevent role inheritance does not affect access to public and gated projects. On selecting **Prevent access**, the user may not be allowed to do project-permissions related tasks in private projects.

6. Click **Create**.  
The site-wide role is created. The **Edit Site-wide Role Permissions** page appears.

 **Note:** You can select the permissions for applications and resources available across all projects.

7. Select the required project permissions listed on the **ROLE PERMISSIONS** tab, to match the tasks you want the user with that role to perform.

 **Tip:** Select "Tracker-Create" permission if you want the user to be able to create new trackers.

The role is created. You can assign it to site members at any time.


### Provide ready-made roles

You might want to create roles that projects across the site can use with minimum effort and maintenance. Using global project roles is an easy way of enforcing role-based similarities and removing role duplication across projects.

You can suggest to the project administrators to use common global project roles while assigning project tasks in CollabNet TeamForge, instead of creating and managing several similar roles for their individual projects.


### Create a global project role

To help project managers get their project members set up quickly, provide ready-made project roles that any project on your site can use.

 **Note:** You need the "Role-Create" permission to create global project roles. All site administrators and some restricted site administrators have this permission.


1. Go to **My Workspace > Admin**.
2. Click **Roles** from the **Projects** menu.
3. Click the **GLOBAL PROJECT ROLES** tab.

All the existing Global Project Roles are listed here. It is a good idea to check this list before you create another role.

 **Note:** You can suggest that the project administrators check the **Project Admin > Permissions > Roles > View: Global Project Roles** list before creating any new roles in their projects.

4. Click **Create**.
5. On the **Create Global Project Role** page, write a name and description for the role.


The role name is case-sensitive.

 **Tip:** Remember that the role name can not be the same as a site-wide role name.

6. To allow inheritance of the role's permissions into private sub-projects, clear the **PREVENT INHERITANCE** option.
7. To allow the project members to be able to request this role, select **PROJECT MEMBERS CAN REQUEST THIS ROLE**.

Project members can submit requests for Available upon Request roles. For a project, the project administrators can set an Available upon Request role to be automatically granted to the project member requesting it.


8. Click **Create**.  
The new global project role is created. The **Edit global project role permissions** page appears.
9. Select the application permissions that are relevant to the role, from those listed on the **ROLE PERMISSIONS** tab.

 **Tip:** You may want to restrain providing project or application administration permissions, until required.

The role is created. The project administrators can assign it to their project members any time.

### Modify a global project role


You may need to add or remove certain permissions from an existing global project role to assign new tasks or change the access permissions given via the role.

 **Note:** Only site administrators or restricted site administrators with "Role-Edit" permission can edit global project roles.

1. Go to **My Workspace > Admin**.



2. Click **Roles** from the **Projects** menu.
3. Click the **GLOBAL PROJECT ROLES** tab.  
All the existing global project roles are listed here.
4. Select the role that you want to edit and click **Edit** or just click the hyperlinked role name.  
The **Edit global project role permissions** page appears.
5. Click **Edit** to make changes to the role details.
6. Modify the **Role Name** or **Description**, if required.  
The role name is case-sensitive and must not be the same as a site-wide role.
7. Change the inheritance setting to prevent or allow inheritance of the role's permissions into private sub-projects.
8. To make the role requestable or non-requestable, change the **Project members can request this role** setting.  
Project members can submit requests for Available upon Request roles. For a project, the project administrators can set a Available upon Request role to be automatically granted to the project member requesting it.
9. Click **Update**.  
The global project role is modified.
10. Select the application permissions that are relevant to the role, from those listed on the **ROLE PERMISSIONS** tab and click **Save**.


 **Tip:** You may want to restrain removing project or application administration permissions as the change impacts existing users too.

The role is modified.

### Delete a global project role

If you no longer need a global project role, you should delete it.

On deleting a global project role, all the user associations in the projects are removed.

 **Note:** Only site administrators or restricted site administrators with "Role-Delete" permission can delete global project roles.

1. Go to **My Workspace > Admin**.
2. Click **Roles** from the **Projects** menu.
3. Click the **GLOBAL PROJECT ROLES** tab.  
All the existing global project roles are listed here.
4. Select the role that you want to delete and click **Delete**.  
A confirmation message appears.
5. Click **OK** to continue with deleting the selected role.

The selected global project role is deleted and all its associations are removed.


### Manage projects as a group

To manage two or more independent projects, create a project group in CollabNet TeamForge 17.4. Similar to projects, a project group provides the platform for sharing project members, roles and permissions across a group of projects. With some well planned settings, you can manage several of your projects and also effectively control the project members accessing each project.

To begin controlling multiple projects at one go, create a project group.

#### Create a project group

To begin managing your several projects together, create a project group in TeamForge 17.4.

 **Note:** When you create a project group, you are granted administration rights for the group and will be able to perform actions such as adding or removing a project from the group. You will be listed as an administrator in the group's **Project Group Details**. (Site administrators are exempt from this, since they can perform all actions on the site without requiring specific roles).


1. Go to **My Workspace > Admin**.
2. Click **PROJECT GROUPS** from the **Projects** menu.

The existing project groups are listed here.

3. To create a new project group, click **Create Project Group**.
4. On the **Create Project Group** page, write a name and description for the project group.  
The project group name must be unique, however, it can be the same as any of the projects.


The project group is created.

In your **Project Group** page, you can add projects, add project group members or specify roles for the project group.

-  **Tip:** You can always come back to this page later to specify projects, users or roles that affect your project group.

### Edit a project group

You can make changes to a project group to keep it updated with the various projects you may be managing with that in CollabNet TeamForge C17.4. You must have the administrator permissions for project groups to make any modifications.

-  **Tip:** You could be either a site administrator or a project administrator, but you must have the project groups administration permissions to manage projects as a group.

1. Go to **My Workspace > Admin**.
2. Click **PROJECT GROUPS** from the **Projects** menu.

The existing project groups are listed here.

3. Click the project group that you want to modify. The **Project Group Details** page appears.
4. Click **Edit** and make the changes as required.

You can update the project group name, description as well as the administrators.

5. To add one or more administrators for the project group, click the Search icon next to **Administrators**.
  - a) On the **Find a User** page, select the required administrators, click **Add** and click **OK**.


The project group administrators are listed on the **Project Group Details** page.

The project group is modified.

From the **Project Group Details** page, you can add the projects, add project group members or specify the desired roles for the project group.

### Delete a project group

If the projects being managed under a project group have achieved their targets you may not need the project group anymore. You can delete a project group in CollabNet TeamForge C17.4. You must have the administrator permissions for project groups.

-  **Tip:** You could be either a site administrator or a project administrator, but you must have the project groups administration permissions to manage projects as a group.


1. Go to **My Workspace > Admin**.
2. Click **PROJECT GROUPS** from the **Projects** menu.

The existing project groups are listed here.

3. To delete a project group, select the project group from the list and click **Delete**.


You may get a warning message as the projects being managed under the group may have active role assignments. Click **OK** on the message to proceed with deleting the project group.

The project group is deleted.

-  **Tip:** The relationship between projects and the project group is broken and you can no longer manage projects using the deleted project group.

### Add projects to a project group

You may be interested to start using the project group that brings your projects together. Add your projects to your project group as the initial step.

-  **Tip:** You could be either a site administrator or a project administrator, but you must have the project groups administration permissions to manage projects as a group.


If you have just created your project group, you might already be on the **Project Group Details** page. Skip the first three steps in that case.

1. Go to **My Workspace > Admin**.
2. Click **PROJECT GROUPS** from the **Projects** menu.

The existing project groups are listed here.

3. Click your project group. The **Project Group Details** page appears.
4. On the **Project Group Details** page, click **Add** to associate your projects with the project group.


A list of projects appears.

-  **Tip:** You can only add the projects for which you are the project administrator.

5. From the **PROJECT** list, select the projects and click **Add**.


The selected projects are added to the project group.

If you add an irrelevant project to the group, you can select it from the **Project Group Details** page and click **Remove** to remove it from the project group.

-  **Tip:** You can always come back to this page later to add more projects, specify users or roles that affect your project group.

### Manage user membership for a project group


Your project group needs to be set up with project group members to facilitate any administrative tasks that you may want to do.

-  **Tip:** You could be either a site administrator or a project administrator, but you must have the project groups administration permissions to manage projects as a group.


If you have just created your project group, you might already be on the **Project Group Details** page. Skip the first three steps in that case.

1. Go to **My Workspace > Admin**.
  2. Click **PROJECT GROUPS** from the **Projects** menu.
- The existing project groups are listed here.
3. Click your project group. The **Project Group Details** page appears.
  4. From the left navigation pane, click the **User Membership** link to add users to the project group.
  5. On the **PROJECT GROUP MEMBERSHIP** tab, click **Add**.
  6. On the **Add User** page, find the users you want by one of these methods:

- Under **Search for Users**, filter the list of site users eligible to join this project group. You can filter by full or partial name or user name.


-  **Note:** Search text is not case-sensitive.

- Browse the list of registered users on the site. Sort them by name, user name, email address or membership status.

-  **Note:** If a site has a great many users, you must filter them first to narrow down the list. This helps avoid slowing down the system.

7. Select the users you want to add.
8. Under **Assign Roles (Optional)**, select the roles you want the users to have.

You can select any available global project role or role created just for this project.

-  **Tip:** If you prefer, you can skip this step and assign roles later on.

9. Save your changes.

- Click **Save** to return to the **Project Group Membership** page.
- Click **Save and Add More** to keep adding users.

The selected users are granted membership to the project group.

If you add a user who may not need to be a member of the project group, you can select the user from the **Project Group Membership** page and click **Remove** to remove the user's membership.

### Manage users' access to project groups

As a project group administrator, you know that the key to any user's access to the project group and the projects belonging to the group, is in the role you assign to the user. You can do almost all the role-related activities with project groups that you could do with individual projects.

You can create site-wide roles and assign the users to those roles and then bring them into your project group. You can create roles specifically for your project group and assign users to those or use global project roles, as required. The permissions inherited through the assigned roles are subject to the "Prevent Inheritance" option setting done while creating the roles.

#### *Give a role to project group members*


A role can be assigned to many project group users at once.

While the user-role matrix provides a convenient way to add project group members to a role, it can become unwieldy if the project has a large number of users or roles. When that is the case, try assigning roles to multiple project group users at one time.

1. Go to **My Workspace > Admin**.
2. Click **PROJECT GROUPS** from the **Projects** menu.

The existing project groups are listed here.

3. Click your project group. The **Project Group Details** page appears.
4. From the left navigation pane, click the **Permissions** link to specify user roles applicable to the project group.
5. On the **Roles** tab, click the **Global Project Roles** or **Roles Created For a Project** option of **View**, to display existing global roles or roles created just for this project.
6. Click the name of the role that you want to assign to project group members.
7. On the **Edit Role** page, click the **Assigned Users** tab.  
The **Assigned Users** page shows all users who currently have the role.
8. Click **Add**.
9. In the **Find a User** window, select the project group members you want to add, and move them from the **Found Users** list to the **Selected Users** list.
  - Click **Add** to move selected users.
  - Click **Add All** to move all users.

 **Note:** You can search by full or partial user name or full name to find the desired project group members.

10. Click **OK**.

The project group members are now assigned the role.

#### *Give roles to a project group member*

A project group member can have any number of roles. As project group administrator, you must assign each project group member's roles with care. The roles would impact not just an individual project, but would also grant same permissions across the projects in a project group.

Permissions are cumulative. The project group member has all of the access permissions allowed by all of the assigned roles, plus any permissions that may have been assigned globally using application permissions.

1. Go to **My Workspace > Admin**.
2. Click **PROJECT GROUPS** from the **Projects** menu.

The existing project groups are listed here.

3. Click your project group. The **Project Group Details** page appears.
4. From the left navigation pane, click the **Permissions** link to specify user roles applicable to the project group.
5. Click the **User-Role Matrix** tab.

Observe the users listed on the left and all the available roles (global and direct) on the right. Users can be assigned global roles and roles created just for this project.

6. Select roles for each project member.
7. Click **Save**.

The roles are now assigned to each project group member.

#### *Assign roles in multiple projects to a user group*

Project managers can assign a role to multiple users at once by assigning the role to a user group that contains all those users. As a site administrator, you can do the same thing across multiple projects, by treating the projects as part of a project group.

1. Go to **My Workspace > Admin**.
2. Click **PROJECT GROUPS** from the **Projects** menu.

The existing project groups are listed here.

3. On the **Project Group** page, click **Permissions**.
4. On the **User Group-Role Matrix** tab, add the user group you want, then select the roles you want to assign to that user group.
5. Click **Finish** or **Finish and Add More**.



**Restriction:** When you give a group access to a CVS or Wandisco Subversion repository, members of the group can view the repository but cannot do repository actions, such as commit and update. You must assign those permissions to users individually.

#### *Assign user groups to a role*

To manage permissions for a lot of groups or roles at once, try assigning user groups to roles.

1. Go to **My Workspace > Admin**.
2. Click **PROJECT GROUPS** from the **Projects** menu.

The existing project groups are listed here.

3. Click your project group. The **Project Group Details** page appears.
4. From the left navigation pane, click the **Permissions** link to specify user roles applicable to the project group.
5. On the **Roles** tab, to display existing global or direct roles, click the **global project roles** or **Direct Roles** option of **View**.

You can assign the project group users or user groups to a global project role or a role created just for this project.

6. Click the name of the role that you want to assign to the user group members in the project group.
7. On the **Edit Role** page, click the **Assigned User Groups** tab.  
The **Assigned User Groups** table shows all user groups who currently have the role.
8. Click **Add**.
9. Type some of the group's name in the **Name (search)** box and click **Find**.
10. From the **Add User Group to Role** table, select the user groups that you want to add, and click **Finish**.




**Restriction:** When you give a group access to a CVS or Wandisco Subversion repository, members of the group can view the repository but cannot do repository actions, such as commit and update. You must assign those permissions to users individually.

The user groups are now assigned the role.

#### *Create a role in a project group*

A role defines the applications that project group members with that role can use, and the specific things project group members can do in each application.

Any project groups administrator can create and assign a role. It is a good idea to check the existing global project roles before creating any new role for a project group.

 **Note:** Any existing site-wide or global project role can be associated with a project group. It is advisable to check the permissions granted via a role before assigning it to users or user groups in a project group as it would impact more than a single project.


1. Go to **My Workspace > Admin**.
2. Click **PROJECT GROUPS** from the **Projects** menu.


The existing project groups are listed here.

3. Click your project group. The **Project Group Details** page appears.
4. From the left navigation pane, click the **Permissions** link to specify user roles applicable to the project group.
5. On the **Roles** tab, click **View: Roles Created For this Project Group**.

You can view global project roles by selecting **View: Global Project Roles** before creating a new role for this project.

6. Click **Create**.
7. On the **Create Role** page, write a name and description for the role.
8. To allow the inheritance of the role into private subprojects, de-select the **Prevent Inheritance** option.


 **Note:** By default, the role inheritance into private subprojects is prevented. For example, you may not want administrator roles to be inherited in subprojects, until required.

 **Note:** Selecting the option to prevent role inheritance does not affect public and gated projects.

9. Click **Create**.

The role is created. The **Edit Role** page appears.

10. For each application listed on the **Role Permissions** page, select the permissions and resources you want to make available to users with this role.

 **Note:** You can specify access to individual top-level folders, but not to specific subfolders.

11. Click **Save**.

The role is created. You can assign it to project group members or user groups associated with the project group at any time.

#### *Edit a role in a project group*

As a project group administrator, you may need to update the permissions granted by a role being used across the projects.

While modifying a role, it is better to be cautious about granting more access than required by the users or a user group. In the case of project groups, as a role could be mapped across projects, consider being restrictive.

1. Go to **My Workspace > Admin**.
2. Click your project group. The **Project Group Details** page appears.
3. From the left navigation pane, click the **Permissions** link to specify user roles applicable to the project group.
4. On the **Roles** tab, click **View: Roles Created For a Project**. The existing roles created for this project are listed.

You can view global project roles by selecting **View: Global Project Roles**, if required.

5. Select the role that you want to modify and click **Edit**.
6. On the **Edit Role** page, make the desired changes.

You can update the role name, description and inheritance settings on clicking **Edit**.

The role's permissions, assigned users or assigned user groups can also be modified as required.


7. Click **Save** after making the changes.

The role in the project group is updated.

### *Delete a role in a project group*

When you no longer need a role that you created for a project, it's a good idea to delete that role.

Any project groups administrator can create and assign a role. It is good to keep the project role tray as small and as manageable as possible.

 **Note:** Any existing site-wide or global project role can be associated with a project group. It is advisable to check the permissions granted via a role before assigning it to users or user groups in a project group as it would impact more than a single project.

1. Go to **My Workspace > Admin**.
2. Click **PROJECT GROUPS** from the **Projects** menu.

The existing project groups are listed here.

3. Click your project group. The **Project Group Details** page appears.
4. From the left navigation pane, click the **Permissions** link to specify user roles applicable to the project group.
5. On the **Roles** tab, click **View: Roles Created For a Project**.

You can view global project roles by selecting **View: Global Project Roles**. However, you can only delete roles created for this project here.

6. Select the role and click **Delete**.

You may get a warning message if the role you are trying to delete is in use. Click **OK** to proceed with deleting the non-required role.

The selected role is deleted.

### **Manage users**

You can help projects succeed by creating and managing user accounts.

#### **Communicate with all site users at once**

To say something to everyone who uses your site, post a site-wide broadcast message.

For example, if you plan to upgrade your site, you may want to let users know a few days before your upgrade that the site will be unavailable for a short time.

The broadcast message can be viewed even without logging into TeamForge


1. Go to **My Workspace > Admin**.
2. Click **SYSTEM TOOLS** from the **Projects** menu.
3. Click **Broadcast Message**.
4. In the **Broadcast Message** page, type the message you want to broadcast and click **Broadcast**.

The message is displayed at the top on all the pages of the site until it is deleted or replaced.


To edit a broadcast message quickly, you do not need to go to the specific broadcast message page. You can edit it from any page, using the Edit icon that appears at the far right end of the broadcast message box.

### Create a new user account

To participate in a TeamForge site, a person must have a user account on that site. TeamForge administrators can create these user accounts. This topic applies to sites with no LDAP authentication.


 **Note:** If your TeamForge site uses LDAP authentication, TeamForge administrators cannot create new user accounts. On a site with LDAP authentication, each user must log into TeamForge using his or her LDAP user name and password.

1. Go to **My Workspace > Admin**.
2. Click **USERS** from the **Projects** menu.
3. Click the drop-down arrow next to **Create** and click **Single User**.  
You can select **Multiple Users** to create several user accounts at one time.
4. On the **Create User** page, enter the field values appropriately.
  - a) Enter a user name for the user.


 **Note:** Your user name must meet these criteria:

- User name is case-sensitive. However, to make usernames case-insensitive set the site-options token `ALLOW_CASE_INSENSITIVE_LOGIN` to True
- Minimum number of characters as specified in the `site-options.conf` file.
- No spaces.
- Should have at least one letter.
- The first character is a letter.

- b) Enter and confirm a password for the user, if you prefer to set the user's password yourself.

 **Tip:** To invite users to create their own password, leave the **PASSWORD** field blank. A *password ticket email* will be sent to users to let them create a password.


- c) Enter the **FULL NAME** and **EMAIL ADDRESS** of the user.

 **Tip:** You can add more email addresses for the user after you finish creating their profile.

- d) Enter the user's organization.


Organization can be a geographic designation, a corporate division, or whatever you want. It's advised to keep it consistent across your site.

- e) Select the language from the **LOCALE** drop-down list.

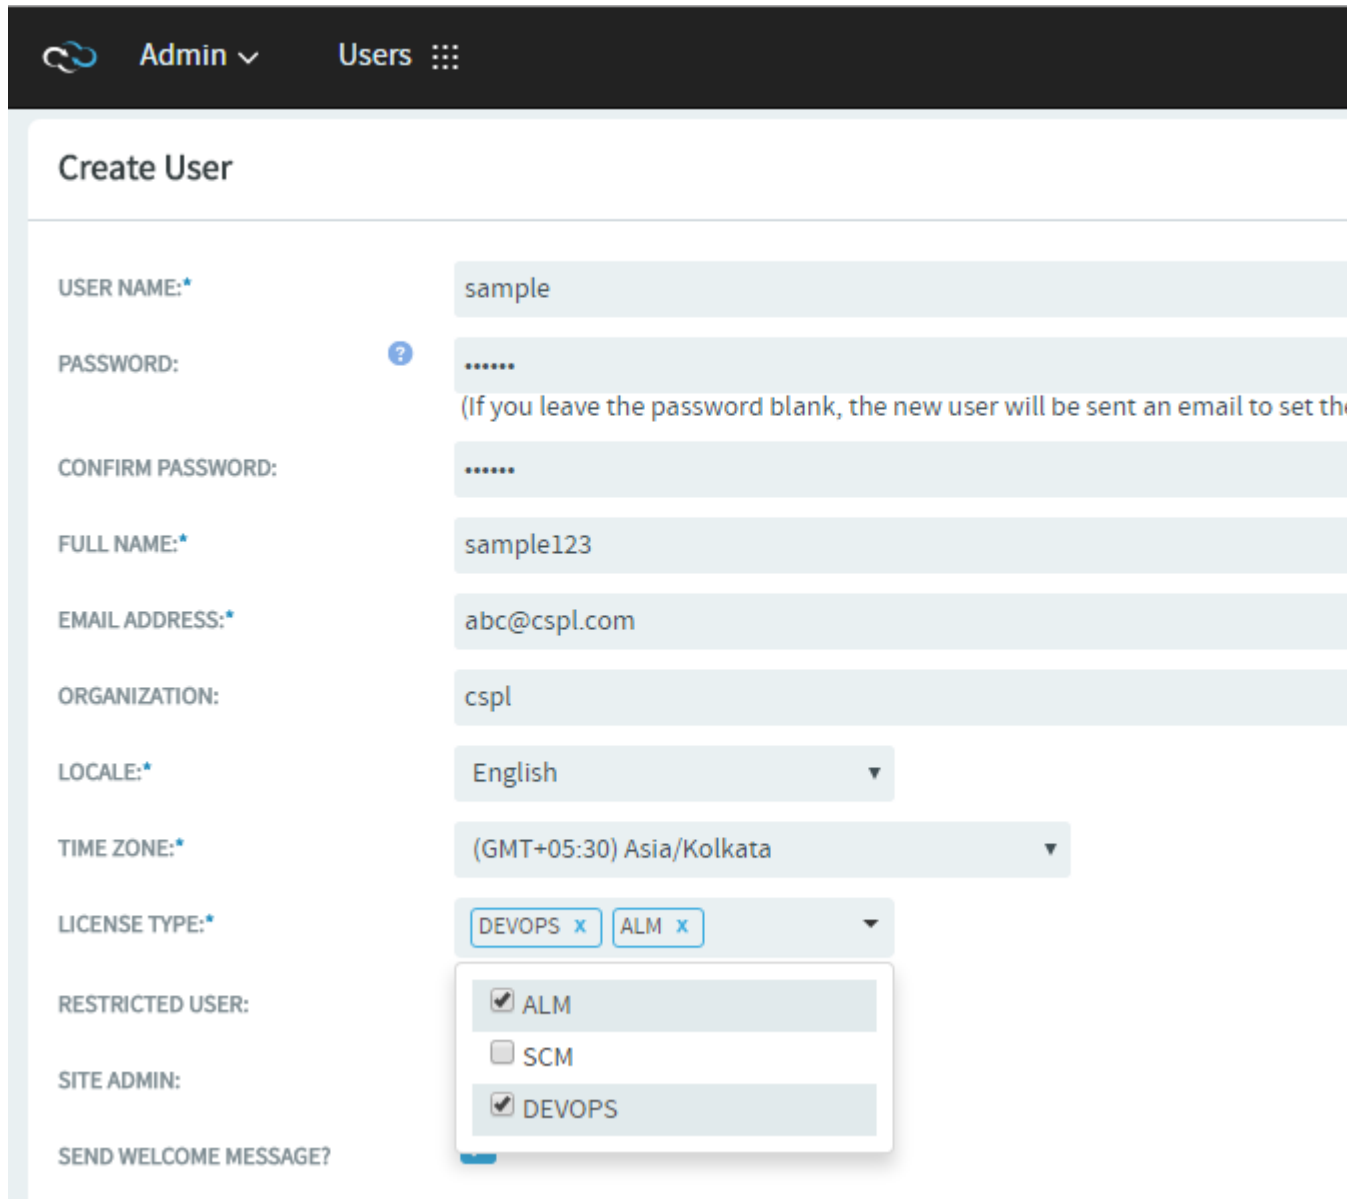
 **Note:** TeamForge supports English, Chinese, Japanese and Korean languages.

- f) From the **TIME ZONE** drop-down list, select the preferred time zone for the user.



 **Important:** Selecting the time zone overrides the default time zone set by the site-options token, *DISPLAY\_TIMEZONE*. It reflects in all the email notifications and TeamForge pages excluding integrated application pages.

- g) Choose the user's TeamForge **LICENSE TYPE**. Multi select option is now enabled. Users can now use combination of license types such as ALM and DevOps.



The screenshot shows the 'Create User' form with the following fields and values:

- USER NAME: sample
- PASSWORD: (masked with dots, with a help icon and a note: "(If you leave the password blank, the new user will be sent an email to set the password)")
- CONFIRM PASSWORD: (masked with dots)
- FULL NAME: sample123
- EMAIL ADDRESS: abc@cspl.com
- ORGANIZATION: cspl
- LOCALE: English
- TIME ZONE: (GMT+05:30) Asia/Kolkata
- LICENSE TYPE: DEVOPS x ALM x
- RESTRICTED USER: (dropdown menu open showing checkboxes for ALM, SCM, and DEVOPS)
- SITE ADMIN: (checkbox)
- SEND WELCOME MESSAGE? (checkbox)

- h) Choose a user type. You can choose only one user type for each user.

- SITE ADMIN: Administrators have unlimited access to all the data in TeamForge.
- RESTRICTED USER: Restricted users can only access projects of which they are members.

 **Important:** If you do not select **RESTRICTED USER**, the user will be unrestricted and able to access all projects that have not been made private by a project administrator.


- i) To send a welcome message to the user, select **SEND WELCOME MESSAGE?**.

##### 5. Click **Create**.

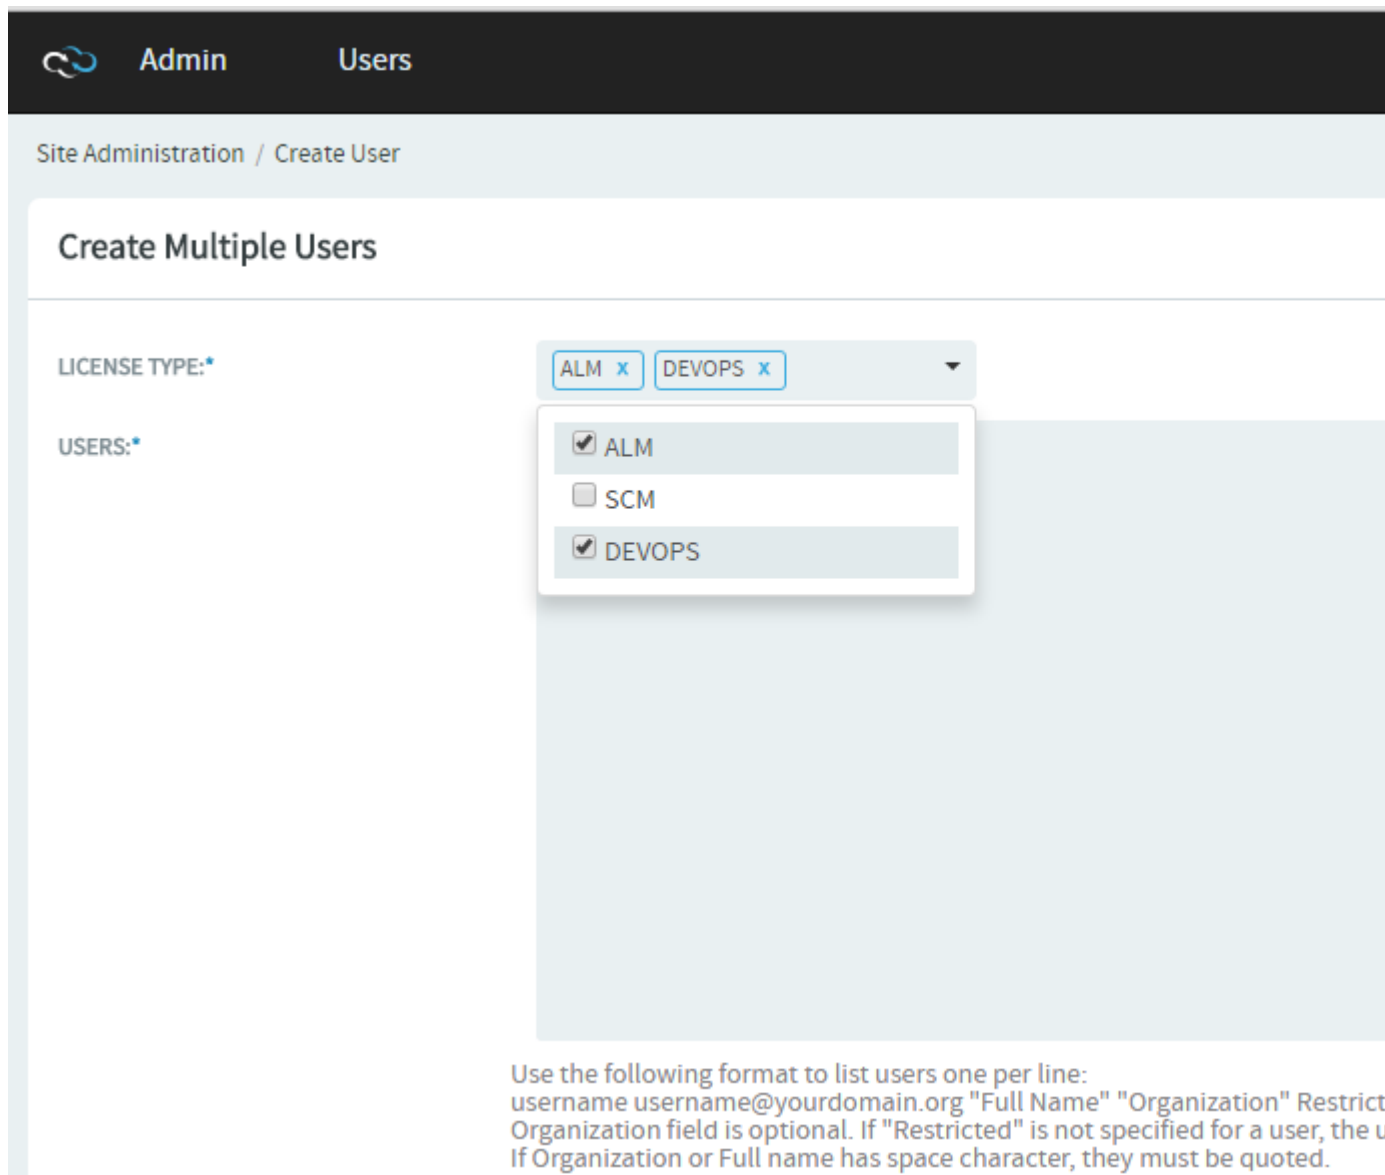
The user account is created.

## Create multiple user accounts

To participate in a TeamForge site, a person must have a user account on that site. TeamForge administrators can provide access to multiple users by creating their accounts together.

 **Note:** If your TeamForge site uses LDAP authentication, TeamForge administrators cannot create new user accounts. On a site with LDAP authentication, each user must log into TeamForge using his or her LDAP user name and password.

1. Go to **My Workspace > Admin**.
2. Click **USERS** from the **Projects** menu.
3. Click the drop-down arrow next to **Create** and click **Multiple Users**.
4. Choose the user's TeamForge **LICENSE TYPE** on **Create Multiple Users** page. .Multi select option is now enabled .Users can now use combination of license types such as ALM and DevOps.



Site Administration / Create User

### Create Multiple Users

LICENSE TYPE:\*

USERS:\*

ALM x DEVOPS x

ALM


SCM

DEVOPS

Use the following format to list users one per line:  
 username username@yourdomain.org "Full Name" "Organization" Restricted  
 Organization field is optional. If "Restricted" is not specified for a user, the u...  
 If Organization or Full name has space character, they must be quoted.

5. On the **Create Multiple Users** page, enter up to 25 lines like this, one user per line:

```
username username@yourdomain.org name organization Restricted
```

 **Note:** Usernames must meet these criteria:

- 1 to 31 characters.

- Only alphanumeric characters.
- No spaces.
- At least one letter.
- The first character is a letter.

**Tip:**

- Organization field is optional.
  - To create an unrestricted user, omit "Restricted." Restricted users can only access projects of which they are members, while unrestricted users can access all projects that have not been made private by a project administrator.
  - Use quotes around the full name or the organization information if it is more than a single word.
  - A maximum of twenty-five user accounts can be created at one time.
6. Choose the user's TeamForge **LICENSE TYPE**. Multi select option is now enabled. Users can now use combination of license types such as ALM and DevOps.

The screenshot shows the 'Create Multiple Users' form in the TeamForge Admin interface. The 'LICENSE TYPE' field is a multi-select dropdown with 'ALM' and 'DEVOPS' selected. The 'USERS' field is empty. A dropdown menu is open showing 'ALM', 'SCM', and 'DEVOPS' with checkboxes.

Use the following format to list users one per line:  
 username username@yourdomain.org "Full Name" "Organization" Restricted  
 Organization field is optional. If "Restricted" is not specified for a user, the u  
 If Organization or Full name has space character, they must be quoted.

7. Click **Create**.

The user accounts are created and password e-mails are sent to all the new users.

## Create a user group

To manage multiple users at once, create a group that represents them.

1. Go to **My Workspace > Admin**.
2. Select **USER GROUPS** from the **Projects** menu.
3. Click **Create** and provide a name for the group and a description of its purpose.



**Note:** If your project is a child of another project, it may have inherited one or more user groups from its parent project. To work with inherited users and user groups, you must go to the project that they belong to.

4. Click **Create**.

## Add a user to a user group

Put together multiple users who share characteristics in a user group.

1. Go to **My Workspace > Admin**.
2. Select **User Groups** from the **Projects** menu.



**Note:** In TeamForge 17.4, **Groups** (in earlier versions of the product) has been renamed to **User Groups** to better distinguish user groups from project groups.

3. Under **User Groups**, click the group to which you want to add the user.
4. On the **Edit Group** page, click **Add**.
5. Use the picker to move users into the group, and click **OK**.  
You can select the inherited project members also from the list.
6. Click **Return**.

## Find a user

To find a user, filter the list all CollabNet TeamForge users on your site.

1. Go to **My Workspace > Admin**.
2. Click **USERS** from the **Projects** menu.
3. Specify the filter criteria in one or more filter fields (at the top of each column) and click **FILTER**.
  - You can find a filter field at the top of each column in most of the tables in the TeamForge application.
  - The filter field could be a text box or a drop-down list with multi-select check boxes.


Stories Page 1 of 40 (600 Items) FILTER  COLUM

PRIORITY	ARTIFACT ID: TITLE	ASSIGNED TO	TEAM	SUBMITTED BY
3 MEDIUM	artf284437 : Create UI for viewing all diffs for entire review on one screen	C. Michael Pilato	SCM	Mark Phippard
3 MEDIUM	artf277197 : Workspace widget for artifacts	Baybora Aksoy	SCM	Mark Phippard
4 LOW	artf246564 : Reports - write protractor tests to cover e2e tests.	None	Reports (Datamart)	Muhammed Ashik
2 HIGH	artf287743 : Include 'Artifact Tags' in Project Templates	Ramkumar Rathinasamy Viveganandan	TF Core > Avengers - Tags	Aishwarya Rajan
2 HIGH	artf287742 : Tags in Edit Artifact Page	Ramkumar Rathinasamy Viveganandan	TF Core > Avengers - Tags	Aishwarya Rajan
2 HIGH	artf283822 : Synchronize SVN Indexer execution	akil roy	Devops/Releng	Johannes Passing
2 HIGH	artf278265 : Cleanup the old EventService	Nithya Pattabiraman	TF Core > Jaguars - Architecture	Sriram
3 MEDIUM	artf288250 : EventQ meta RPM	Sinhuja Harikrishnan	Devops/Releng	Sinhuja Harikrishnan
3 MEDIUM	artf287737 : Update the default/trial license to include DevOps (and perhaps new license types)	Gowthaman Ravindran	Mavericks - CRUS	Mark Phippard
3 MEDIUM	artf287193 : SVN load balancing	akil roy	Devops/Releng	Johannes Passing
3 MEDIUM	artf286210 : Monit Snapshot before restart	h.karthik	Devops/Releng	Johannes Passing

1 2 3 4 5 6 7 8 9 10 >

Monitor Import Export Cut Edit Edit Inline Mass Update

- You can type your filter criteria in the text boxes. The search text is case-insensitive.
- You can also select the filter values from one or more drop-down lists. By default, you can only select up to 10 filter values in a drop-down list. However, you can set a value that suits your requirement for the `FILTER_DROPDOWN_MAX_SELECTION` token in the `site-options.conf` file to increase or decrease the count.
- Filter-as-you-type:** You can find the **Enter keywords** text box in all filter drop-down lists. As you type your filter keyword, instant search results are shown in the drop-down list. For example, in the following illustration, typing "R" instantly shows all statuses having the alphabet "R". The search text is case-insensitive.



- Some search filters may not appear if your site administrator has not enabled them.

4. After filtering, if you want to clear the filters, click **FILTER** and select **Clear** from the drop-down list.

All users meeting your filter criteria are displayed.

### Find a user's groups



You can get a consolidated view of all user groups of which a user is a member.

- Go to **My Workspace > Admin**.
- Click **USERS** from the **Projects** menu.
- On the **Users Details** page, click the **USER GROUP MEMBERSHIP** tab.

The groups listed on this tab are the groups that this user is a member of.

### Edit a user account

When a user has trouble accessing the site, you may need to reset the user's password or change the user's account status.

-  **Note:** If your TeamForge site uses LDAP for single-sign-on, passwords must be reset in the LDAP system, not on the Web administration pages. Ask your system administrator for help.
-  **Important:** To avoid disasters, TeamForge makes it impossible to delete or deactivate the TeamForge admin account. You also can't remove the TeamForge admin flag or mark the admin user as a restricted user.


1. Go to **My Workspace > Admin**.
2. Click **USERS** from the **Projects** menu.
3. On the **USERS** tab, click the name of the user whose account you want to edit.
4. On the **User Details** page, click **Edit**.
5. On the **Edit User Information** page, make your changes and click **Update**.


You can specify up to a maximum of three alternate email addresses, if required.

### Act on multiple user accounts at once

A TeamForge administrator can edit the status of multiple user accounts simultaneously.

For example, if you have multiple pending new accounts to approve, you can approve them in a batch instead of individually editing each account.

-  **Note:** A pending user is a user who has requested an account but has not yet confirmed his or her email addresses.

-  **Important:** In the case of TeamForge admin accounts, you cannot make any of these edits:

- Delete the account.
- Change the account status to anything but active.
- Remove the TeamForge admin flag.
- Mark it as a restricted user.

1. Go to **My Workspace > Admin**.
2. Click **USERS** from the **Projects** menu.
3. On the **Users** page, select the users whose status you want to edit.
4. Click the desired status change.
  - **Delete** - Deleted users are removed from all projects. All assigned items are removed from the user. Deleted users do not count against your TeamForge license count.
  - **Disable** - Disabled users cannot log in to TeamForge and do not receive notification messages, but they remain members of projects and selection lists.
  - **Activate** - Active users have full use of TeamForge, subject to RBAC permissions.

### Reset the admin account password

If your TeamForge installation authenticates against an LDAP directory, follow these instructions to reset your admin account password.

If your installation does not validate against LDAP, click **Forgot Your Password** on the TeamForge home page to reset the password for the admin account.

1. With a web browser, go to the URL `http://<host>sf/sfmain/do/forgotAdminPassword`.
2. On the **Admin Account Password Retrieval** page, Click **Send Email**.  
TeamForge sends an email to the address specified for the admin user.
3. Check your email and click the link provided to reset your password.
4. On the **Reset Password** page, enter and confirm a new password.
5. Click **Reset Password**.

You can now log into CollabNet TeamForge with your new password.

### License information

You can obtain a summary of the license information from the **License Info** page.

The **License Info** page provides you with all the basic information about the licenses you purchased for your TeamForge site. This includes details such as the number of TeamForge licenses you had obtained, how many you have used, expiration date and so on.

1. Go to **My Workspace > Admin**.
2. Select **LICENSE INFO** from the **Projects** menu.

### View all roles assigned to a user

A user can have multiple roles in different projects either by being directly assigned those roles or by inheriting them. You might find it useful to see all the roles assigned to a user in a TeamForge site before adding or removing a role.


1. Go to **My Workspace > Admin**.
2. Click **USERS** from the **Projects** menu.
3. In the **ROLES** tab, select a role type in the **View** drop-down - Roles created for a Project, Roles Inherited From Parent Project or Site-wide Roles.
  - Roles created for this project include the roles the user is directly assigned or assigned through a user group in projects and project groups.
  - Inherited roles include the roles the user inherits from parent projects and project groups.

### Support projects

CollabNet TeamForge administrators can do a variety of things to help projects on the site be successful.


#### Create a new project

TeamForge administrators can create new projects without having to submit them for approval.

 **Note:** When a TeamForge administrator creates a new project, he or she is not made a member of the project, and the Founder Project Admin role is not created. To designate a project administrator, you must add the user to the project, then create and assign a project administrator role manually.

1. Go to **My Workspace > Admin**.
2. In the list of TeamForge projects, click **Create Project**.
3. On the **Create Project** page, provide a name for the project.  
This is the name that will appear in all project lists and on the project home page.
4. Enter a URL name for the project, if appropriate.  
This is the name that will appear in the project's URL.
5. Write a description of the project.
6. Select a project template.

A project template is used to pre-populate new projects with the structure and configuration of an existing project. If you do not want to use a project template, choose **None**.

 **Note:** If you create a project from a template that contains an integrated application, you may have to provide some information specific to the integrated application. For example, for Project Tracker you must set a new artifact prefix that is different from the prefix in the template.

7. Click **Create**.

The project is created.

#### Approve a new project

Any registered TeamForge user can request a new project. A new project is activated only after a site administrator approves it.

Before approving a new project, you have the option to review the project details.

1. Go to **My Workspace > Admin**.
2. Click the **PENDING PROJECTS** tab.
3. Select the projects you want to approve from the pending projects list.
  - Click **Approve** to approve the project and move it to the **All Projects** page with status **Active**.
  - Click **Reject** to reject the project and remove it from the list.

The project requester receives an email notification when the project is approved or rejected. If you entered a comment, that also appears in the email notification.

### Rename a project

As the focus of a project shifts, its name or description can become obsolete. A site administrator can update the name or description to help keep the project current.

1. Go to **My Workspace > Admin**.
2. In the TeamForge project list, click the project you want to edit.
3. Click **PROJECT ADMIN** from the **Project Home** menu.
4. On the **Project Admin Menu**, click **Project Settings** page, make the changes you need.
5. Click **Save**.

### Delete a project

If you no longer need a project or any of the data in it, you should delete it.

Deleting a project deletes all of the data within it, with the exception of source code data, which is maintained separately from the rest of the site's content.



**Caution:** Delete a project only if you are sure that you no longer need any of the data within it. Move any items that you want to save.



#### Note:

- You can delete a parent project only when its members, user groups and roles are not in use in any other project.
- When you delete a parent project, the direct subprojects are moved one level up; that is, under the immediate parent of the deleted project.
- If the deleted project has no parent, its subprojects become root-level projects.

1. Go to **My Workspace > Admin**.
2. From the TeamForge project list, choose the project that you want to delete and click **Delete**.

The project is deleted.

### Lock or unlock a project


To ensure that no changes occur in a project while you are collating or migrating project data, lock the project. You must have project administration permissions or be a site administrator to lock or unlock a project.

To lock or unlock a project in TeamForge, go to **Project Settings** and lock/unlock the project.



**Note:** A locked project does not allow any member (including project administrators and site administrators) to make any changes to the project. Besides that, a locked project can not be set as the parent project for any other project and tasks like adding, editing or deleting integrated applications are also not allowed.

Click **PROJECT ADMIN** from the **Project Home** menu.

The project is locked or unlocked as desired. The lock icon  **Locked** appears on all the project pages while the project is locked.



**Note:** If a locked project has an integrated application, for example, project tracker, all the project tracker pages are also non-editable while the project is locked. The user who has access permissions for the integrated application can only view the pages.

### Update a project template

To revise or correct an existing project template, overwrite it with a template of the same name.

You must be a site administrator to overwrite an existing project template.


1. Revise the project that will serve as the basis for the new project template.



**Tip:** You can disable a project template while you make changes to the project. To disable a project template, use the **My Workspace > Projects > Templates** page.



2. Select **PROJECT ADMIN** from the **Project Home** menu.
3. On the **Project Settings** page, click **Create Project Template**.
4. Select **Replace Existing Template** and choose the template you want to overwrite.
5. Provide the description for the template.  
(If you want to change the template name too, create a new template from the same project and disable the existing template.)

 **Tip:** It's a good idea to use the description to note the changes from the previous version of the template.

6. Select the items you want to be available when new projects are created from this template.

The replacement project template is created. Its name and description appear on the **Template** tab of the **Projects** list, accessible from your personal navigation bar.

### Enable or disable project templates


Site administrators or users with site-wide roles with project administration permissions can enable/disable project templates.

In TeamForge, enable a project template to make it available for use for creating new projects.

1. On **My Page**, click **Projects** and select the **PROJECT TEMPLATES** tab.
2. To make the project template available for use, select the template and click **Enable**.
3. To stop making the project template available for use, click **Disable**.

### **Important:**

- You can create new projects only from the list of enabled project templates.
- The **Template Name** field displays the name of the template using which the project was created. However, post upgrade to TeamForge 16.7 (or later), you will see the Template Name as "not available" for projects created in TeamForge 16.3 or earlier.
- The **Template Name** field shows a hyphen (-) in cases where projects are created not from a template.
- The template name is struck through in cases where the template used to create the project was deleted.

 **Note:** Projects created using a template will now have the information about the template used for project creation.

### Require documents to be associated with artifacts

To help reduce the problem of "orphan" documents, require users who create a document anywhere on the site to associate the document with an artifact.

"Orphan" documents are documents that are abandoned because they are not connected to any tracked activity.

1. Open the `conf/site-options.conf` file in a text editor.
2. Change the value of the `sf.requireAssociationOnDocumentCreate` variable to true
3. Change the value of the `sf.allowedAssociationTypeOnDocumentCreate` variable to `[TrackerArtifact]`
4. If you want to prevent users from associating documents with closed artifacts, change the value of the `sf.requireArtifactToBeOpenOnDocumentAssociation` variable to true
5. Save `conf/site-options.conf`.
6. Recreate the runtime environment.

```
./install.sh -V -r -d /opt/collabnet/teamforge
```

### Manage project categories

To help users navigate your site, help them sort projects into categories that make sense.

#### Add a project category

When you set up project categories for your site, project administrators can use this taxonomy to organize their projects.

You can create any number of top-level categories and any number of sub-category levels.

1. Go to **My Workspace > Admin**.

2. Click **CATEGORIES** from the **Projects** menu.
3. In the **Project Categories** tree, find the location where you want to create the new category.
  - Highlighting **Project Categories** creates a new top-level category.
  - Highlighting any category creates a sub-category beneath it.
4. Click **New**.
5. In the **Create Category** window, write a name and description for the category.
6. Click **Save**.

The category is created. It appears in the **Project Categories** navigation tree, and is available for use by all project administrators when categorizing their projects.

#### *Edit a project category*

A project category's membership and function may change over time. If it does, you can update the category's name or description.

1. Go to **My Workspace > Admin**.
2. On the site administration navigation bar, click **CATEGORIES**.  
The **Project Categories** tree displays the hierarchy of existing categories.
3. In the **Project Categories** tree, find the category that you want to edit.
4. Make the changes you need and click **Update**.

#### *Move a project category*

You can reorganize projects by moving a project category to another place in the project category hierarchy.

You can move a project category in the following ways:

- From a top-level category to a sub-category
- From a sub-category to a top-level category
- From a sub-category to another sub-category

When you move a project category, any sub-categories that it contains are also moved to the destination category.

1. Go to **My Workspace > Admin**.
2. Click **CATEGORIES** from the **Projects** menu.
3. On the **Edit Category** page, in the **Project Categories** section, click the project category you want to move.
4. On the **Edit** menu, click **Cut**.
5. Find the location to which you want to move the selected project category.  
You can move a project category either to the root category or into any other project category.
6. Choose the **Paste** option from the **Edit** menu.

The project category is now moved to the selected destination.

#### *Delete a project category*

If you no longer need a project category, you should delete it.

When you delete a project category, all of its sub-categories are also deleted.


1. Go to **My Workspace > Admin**.
2. Click **CATEGORIES** from the **Projects** menu.  
The **Project Categories** tree displays the hierarchy of existing categories.
3. Using the document tree, find the project category that you want to delete.
4. Choose **Delete** from the **EDIT** menu.

The project category and all of its sub-categories are deleted.

#### *Stop using project categories*

If you do not need to sort projects into categories, remove the ability to do so on your site.

By default, project categorization is disabled for new TeamForge installations.


 **Note:** Disabling project categorization does not delete categories you have already set up.


1. Go to **My Workspace > Admin**.
2. Click **CATEGORIES** from the **Projects** menu.
3. On the **Edit Category** page, select **Disabled** for **SITE-WIDE CATEGORIZATION** and click **Update**.

Project categorization is disabled for your TeamForge installation.


### Add a parent project to your project

You can coordinate work among multiple projects; enable the user, user group and role inheritance by adding a common parent project to several projects.

 **Note:** A subproject can have only one parent. You can change or remove that parent at any time.

 **Note:** You must be a project administrator or site administrator to add/edit or remove the parent projects.


1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Settings** page, click **Add Parent**.
3. Choose a parent project that it makes sense for your project to belong to, and click **Update**.

 **Note:** You cannot add a parent project for the 'Look' project as it is a special project in itself.

### Separate a subproject from its parent project


When a subproject grows beyond its original scope, you may want to make it stand-alone project or move it to a different project hierarchy.

By removing the association with a parent project, you can manage the subproject as a separate project.


 **Note:** Only one parent project can be selected for a subproject. However, the parent project can be changed or removed, as required.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. On the **Project Settings** page, click **Edit Parent**.
3. Change the parent project as required.

You can be a project administrator or a site administrator to change or remove a parent project. As a project administrator, you can remove or change a parent project only if you have administrator permissions for both the projects that are being linked.

 **Note:** A parent project can be removed or changed only when its members, user groups and roles are not in use in any other project. In other words, you can not remove/change a parent project while its members, user groups or roles are in use in any other project.

- • On the **Choose a Parent Project** page, select the desired parent project and click **Update**.

 **Note:** If the project hierarchy exists, the project and its subprojects are moved only under the project from which members, user groups or roles are inherited. If project hierarchy does not exist and no inheritance is in use, the project is made a 'Root' project.

- • Click **Remove Parent** in the **Project Settings** page to remove the association with parent project.
- On the pop-up message box, click **OK**, if you wish to make the project a 'Root' project.

The project hierarchy is changed or removed in accordance with role based access control and inheritance rules.

### Provide source control services

You can add a source control server to the CollabNet TeamForge environment at any time after CollabNet TeamForge installation is complete.

### Integrate a source code server

A site must have one or more servers to handle source code repositories and users. The source code server can be the same server as the application server or a separate server.

When you set up a managed software configuration management (SCM) server, you enable users to create, manage and share repositories through CollabNet TeamForge.

- ☞ **Note:** The ability to add integration servers depends on the value of the `DISABLE_CREATE_INTEGRATION_SERVERS` flag in the `site-options.conf` file. You can add new integration servers when the flag is set to its default value of "false".

You can integrate more than one source code server of a given type. For example, you can have two or more Subversion servers on your site. Consult a system administrator about the requirements for setting this up.

- ☞ **Tip:** If you use a source code solution other than Subversion or CVS you can integrate it using the CollabNet TeamForge SOAP APIs. This enables you to exchange commit data with any SCM application. Consult your CollabNet TeamForge system administrator.
- ☞ **Note:** CVS servers that integrate with CollabNet TeamForge must use the native UNIX/Linux authentication method, and not external authentication mechanisms such as NIS, NIS+, Winbind, Active Directory, or LDAP. TeamForge creates and manipulates local system accounts using the default `useradd`, `usermod`, `groupadd`, `groupdel`, and `userdel` commands. It expects to find any accounts or groups it created in `/etc/passwd` and `/etc/group`.

1. Go to **My Workspace > Admin**.
2. Click **Projects > INTEGRATIONS**.
3. On the **SCM INTEGRATIONS** page, click **Create**.
4. On the **Create Integration** page, write a name and description for the integration.
5. Choose the type of SCM server you want.

- ☞ **Restriction:** When you give a group access to a CVS or Wandisco Subversion repository, members of the group can view the repository but cannot do repository actions, such as commit and update. You must assign those permissions to users individually.

- ☞ **Note:** The SCM Adapter option only works if you have created your own SCM integration using the CollabNet TeamForge SOAP APIs.

6. Supply the host name for the **Soap Service Host**.

This is the network address of the machine on which the integrated service, such as Subversion, is running.

- ☞ **Note:** The default `localhost` will work only if the integration server is on the same server as the CollabNet TeamForge server.

7. Leave the default values in the **SOAP Service Port** field.

8. Specify whether users will use SSL to connect to their repositories.

9. Change the **Repository Root** value if you want to store the repository on your server in a different location. The repository root is the top-level directory under which all source code repositories reside.

10. Select the **Requires Approval** check box if you want an administrator to approve all repositories created on the server. By default, unmanaged servers require approval for all repositories, because repositories must be created and integrated manually.

- ☞ **Note:** By default, repositories created (or deleted) by site administrators and users with site wide role (with Integrations - SCM INTEGRATIONS permission) need no approval.

11. Supply the URL by which your users will access the service.

This will be of the form `http://<myscmserver.com>/integration/viewvc/viewvc.cgi`

- ☞ **Important:** If your system administrator has upgraded your site from SourceForge Enterprise Edition 4.4 or earlier, remove the port number in the SCM Viewer URL.

- ☞ **Tip:** If you are working with a CVS server that uses Pserver authentication, ask your system administrator for the right URL.

12. The **Use Internal Code Browser** option is selected by default to allow the users to access the TeamForge code browser for a Subversion or Git server. For more information about this feature, see [Get the code](#).


 **Important:**

- Software requirements for using TeamForge code browser, if Subversion or Git is running on a separate server: Subversion Edge 5.1.0 and TeamForge - Git integration 8.4.4 and later.
- You need to specify the **SCM Viewer URL** (*Step 11*), the way you would for ViewVC or Gitweb. The 'http://' or 'https://', the host name and the port number (if any) need to be set correctly. As a minimum requirement, the URL should point to the root of the SCM http server and the domain name.
- If the SCM server is not the same as the TeamForge server, make sure you install SSL certificates or visit the URL for the site directly from your browser, and import and trust the certificate into the browser.

**13. Click Save.**

CollabNet TeamForge attempts to validate the SCM viewer URL. If it cannot validate the URL, you can:

- Correct it if you have entered it incorrectly.
- Select **Save with errors** if the URL is different for an end user than it is for the CollabNet TeamForge server; for example, if you have a firewall in place.
- If you are adding a managed source code server, it is now added. All projects can now establish repositories on the server.
- If you are adding an unmanaged CVS server, all projects can now request repositories on the server. A CollabNet TeamForge administrator must create and integrate them manually.


 **Note:** Only CVS servers can run unmanaged.

**Approve or reject a source control repository request**


When a user requests a source code repository on a source control server for which you have required approval, a CollabNet TeamForge administrator must approve the request before the repository is created.

When adding a source control server integration, you have the option to require CollabNet TeamForge administrator approval for all repositories created on the server.

When a user requests a source code repository on a managed SCM server, the repository is created automatically after it is approved by a CollabNet TeamForge administrator.

 **Note:** Before approving a source code repository request for an unmanaged SCM server, a CollabNet TeamForge administrator must create and integrate the repository manually.

1. Click **Admin** in the TeamForge navigation bar.
2. Click **Integrations** from the **Projects** menu.
3. Click the **PENDING SCM INTEGRATIONS** tab.

 **Note:** Non-site administrators can now access the SCM Integrations tab if they have permission to manage SCM integrations.

4. From the list of pending SCM repository requests, select the SCM repositories that you want to approve.
  - Click **Approve** to approve the repository.
  - Click **Reject** to reject the project and remove it from the list.

The person who requested the repository receives an email notification when the repository is approved or rejected. If you entered a comment, that also appears in the email notification.

**Approve or reject an unmanaged SCM server access request**

When a user asks for access to an unmanaged SCM server (such as an unmanaged CVS server), an administrator must approve or reject the request.

When a project administrator assigns a role that provides SCM access to a project member, a CollabNet TeamForge administrator must manually create the user account on the unmanaged SCM server. Because user creation on unmanaged SCM servers is not managed by CollabNet TeamForge, TeamForge cannot verify that the user account has been created. A TeamForge administrator must confirm that he or she has created the account.

Requests for SCM access removal are also submitted for approval by the TeamForge administrator.

1. When you get a repository access request on an unmanaged SCM server, log on to the server and create the requested user account on the unmanaged SCM server.
2. Log on to the TeamForge application server.
3. Go to **My Workspace > Admin**.
4. Click **Integrations** from the **Projects** menu.
5. Click the **SCM ACCESS REQUESTS** tab.
6. As the user account is created, select the repository access request from the **Repository Access Requests** section and click **Approve**.

You may also click **Reject** to reject the repository access request.

The user receives an email notification that the user account has been created and that the repository access request has been approved or rejected.

#### **Approve or reject repository delete requests**

SCM system maintenance requests (such as a repository delete request) must be approved (or rejected) by an administrator.

1. Go to **My Workspace > Admin**.
2. Click **Integrations** from the **Projects** menu.
3. Click the **SCM ACCESS REQUESTS** tab.
4. Select the repository deletion request from the **System Maintenance Requests** section and click **Approve**.

You can click **Reject** to reject the repository deletion request.

The user receives an email notification that the repository deletion request has been approved or rejected.

#### **Edit SCM integration details**

You can move or reconfigure a source control server without having to reintegrate the server into CollabNet TeamForge .

 **Tip:** If you edit or lose the permissions on your SCM server, use the **Synchronize Permissions** button on the **SCM Integrations** page to recreate the correct permissions on your SCM server from CollabNet TeamForge .

1. Click **Admin** in the CollabNet TeamForge navigation bar.
2. Click **Integrations** from the **Projects** menu.
3. On the **SCM INTEGRATIONS** tab, click the name of the SCM integration you want to edit.
4. On the **Edit Integration** page, make the changes you need and click **Save**.

#### **Move a source code repository**

When the existing code base for an application may need to be managed by a different team or project, you can move the source code repository from the first project into the other one.

1. Open the project from which you want to move a source code repository.
2. Click **SOURCE CODE** from the **Project Home** menu.
3. Select the repository you want to move and click **Cut**.
4. Open the project into which you want to move the repository.
5. Click **SOURCE CODE** from the **Project Home** menu.
6. Click **Paste**.

#### **Connect to a Subversion Edge console**

When a Subversion Edge server has been converted to a SCM Integration server in TeamForge, you can log into its management console from within TeamForge.

1. Go to **My Workspace > Admin**.
2. Click **Integrations** from the **Projects** menu.
3. In the list of servers on the **SCM INTEGRATIONS** page, click the **Open Console** link for the Subversion Edge server you want to connect to.



**Note:** Only Subversion Edge servers have this link.

The login page for the server's management console appears. You can log in using your TeamForge administrator credentials and view statistics such as network throughput and disk space usage for the server.

### Manage replicas

A replica server in TeamForge is a Subversion Edge server that replicates the content of an existing core SCM integration server.

#### *Approve a replica server request*

When there is a request for a replica of a core SCM integration server, a TeamForge administrator must approve the request before the replica is created.

Replica requests from a TeamForge admin user or site administrator are automatically approved. Replicas requested by other users need approval by a TeamForge administrator.

1. Go to **My Workspace > Admin**.
2. Click **Integrations** from the **Projects** menu.
3. On the **SCM INTEGRATIONS** tab, click the **Pending SCM Replicas** tab.
4. From the list of pending SCM replica requests, select the SCM replicas that you want to approve.
  - Click **Approve** to approve the replica.
  - Click **Reject** to reject the replica and remove it from the list.

When a replica is approved, it is listed in the **SCM INTEGRATIONS** tab beneath the master SVN server.

#### *Edit replica settings*

As a TeamForge site administrator, you can configure replica settings for the polling frequency of the master, and repository initialization and synchronization events.

Replication events, such as creating new repositories and synching commits, are stored in a queue on the TeamForge Application Server. The replica server polls the TeamForge Application Server for new events and then processes those events on the replica server.

These events are divided into two separate pools:

- New repository initializations - this includes creating the repository and performing the initial synchronization of the content.
- All other events

When existing repositories are selected for replication, this can take a long time. It could take many hours or even days to fully replicate an entire repository across a WAN. These big events are processed in their own thread pool, so that other repositories which are already synchronized don't have to wait in line for them to finish.

For each pool, you can define how many simultaneous events will be processed. The higher the number, the greater the potential load on both the TeamForge replica server and the Subversion master. However, this can also decrease the wait time for a given commit to appear on the replica server.

1. Go to **My Workspace > Admin**.
2. Click **Integrations** from the **Projects** menu.
3. On the **SCM INTEGRATIONS** tab, click the name of the Subversion Edge replica you want to edit.

The **Edit System** page for the replica appears. Here's an example:

### Subversion Edge Replica

Name: \*

Description: \*

Command polling interval: \*  seconds

Maximum simultaneous new repository initializations: \*

Maximum simultaneous repository synchronizations: \*

Hostname: cu135.cloud.sp.collab.net

Master Integration Server: [Subversion](#)

Managed By: Subversion Replication (svnedge)

Approval Status: Approved

Last Contact: Wed Mar 23 15:13:28 PDT 2011

---

#### Recent Replication Command History

Status	Replica Command	Repository	Revision	Created Date	Last Contact
	Sync repo	reps1004: i18n	8	Wed Mar 23 12:29:25 PDT 2011	Wed Mar 23 12:29:29 PDT 2011
	Start replicating repo	reps1005: Docs	-	Wed Mar 23 12:16:36 PDT 2011	Wed Mar 23 12:22:35 PDT 2011
	Start replicating repo	reps1004: i18n	-	Wed Mar 23 12:07:43 PDT 2011	Wed Mar 23 12:07:58 PDT 2011
	Update server properties	-	-	Wed Mar 23 12:07:26 PDT 2011	Wed Mar 23 12:07:58 PDT 2011
	Replica server request approved	-	-	Wed Mar 23 12:07:11 PDT 2011	Wed Mar 23 12:07:21 PDT 2011

4. Change the replica name or description if required.



**Tip:** Including the geographic location would help users select a nearby replica.

5. Set **Command polling interval** to define how frequently the replica polls the master looking for new events. The replica will process all new events when it polls.  
The default value for this setting is 60 seconds, but it can range from 5 to 1000000 seconds.
6. Set **Maximum simultaneous new repository creations** to a low value.  
New repository initializations can take a long time and generate a lot of load. So you wouldn't want to allow too many of them to run at once. This value can range from 1 to 100, but we suggest you keep it at 3 or less.
7. Set **Maximum simultaneous repository synchronizations** taking into account how many repositories you will be replicating and how many you think are likely to have commits occurring within the polling interval.  
This value can range from 1 to 100. You may want to set this higher than the previous field, but we suggest you keep it at 10 or less. There's no reason to enter too high a number because you are merely specifying how many synchs can run at the exact same time -- and it never runs more than one per repository.
8. When you've made your changes, click **Save**.

#### *Remove a replica*

When you remove a replica from TeamForge, it is restored to a Subversion Edge server in standalone mode.

1. Go to **My Workspace > Admin**.
2. Click **Integrations** from the **Projects** menu.
3. In the **Edit System** page for the replica, click **Delete**.

The replica is removed from TeamForge. The repositories that existed on the replica are deleted.

#### **Integrate an unmanaged CVS repository with CollabNet TeamForge**

After creating the repository, you must modify two triggers, `verifymsg` and `loginfo`, to complete the integration with CollabNet TeamForge.

The triggers must be modified manually, because much of the information in the triggers is specific to your CollabNet TeamForge installation.

Each trigger must be modified to run a `.java` program. These topics provide an annotated example and instructions on modifying and installing the trigger files.



**Note:** You must create the CVS repository before you can modify and install the trigger files.



**Note:** CVS servers that integrate with CollabNet TeamForge must use the native UNIX/Linux authentication method, and not external authentication mechanisms such as NIS, NIS+, Winbind, Active Directory, or LDAP. TeamForge creates and manipulates local system accounts using the default `useradd`, `usermod`,



groupadd, groupdel, and userdel commands. It expects to find any accounts or groups it created in /etc/passwd and /etc/group.

### Create a logininfo trigger file

Here's how to create the logininfo trigger file.

1. Check out the module CVSROOT from the CVS repository.

The CVSROOT module is automatically created in all new CVS repositories, and contains the `verifymsg` and `logininfo` files. When created, the files contain only comments.

2. Begin modifying the file by inserting ALL followed by a tab.

This enables CollabNet TeamForge to see all modules in the repository.



3. Modify the path to your Java home directory as in this example and add it to the file, followed by a single space.

Example	Modification
<code>/usr/local/java/j2sdk1.4.1_01/bin/ java -cp</code>	Replace <code>/usr/local/java/j2sdk1.4.1_01/</code> with the directory in which you have installed Java.

4. Modify the paths to the `.jar` files as shown here, then add them to the file separated by colons.

Examples	Modifications
<code>/yourdirectory/sourceforge_home/ lib/saa.jar /yourdirectory/ sourceforge_home/lib/axis-1.1rc1- va.jar /yourdirectory/ sourceforge_home/lib/ saturn.jar /yourdirectory/ sourceforge_home/lib/jaxrpc.jar / yourdirectory/sourceforge_home/ lib/externalintegration.jar / yourdirectory/sourceforge_home/ lib/commons-discovery.jar / yourdirectory/sourceforge_home/lib/ commons-logging.jar /yourdirectory/ sourceforge_home/lib/log4j-1.2.8.jar</code>	Replace <code>/yourdirectory/</code> with the directory in which you have installed CollabNet TeamForge .

5. Modify the parameters as follows:

Example	Modification
<code>-Dintegration.name=exsy1001</code>	Replace <code>exsy1001</code> with the ID of your SCM server integration with CollabNet TeamForge .   <b>Note:</b> To find the ID, click <b>Admin &gt; Integrations &gt; Integration Name</b> . The ID appears at the end of the URL.
<code>-Dcvs.cvsroot=\$CVSROOT</code>	No change
<code>-Dappserver.url=http://localhost:8080/ce-soap60/ services/ScmListener</code>	Replace <code>localhost</code> with the name of your machine running CollabNet TeamForge .   <b>Note:</b> If needed, replace <code>8080</code> with your SOAP server port.
<code>-Dcvs.username=\$USER</code>	No change
<code>-Dlog4j.configuration=file:/yourdirectory/ sourceforge_home/etc/externalintegration/log4j.xml</code>	Replace <code>/yourdirectory/</code> with the directory in which you have installed CollabNet TeamForge .

Example	Modification
-Dexternalintegration.triggers.log.dir=/tmp	No change


6. Add the following Java class containing the body of the trigger.

```
com.vasoftware.sf.externalintegration.triggers.cvstriggers.LogInfo
    %{sVv}
```

7. Check your file against the completed example shown on Completed loginfo trigger file.

8. When you are finished, check the modified file back into your CVS repository.

If you have not already done so, repeat the process for the verifymsg trigger file.

 **Note:** If you make an error in the file, commits made to the repository will not fail, but you will receive error messages.

#### Create a modified verifymsg trigger file

Here's how to modify the verifymsg trigger file.

1. Check out the module CVSROOT from the CVS repository.

The CVSROOT module is automatically created in all new CVS repositories, and contains the verifymsg and loginfo files. When created, the files contain only comments.

2. Begin modifying the file by inserting `.*` followed by a tab.

This enables CollabNet TeamForge to see all modules in the repository.

3. Modify the path to your Java home directory as indicated below.

Add it to the file, followed by a single space.


Example	Modification
/usr/local/java/j2sdk1.4.1_01/bin/java -cp	Replace /usr/local/java/j2sdk1.4.1_01/ with the directory in which you have installed Java.

4. Modify the paths to the `.jar` files as shown in the following table.

Then add them to the file separated by colons.


Examples	Modifications
/yourdirectory/sourceforge_home/lib/saaj.jar / yourdirectory/sourceforge_home/lib/axis-1.1rc1- va.jar /yourdirectory/sourceforge_home/lib/ saturn.jar /yourdirectory/sourceforge_home/lib/ jaxrpc.jar /yourdirectory/sourceforge_home/ lib/externalintegration.jar /yourdirectory/ sourceforge_home/lib/commons-discovery.jar / yourdirectory/sourceforge_home/lib/commons- logging.jar /yourdirectory/sourceforge_home/lib/ log4j-1.2.8.jar	Replace /yourdirectory/ with the directory in which you have installed CollabNet TeamForge .

5. Modify the parameters as follows:

Examples	Modifications
-Dintegration.name=exsy1001	Replace exsy1001 with the ID of your SCM server integration with CollabNet TeamForge .   <b>Note:</b> To find the ID, click <b>Admin &gt; Integrations &gt; Integration Name</b> . The ID is at the end of the URL.

Examples	Modifications
-Dcvs.cvsroot=\$CVSROOT	No change
-Dappserver.url=http://localhost:8080/ce-soap60/services/ScmListener	Replace localhost with the name of your machine running CollabNet TeamForge . If needed, replace 8080 with your SOAP server port.
-Dcvs.username=\$USER	No change
-Dlog4j.configuration=file:/yourdirectory/sourceforge_home/etc/externalintegration/log4j.xml	Replace /yourdirectory/ with the directory in which you have installed CollabNet TeamForge .
-Dexternalintegration.triggers.log.dir=/tmp	No change

6. Add the following Java class containing the body of the trigger:  
com.vasoftware.sf.externalintegration.triggers.cvstriggers.VerifyMessage
7. Check your file against the completed example shown on Completed verifymsg trigger file.
8. Check the modified file back into your CVS repository.
9. If you have not already done so, repeat the process for the logininfo trigger file.

 **Note:** If you make an error in the file and check it back into your CVS repository, all commits made to the repository will fail. To correct the file, you cannot simply check it out and fix it. You must fix it on the repository. After fixing the file on the repository, you must then check it out, fix it again, then check it back in.

### Change your site's look and feel

You can redesign some aspects of your site to suit your organization's needs and preferences.

These instructions support only some basic types of customization. Almost infinite varieties of customization are possible. To get into specific customization options in more detail, search or post a question on the TeamForge [discussion forum](#) or talk to your CollabNet representative.

### Customize TeamForge

You can customize your TeamForge site by building a maven project and uploading the customization jar file that extends or customizes TeamForge.

Maven projects are built and packaged to generate TeamForge customization jar and MANIFEST.MF files. The generated customization jar file is then uploaded to TeamForge. When you upload a customization jar, it is processed and if it has a custom event, it is registered. Later, if it has customizations, they are cached by the customization mechanism for a cost-free access at every request. Cached customizations are then served by the following three servlets:

Servlet	Description
/ctf/api/main/js-customization	Retrieves all the Javascript customizations.
/ctf/api/main/css-customization	Retrieves all the CSS customizations.
/ctf/js/modules/customization- <customization-name>/<resource-name>	Resolves the resource relative to the main folder configured for the given customization name.

The customization mechanism provides access to all the enabled customizations in the cache.

A customization jar can contain:

- Custom events
- Javascript customizations
- CSS customizations
- Custom bundles
- Here's a sample jar file structure:

```
JAR
|
+-- META-INF
|   +-- MANIFEST.MF
+-- js/
|   +-- custom.js
+-- css/
|   +-- custom.css
+-- bundles/
|   +-- bundle_en.html
|   +-- img/
|       +-- footer.png
```

- Here's a [sample customization jar file](#).

While custom events are configured through an `events.xml` file in the `META-INF` folder in the `jar` file, Javascript, CSS and custom bundles are configured through `META-INF/MANIFEST.MF` entries.

Here's a list of `META-INF/MANIFEST.MF` entries:

MANIFEST.MF entries	Description
CTF-Customizations-Enabled	The entry to enable or disable a customization. This entry applies to custom event and customizations. <ul style="list-style-type: none"> <li>• Custom event and customizations are applied if this entry is set to True (default value).</li> <li>• Custom event and customizations are not applied if this entry is set to False.</li> </ul>
CTF-Customization-Name	The entry to set the name of the customization to be used for getting bundles.
CTF-Customizations-Priority	The entry to set the priority for customizations. Allows you to specify the 'priority' of the customization. Customizations are sorted by the servlets based on the 'priority'. Customizations with low priority are included at the end. The priority value could be from 1 to 100, 100 being the default value.
CTF-JS-Customization	Path to a Javascript file.
CTF-CSS-Customization	Path to CSS stylesheet.
CTF-Bundle-Customization	Path to the main bundles directory.

### An illustration of how to add a CSS customization

#### 1. Build a customization project.

Maven project structure that includes a custom stylesheet file as a resource:

```
css-customization% find -type f
./pom.xml
./src/main/resources/custom/custom.css
css-customization %
```

```
css-customization % cat src/main/resources/custom/custom.css
div.core-footer {
    background-image:url('/ctf/js/modules/customization-mybundles/img/footer.png');
}
css-customization %
```

The pom.xml descriptor, uses the packaging plugin for setting the needed MANIFEST.MF properties:

```
<project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://
www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://
maven.apache.org/xsd/maven-4.0.0.xsd">
  <modelVersion>4.0.0</modelVersion>

  <groupId>com.ctf.customizations.samples</groupId>
  <artifactId>css-customization</artifactId>
  <version>1.0-SNAPSHOT</version>
  <packaging>jar</packaging>
  ...
  <build>
    <plugins>
      <plugin>
        <groupId>org.apache.maven.plugins</groupId>
        <artifactId>maven-jar-plugin</artifactId>
        <version>2.4</version>
        <configuration>
          <archive>
            <manifestEntries>
              <CTF-Customizations-Enabled>True</CTF-Customizations-
Enabled>
              <CTF-Customization-Name>mystyles</CTF-Customization-Name>
              <CTF-CSS-Customization>custom/custom.css</CTF-CSS-
Customization>
            </manifestEntries>
          </archive>
        </configuration>
      </plugin>
    </plugins>
  </build>
  ...
</project>
```

## 2. Package the Maven project.

```
css-customization % mvn package
[INFO] Scanning for projects...
...
[INFO] --- maven-jar-plugin:2.4:jar (default-jar) @ stylesheet ---
[INFO] Building jar: /home/matias/workspaces/ctf/css-customization/target/
mystyles.jar
[INFO]
-----
[INFO] BUILD SUCCESS
[INFO]
-----
[INFO] Total time: 2.959s
[INFO] Finished at: Wed May 21 20:26:28 ART 2014
[INFO] Final Memory: 8M/105M
[INFO]
-----
css-customization %
```

Generated jar:

```
css-customization % jar tvf target/mystyles.jar
 207 Wed May 21 20:26:28 ART 2014 META-INF/MANIFEST.MF
  17 Wed May 21 20:26:26 ART 2014 custom/custom.css
```

```

1092 Wed May 21 20:16:06 ART 2014 META-INF/maven/
com.ctf.customizations.samples/css-customization/pom.xml
132 Wed May 21 20:26:28 ART 2014 META-INF/maven/
com.ctf.customizations.samples/css-customization/pom.properties
css-customization %

```

Generated MANIFEST.MF:

```


Manifest-Version: 1.0
Built-By: matias
Build-Jdk: 1.7.0_55
CTF-Customization-Name: mystyles
CTF-CSS-Customization: custom/custom.css
CTF-Customizations-Enabled: True
Created-By: Apache Maven 3.0.4
Archiver-Version: Plexus Archiver

```

3. Upload the generated jar file as a custom event so that your customization is applied to TeamForge pages.
  - a) Log on to TeamForge as a site administrator.
  - b) Select **My Workspace > Admin**.
  - c) Select **Projects > System Tools > Customizations** and click **Create**.
  - d) Click **Choose File**, select the customization jar file and click **Add**.
4. Enable, disable, delete or download a customization.
  - a) Select **My Workspace > Admin**.
  - b) Select **Projects > System Tools > Customizations**.
  - c) Select one or more customizations (check boxes) you want to enable or disable and click **Enable** or **Disable**.
  - d) Select one or more customizations (check boxes) you want to delete and click **Delete**.
  - e) Select one or more customizations (check boxes) you want to download and click **Download**.

### Customize anything on your site


Follow these general instructions to customize a page, picture, text string, or other element on your site.

 **Important:** Custom branding changes can be overridden when your site is upgraded to a new version. You may have to reapply any look-and-feel modifications after an upgrade.

1. Download the sample branding files.

Choose one of these files:

- *Basic branding package*: Contains the files you need to do most of your branding tasks. Safest to use this file if you are doing your own branding.
- *Advanced branding package*: Contains all the files that can be customized. For use when someone from CollabNet is helping you with your branding.

 **Note:** It is important that you have the most recent available version of this archive as a starting point. Check that the version number at the top of the `readme.txt` file in your copy of the branding package is the same as your version of the application. If it is not the same, check [www.collab.net](http://www.collab.net) to see if there is a more recent version.

2. In the `look` project, check out the branding repository.
3. Copy the default version of the appropriate file from the branding zip file to the equivalent directory in your local copy of the branding repository.
4. Change the file to produce the results you want.  
For example:
  - To change a logo on your site's home page, overwrite the `home.gif` file with a new file of the same name.
  - To change a logo on a project home page, overwrite the `project.gif` file with a new file of the same name.
5. Commit the changed files into your site's branding repository.

- 👉 **Important:** Your branding repository does not have to contain all the files that are in the sample branding zip file, but the structure of your repository must be an exact mirror of the structure of the sample file set.

### Customize the home page of your site

To change the content of your site's main page, replace the `home.vm` file or add either `domain_home.html` or `DomainHome.html.html` file to the `html` folder in the branding repository.

- 👉 **Note:** For the general steps for changing the look and feel of a page, see [Customize anything on your site](#) on page 358.

The `\branding\templates\sfmain\home.vm` template controls the look, feel and structure of the standard home page. The default version allows users to log in and sign up for new user accounts, if CollabNet TeamForge is configured to allow user self-creation.

If the `DomainHome.html.html` or `domain_home.html` file is checked into the branding repository, the contents of the file are displayed as the site home page.

- 👉 **Tip:** If both `DomainHome.html.html` and `domain_home.html` files exist in the repository, the contents of the `DomainHome.html.html` are displayed.

Edit the `home.vm` template to produce the page you want.

You can change these objects on the site home page:

Object	Description
siteNews	<p>The html block that shows site news.</p> <ul style="list-style-type: none"> <li>👉 <b>Note:</b> <ul style="list-style-type: none"> <li>The "siteNews" html block itself is not customizable.</li> <li>By uncommenting this "siteNews" object and commenting out "communityNews" object, site news can be displayed across all projects in the site.</li> <li>In addition to enabling the "siteNews" html block, you must set <code>ENABLE_SITE_NEWS</code> token to true if you want site news published on your site's home page.</li> </ul> </li> </ul>
communityNews	<p>The html block that shows community news.</p> <ul style="list-style-type: none"> <li>👉 <b>Note:</b> By uncommenting this object, and commenting out "siteNews" object, community news (news from the look project) can be displayed across all projects in the site.</li> </ul>
mostActiveProjects	<p>The html block that shows the most active projects. The html block itself is not customizable.</p>
displayActivityGraph	<p>A flag that indicates that the activity graph should be displayed.</p>
displayTeamForgeLinks	<p>A flag that indicates that CollabNet TeamForge quick links should be displayed.</p>

### Customize the home page of projects

To change the default main pages of the projects on your site, edit the `project_home.vm` file.

- 👉 **Note:** For the general steps for changing the look and feel of a page, see [Customize anything on your site](#) on page 358.


Edit the `project_home.vm` template to produce the project page you want.

You can change these objects on the project home page:

Object	Description
projectData	The object that contains the information about the project. It implements the interface <code>com.collabnet.ce.customization.IProjectData</code> .
adminList	The list of project administrators. Each object of the list implements the interface <code>com.collabnet.customization.IUserRow</code> .
memberList	The list of project members. Each object of the list implements the interface <code>com.collabnet.customization.IUserRow</code> .
projectMember	A flag that indicates that the user is a member of the project.
joinProjectButton	The button that contains the link to the Join Project page. It returns a <code>com.collabnet.ce.customization.widgets.Button</code> .
useCustomHomePage	A flag that indicates that the page shows the Wiki Home page instead of the standard Home page.
customHomePage	The html that displays as the Project Home page.
editCustomHomePageButton	The button that is used to edit the custom Home page. It returns a <code>com.collabnet.ce.customization.widgets.Button</code> .
projectAdmin	A flag that indicates whether or not the current user is a Project Admin.
useCustomProjectLogo	A flag that indicates that the Wiki project logo image will be used instead of the standard project logo.
customLogoPathString	The url from where the custom project logo image can be loaded.

### Change your site's outgoing emails

When your site sends out automated emails, the text of the emails can be customized to fit your site's specific needs.

 **Note:** Before customizing your site, download the branding files. See [Customize anything on your site](#) on page 358.

You control screen labels and messages by overriding the resource bundle keys that specify the text strings that appear in Velocity macros and JSPs.

1. In your local copy of the branding repository, create a directory called `templates/mail`.
2. In the `templates/mail` directory, create a file containing the custom content for an email that the system sends out. Give the file the same name as the equivalent sample email file in the branding files package. For example, to override the email that is sent out to new members of the site, name the file `templates/mail/user_welcome.vm`. Use Velocity syntax to identify the parts of the email, like this:

```
##subject
Welcome to our TeamForge site!
##subject
##body
Here is the content that I want to appear in emails coming from my site...
##body
```





**Note:** To customize a template in a specific language, identify the locale as an extension to the file name. For example, to create a user welcome file in Japanese, name the file `templates/mail/user_welcome_ja.vm`

3. Commit your new and changed files into the repository.

### Provide external applications

If your site's users need access to an application or web site that is not part of CollabNet TeamForge, you can make it available by linking or integrating from within your CollabNet TeamForge site.

#### Provide a linked application

To bring an external tool into your environment quickly and easily, set it up as a linked application.

##### *Create a site-wide linked application*

When you create a site-wide linked application, it appears in all projects on your CollabNet TeamForge site.

Site-wide linked applications are especially useful for incorporating corporate standard external applications, such as a company intranet sites, into your CollabNet TeamForge installation.

1. Go to **My Workspace > Admin**.
2. On the site administration navigation bar, click **INTEGRATIONS > SITE-WIDE LINKED APPLICATIONS**.
3. Click **Create**.
4. On the **Create Site-wide Linked Application** page, provide a name for the linked application. This name appears on the link in the CollabNet TeamForge navigation bar.
5. Enter the server location or URL for the linked application.
6. Select whether you want to enable single sign-on for the linked application.
  - If you use single sign-on, access to the linked application is managed through the CollabNet TeamForge authentication system. Users are not required to log into the linked application after they have logged into CollabNet TeamForge.
  - If you do not use single sign-on, users will be required to log in to the linked application using its native authentication system.
7. Choose how you want the linked application to appear when a user clicks it.
  - In the same window. The linked application takes over the entire browser window, replacing whatever the user was looking at.
  - In a new window. The linked application launches in a separate browser window.
  - In an iframe. The linked application appears in a box in the same window, framed by the CollabNet TeamForge site's header and navigation controls.
8. Click **Save**.

A link for the site-wide linked application is added to your CollabNet TeamForge navigation bar. Clicking the link displays the application in the main CollabNet TeamForge window.


##### *Edit a site-wide linked application*

When the use patterns of a linked application change, you may need to change the way the application integrates with CollabNet TeamForge .

1. Go to **My Workspace > Admin**.
2. On the site administration navigation bar, click **INTEGRATIONS > SITE-WIDE LINKED APPLICATIONS**.
3. Click the name of the site-wide linked application that you want to edit.
4. On the **Edit Site-wide Linked Application** page, make the changes you need.

You can edit these elements:

- The name of the application.
- The application's URL.
- Whether the application uses single sign-on.
- Whether the application is displayed in a new window, in the same window, or in an IFrame.

 **Note:** You cannot change the application icon.

5. Click **Save**.

### Provide integrated applications for use in TeamForge projects

To make a tool comprehensively available to your users, set it up as an integrated application.

Before you can make an integrated application available to project admins, your system administrator must integrate the application with your site. This may involve modifying the application. How this is done depends in part on the application.

#### *Integrate an external application into a TeamForge site*

When you integrate an external application into your CollabNet TeamForge site, your site's project administrators can choose to include the integrated application alongside the built-in tools in their projects.

When you have integrated the application, project administrators on your site can add it to their set of collaboration tools. Objects they create will share the core TeamForge features, such as authorization, authentication, go-urls, association, linkification, templating, Project Pages components, search, and source code management support.

1. Find the XML files that describes your integrated application.

For an example of a working XML file, see [pebble-app.xml](#) on page 776.

2. Log into TeamForge as an admin user.


3. Go to **My Workspace > Admin**.

4. Click **INTEGRATED APPS** from the **Projects** menu.

5. Click **Create**.

6. Use the **Browse** window to find the configuration file you created, then click **Next**.

7. On the **Preview** screen, review the parameters you set in the configuration file.

 **Note:** You may have to revise one or more values to ensure they are valid.

8. Click **Save**.

The application is now available for all projects on your site. You can direct project administrators to [these instructions](#) to add it to their own project toolbars.

#### *Enable or disable integrated applications*

Site administrators can enable or disable integrated applications.

In TeamForge, enable an integrated application to make it available for use in the projects.

 **Tip:** You can also disable an integrated application to restrict it from being added to projects.

1. Go to **My Workspace > Admin**.

2. Click **INTEGRATED APPS** from the **Projects** menu.

3. To make the integrated application available for use, select the integrated application and click **Enable**.


4. To stop making the integrated application available for use, click **Disable**.

The system asks for confirmation before disabling an integrated application. Click **OK** to disable the integrated application.

A disabled integrated application can be re-enabled when there is a need to use that integrated application in a project.

#### *Hide / Remove a tool (linked or integrated application)*

To stop making an application or site outside of TeamForge available to your users from inside your TeamForge projects seamlessly, hide or disintegrate an application.

 **Note:** You may have integrated several external applications per project to maximize the integrated applications feature's utility. However, because each integrated application adds an icon to the project navigation bar, a large number of integrated applications can cause horizontal scrolling. Consider removing the linked or integrated applications that are not in use.

1. Click **PROJECT ADMIN** from the **Project Home** menu.

2. On the **Project Admin Menu**, click **Tools**.
3. To hide an integrated application from the list of tools displayed, clear the **Visible** check box and click **Save**. The selected integrated application is hidden for the project. The icon of the hidden linked or integrated application disappears from your **Project Home** menu.



**Warning:** Hiding a linked or integrated application, hides all associations to the integrated application too. It also hides the component type in the Project Home's **Create Component** page.

4. You can also click the **Delete** icon to delete a linked or integrated application permanently from your project.



**Warning:** Deleting an integrated application (such as Review Board, EventQ or Binaries) means a permanent removal of the application including all related integrated application data from your project. Exercise caution before deleting an integrated application.

#### *Set site-level permissions for an integrated application*

In CollabNet TeamForge, as the site administrator you can set the site-level permissions for an integrated application.

The default access permissions are usually set using the configuring xml file for each integrated application.

However, you may want to provide permission access at site-level for some of the users.

1. Go to **My Workspace > Admin**.
2. Click **INTEGRATED APPS** from the **Projects** menu.
3. Set the permissions for the integrated application and click **Submit**.  
The permissions are set for the selected integrated application.



**Note:** It is possible that an integrated application may be added to a project by a restricted user whose project administration permissions come from a global project role. In such a case, you could be left with no one authorized to administrate the integrated application. TeamForge handles that situation in one of two ways:

- If one or more users with the "Founder project administrator" role is a member of the project, then the restricted user who did the integration gets that role too.
- If there is no project member with the "founder project administrator" role, the restricted user who did the integration gets a new administrator role called "<integrated\_application\_name> Administrator."

#### *View integrated application information*

TeamForge site administrators, and users with site-wide roles and the IAF permission, can view information about an integrated application's configuration and other details.

1. Go to **My Workspace > Admin**.
2. Click **INTEGRATED APPS** from the **Projects** menu.
3. Click the name of the integrated application that you want to view.  
The default category tab, **GENERIC**, displays all the existing integrated applications associated with the site. You can have your integrated application displayed under a separate **CATEGORY** tab by defining it in the deployment configuration file.

The configuration details of the integrated application are displayed.

#### *Edit an integrated application*

TeamForge site administrators can update an integrated application's URLs (SOAP endpoint URL, administration URL, browser URL) and other parameters by uploading XML files containing the changes.

1. Go to **My Workspace > Admin**.
2. Click **INTEGRATED APPS** from the **Projects** menu.
3. Select the integrated application you want to edit.
4. Click **Edit**.
5. On the **Edit Integrated Application** page, make the changes you need.

You can edit these elements:

- The Deployment Configuration File.
- The Application Configuration File.

6. Click the **Browse tab** and select the files to attach.
7. Click **Next**.

### *Export an integrated application*

TeamForge site administrators can export an integrated application's variables in XML format and have them available for editing.

1. Go to **My Workspace > Admin**.
2. Click **INTEGRATED APPS** from the **Projects** menu.
3. To export the integrated application, select the integrated application name.
4. From the **View Integrated Application** page, click **Export**.

The integrated application will be exported.

### **Add a custom event handler to your TeamForge site**

An event handler is a program that watches for events on a TeamForge site and communicates them to another system. You can add your own event handlers to the set that are built into TeamForge.

For example, some of your site's members may be using TeamForge alongside a legacy issue tracking system. You may want to write an event handler that listens to the "Artifact Create" event on your TeamForge site and sends the details about any newly created artifact to the legacy system through a webservice.

1. Create your custom event handler and package it as a `.jar` file.
2. Check with your system administrator that the `ENABLE_UI_FOR_CUSTOM_EVENT_HANDLERS` variable in the site configuration file is set to true.
3. Go to **My Workspace > Admin**.
4. Click **SYSTEM TOOLS** from the **Projects** menu.
5. Click **Customizations**.
6. Click **Create** and use the **Browse** control to locate your `.jar` file.
7. Click **Add**.



**Note:** If the system reports "Error Parsing Event Jar File," debug your event handler until the error message no longer appears.

Your `.jar` file is uploaded to your TeamForge site and the event cache is cleared. All the events you specified in your event handler are now captured and sent to the external web service.

### **System tools to monitor your site**

To facilitate troubleshooting, keep an eye on key data about your CollabNet TeamForge site.

#### **Check your server's status**

As a first troubleshooting step, check the **Server Status** page to see if your servers are running and connected to the main CollabNet TeamForge server.



**Note:** The **Server Status** page provides only the current status of each server. It is not a diagnostic tool, and you must correct any connection or configuration issues on the server in question.

1. Go to **My Workspace > Admin**.
2. On the site administration navigation bar, click **SYSTEM TOOLS**.
3. On the **System Tools** menu, click **Server Status**.
4. Find the status of the server you are interested in.

For each server, you see one of the following status messages:

- **OK** - The server is running and connected to the main TeamForge server.
- **Could Not Connect** - The server is not running or is not connected to the main TeamForge server.

### **Read your site's logs**

Inspecting the system logs for your CollabNet TeamForge site may product useful information for solving problems.

1. Go to **My Workspace > Admin**.
2. On the site administration navigation bar, click **SYSTEM TOOLS**.
3. On the **System Tools** menu, click **System Logs**.  
All the logs the server has written are listed.
4. On the **Site Configuration** page, click the log file you are interested in.  
For explanations of some of the most useful logs, see [Log files](#).

### Configure your site's log level

Change your site's log level in the **Configure Logging** page.

The **Configure Logging** page lets you change the application server (JBoss) log level. Changing the application server log level affects only the `server.log` and `vamessages.log` files.

1. Go to **My Workspace > Admin**.
2. On the site administration navigation bar, click **SYSTEM TOOLS**.
3. On the **System Tools** menu, click **Configure Logging**.
4. On the **Configure Logging** page, select either **INFO** or **DEBUG** from the **LOG LEVEL** drop-down list.



#### Note:

- The default log level is **INFO**.
- The `server.log` and `vamessages.log` files grow in size if you change the log level from **INFO** to **DEBUG**.
- JBoss restart is not required after changing the log level using the **Configure Logging** page.

5. Click **Save**.

### Configure your site's settings

Use the **Configure Application** tool to define your site level TeamForge settings.

The **Configure Application** tool, added in TeamForge 17.4, makes TeamForge site level settings configurable via the user interface. This comes in handy for site administrators, who otherwise would have to work with the `site-options.conf` file and then recreate the TeamForge runtime for site level configuration changes.

- To modify your site settings, select **My Workspace > Admin** and select **Projects > System Tools > Configure Application**, modify the available site settings as required and click **Save**.

#### Available settings

- **Project**

These are project settings that apply globally to all the projects in your TeamForge site.

- **NOTIFY ASSOCIATION AND DEPENDENCY UPDATES:** Notification emails can be sent to all monitoring users when an association or dependency is added to an object. You can select the **NOTIFY ASSOCIATION AND DEPENDENCY UPDATES** check box to let TeamForge send notification emails to all monitoring users when an association or dependency is added. Clear this check box otherwise.

- **Tracker**

These are tracker settings that apply globally to all the trackers in your TeamForge site.

- **MASS IMPORT ARTIFACTS LIMIT:** You can restrict the number of artifacts that can be mass-imported. Type the maximum number of artifacts that you want to allow via mass import in the **MASS IMPORT ARTIFACTS LIMIT** text box.

### Get build information

For solving some kinds of problems, it may be useful to know which software components are installed, the exact build number, and other technical information about your CollabNet TeamForge site.

1. Go to **My Workspace > Admin**.
2. On the site administration navigation bar, click **SYSTEM TOOLS**.
3. On the **System Tools** menu, click **Build Information**.

The **Build Information** page consists of the following information: TeamForge version number, build number, operating system name and version number, instance (any packages of customizations that have been applied) and a list of installed RPMs.

### Get reports on site activity

You can keep track of various kinds of user activity, such as the number of user logins, source code commits, and tracker activity on your site.

- Monitor the total number of commits on your site by all users, in all projects, per day. You can track SCM commits to Subversion and CVS repositories, and to Git repositories as well, if you have the [TeamForge Git integration](#) set up.
- Check how many users log into your site in a given period. For example, to assess the impact of changes you've made to your site's look and feel (see [Customize anything on your site](#) on page 358), look for trends in the user login numbers for the week following the changes.
- Check the artifacts that are created and closed in a given period.

1. Go to **My Workspace > Admin**.

2. Click **REPORTS** from the **Projects** menu and select one of the available reports.

3. Click a **Zoom** option to see a week, a month, a quarter, 6 months, or a year of activity.

The **Max** option shows all the data that's ever been collected.

4. Use the **From** and **To** date fields to zero in on the exact period of activity you are interested in.

5. Roll your cursor over any data point on the graph to see the exact numbers that determine that point's position.

6. Click **Grid View** to see the data as a table.

You can toggle between chart and grid view as needed.

### Keep your site secure

As site administrator, you are responsible for making sure passwords are effective and hackers can't get in.

#### Set password requirements


For security reasons, you may want to enforce a password policy that requires certain password characteristics.

You can enforce the following password requirements:

- Minimum number of characters
- Maximum number of characters
- At least one number
- At least one non-alphanumeric character
- Mixed upper and lower-case characters
- Allowed invalid login attempts
- Disallow old passwords
- Use of dictionary words
- Default password expiry
- Disable/delete user accounts
- In the `site-options.conf` file, set the following variables.


 **Note:** Default values are inside square brackets.

- `MINIMUM_PASSWORD_LENGTH=6` | [6]
- `MAX_PASSWORD_LENGTH=12` | [256]
- `PASSWORD_REQUIRES_NUMBER=true` | [true]
- `PASSWORD_REQUIRES_NON_ALPHANUM=true` | [true]
- `PASSWORD_REQUIRES_MIXED_CASE=true` | [true]
- `REQUIRE_PASSWORD_SECURITY=true` | [true]
- `LOGIN_ATTEMPT_LOCK=3` | [3]

 **Tip:** This option controls locking out the user account after "n" invalid login attempts. Set this to zero or a negative number to lock the user account when the user provides an incorrect password for

the first time. Set this to a positive number, say "2", to allow the user two wrong password attempts. The user account would be locked at the "x+1" (here, third) attempt. When a user's account is locked, either an administrator must unlock it or the user can use the "Forgot Your Password?" link to reset the password.

- `PASSWORD_HISTORY_AGE=5` | [5]

 **Tip:** The maximum allowed value of `PASSWORD_HISTORY_AGE` token is 10. This option disallows the previous "n" passwords, while setting a password. However, if this option is set to zero, a negative number or it is left empty, the user can use any previous password. The password being set must satisfy the existing password policy each time.

- `ALLOW_PASSWORD_DICTIONARY_WORD=true`

 **Note:** You must set the `REQUIRE_PASSWORD_SECURITY` token to true in the `site-options.conf` file, for `LOGIN_ATTEMPT_LOCK`, `PASSWORD_HISTORY_AGE` and `ALLOW_PASSWORD_DICTIONARY_WORD` security setting to take effect.

### Prevent cross-site scripting (XSS)

By using the upload document function, an attacker could potentially upload an HTML page to TeamForge that contains active code, such as JavaScript. This active code would then be executed by clients' browsers when they view the page.

To prevent an attack of this sort, you can specify whether or not HTML code is displayed in CollabNet TeamForge. This flag applies to all documents, tracker, task, and forum attachments, and files in the file release system.

1. Set the `SAFE_DOWNLOAD_MODE` token according to your requirements. For more information, see [SAFE\\_DOWNLOAD\\_MODE](#) on page 769.
2. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```


 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.


### Support hardware provisioning

You can provide project managers with a TeamForge Lab Management site where project members can remotely build and test their code.


 **Tip:** For more information about the new features in TeamForge Lab Management 2.2, see [this introductory video](#). (You'll need a free open.collab.net membership to watch it.)

### Set up a new TeamForge Lab Management project

To start using TeamForge Lab Management project, create or select a TeamForge (or CollabNet Enterprise Edition) project to link to.


 **Note:** You can use Lab Management with either CollabNet TeamForge or CollabNet Enterprise Edition. These instructions assume you are using TeamForge. For information about working with CollabNetEnterprise Edition projects, see the [help for CEE project managers](#).

1. Log into TeamForge Lab Management as an administrator.
2. Click the **Administration** tab.
3. On the **Administration** menu, click **Manage Projects**.
4. Select the **Add New Projects** tab.
5. If Lab Management is connected to a TeamForge server, the **Name** field displays a drop-down list of projects instead of requiring a text entry. Select the project you want to add to Lab Management.
6. In the **Summary** field, write a brief summary of the project.

 **Note:** The TeamForge Lab Management project summary may be different from the CEE or SFEE project summary.

7. Click **Add New** button.

Your project is added. You can now pass the project on to the Project Admin of the project to begin administrating and configuring.


 **Note:** If your project is an SFEE project, you will see that TeamForge Lab Management has added some Roles to your project.

### Move hosts from one project to another

When a project has too many machines or needs more, you may wish to reallocate hosts between projects.

You must be a domain administrator to move hosts between projects.


1. Notify the project administrator of the project you are taking hosts from, so their project members won't lose important work which they may not have saved on that system.
2. Go to the **Project Home** page for the project which contains the host you are moving.
3. Select the **Hosts** tab to bring up the hosts list and click the host you want to re-allocate.
4. Click the **Admin** tab for the host.
5. From the **Project** drop-down box in the **Allocation Details** section, select the project that you want to move to.

 **Note:** If you do not see a drop-down box here, it means:

- The host is a virtual guest. Virtual guests must remain in the same project as their parents.
- The host contains virtual guests that are not all allocated to the same user. From the web interface, you can only change the project of a virtual host if the parent and all its guests are allocated to the same user.
- The host is a cloud source.


6. Click **Update** to move the host (and its guests, if any).
7. If you are concerned about the security or contents of the machine that has been moved to the new project, rebuild the machine as soon as possible.

This is the only way to ensure that only the exact software and users that you specified will be on the machine.

 **Note:** For a more flexible way of moving hosts between users and projects, you can use a TeamForge Lab Management API client to perform the move.

### Forward a port

To make a service accessible to users, sometimes you may need to forward a port on a Lab Management box to an external-facing port.

 **Note:** Port forwarding is only needed if Lab Management has been setup in an isolated deployment mode. If you can access the nodes directly, without tunneling through the manager, then you do not need to worry about setting up port forwards.

1. In the left-hand column of the **Administration** screen, click **Port Forwarding**.
2. Select a host and specify the type of traffic you want to forward.

#### Host

The TeamForge Lab Management host that you want to access the network service from. For example, `cu014.lmdomain.yourdomain.net`

#### Host Port


The port on the Host -- the hostname you selected above -- that you want to access the network service on. For example, if you would like to forward a web server, you would select Host Port 80.

#### External Port

The port on the TeamForge Lab Management Manager host that you would like to forward the traffic to. For example, if you would like




to forward the above web server on Host `cu014.lmdomain.yourdomain.net` and Host Port 80 to port 8080 on the TeamForge Lab Management Manager host, you would enter 8080 in this field.

 **Note:** Only certain ports may be selected for the External Port. The list of ports is given at the top of the screen.

### Protocol

The protocol or protocols to support when forwarding traffic. Choose one of these:

- TCP
- UDP
- Both TCP and UDP

 **Note:** If you're unsure, ask the requestor for clarification. Most network services will only require TCP.


### Description

A textual description of the service that is being forwarded. All users in the project will see this description when they view the Summary page for the Host.

 **Note:** The Description is mandatory.

### 3. Click **Add New** to add your new forwarding.


The new forwarding is active immediately. To test your new forwarding, try accessing the external port on the TeamForge Lab Management Manager from your desktop.

 **Tip:** If the network service requires special client software, especially if the forwarding is using the UDP protocol, you may not be able to easily test it. In that case, contact the requestor and ask them to verify that the forwarding works properly.

### Stop port forwarding

When you don't want to forward a port, you can delete or suspend forwarding for that port.


1. In the left-hand column of the **Administration** screen, click **Port Forwarding**.
2. If you want to remove the port forwarding permanently, click **Delete** for the forwarding you wish you delete. Forwarding for this port is immediately disabled and deleted forever.

 **Note:** You can always create a new forwarding with the same properties in the future.


3. If you think you might need this forwarding again soon, clear **Enable** and click **Update** for the forwarding you wish to disable. Enable the forwarding again by selecting **Enable** again and click **Update** again.

### Set up host URL mapping

If you have real users accessing services on a box you have set up virtually with TeamForge Lab Management, consider giving them a human-friendly URL to go to. You can do this by mapping an arbitrary URL to the host's real URL.

 **Note:** Only TeamForge Lab Management domain administrators can administer host URL mappings.

1. On the **Administration** tab, click **Host URL Mapping** in the navigation bar.
2. Provide the following information about the service you wish to map:


<b>Host</b>	The TeamForge Lab Management host that you want to access the network service from. For example, <code>cu014.lmdomain.yourdomain.net</code>
<b>Host Port</b>	The port on the Host -- the hostname you selected above -- that you want to access the web service on. For example, if you would like to map a web server running on port 8080, you would select Host Port 8080.
<b>Host URL</b>	The URL path on the Host that you want to map to. For example, if you want to map the URL <code>http://cu014.lmdomain.yourdomain.net:8080/build/</code> , you would enter <code>/build</code> for the Host URL.
<b>Mapped URL</b>	<p>The location (URL path) on the TeamForge Lab Management manager to map the Host URL to. The location will be under the directory <code>/apps</code>. For example, if you enter <code>mybuild</code> for Mapped URL in our example, you will then be able to reach <code>http://cu014.lmdomain.yourdomain.net:8080/build/</code> as <code>http://mgr.lmdomain.yourdomain.net/apps/mybuild/</code>.</p> <p>Valid characters for Mapped URL's are: Letters, numbers, spaces, underscores, and hyphens.</p>
<b>Use SSL</b>	Whether or not to force requests for the mapped URL to use SSL. Mapped URL's will always be accessible via SSL, even if this box is not checked. But, if you check this box, it will be impossible to access this URL without SSL. In general, unless there is a very good reason for turning the <code>Use SSL</code> option off, we recommend keeping it on.
<b>Dynamic Rewriting</b>	Set to either <b>Safer</b> or <b>More Aggressive</b> . Using the <b>Safer</b> option, only HTML attributes which are known to contain URL's are dynamically rewritten. But in some cases, especially with applications that use embedded Javascript code, this will not be sufficient and some parts of your application will not work properly. In this case, try the <b>More Aggressive</b> option. This option still attempts to be careful about what it rewrites, but looks into the HTML more deeply and will rewrite a wider variety of HTML tags and attributes.
<b>Description</b>	<p>A textual description of the web service that is being mapped. All users in the project will see this description when they view the Summary page for the Host.</p> <p> <b>Note:</b> The Description is mandatory.</p>

### 3. Click **Add New** to add your new URL mapping.

The new URL mapping is active immediately. To test it, try accessing the resulting URL from your desktop.

## Stop URL mapping


When you don't need a URL mapped, it's a good idea to remove the mapping.

 **Note:** Only TeamForge Lab Management domain administrators can administer host URL mappings.

1. On the **Administration** tab, click **Host URL Mapping** in the navigation bar.
2. Click **Delete** for the URL mapping you want to stop.

The URL mapping is immediately disabled and deleted forever.

 **Note:** You can always create a new URL mapping with the same properties in the future.

 **Tip:** To temporarily disable a URL mapping, clear **Enable** and click **Update** for the URL mapping you wish to disable. To enable the URL mapping again, select **Enable** and click **Update** again.

## Delete a project

To conserve resources on your TeamForge Lab Management site, it's a good idea to delete projects that are not in use.

Only TeamForge Lab Management Domain Admins can delete and undelete projects. When a project has been deleted, only TeamForge Lab Management Domain Admins can access it, and then only to undelete it.

Projects that contain hosts cannot be deleted.

1. Check that there are no hosts allocated to this project.  
You cannot delete the project if there are hosts currently allocated to it.
2. On the **Projects** tab, find the project you want to delete, and select **Edit** in the Actions column on the right for that project."
3. On the **Edit Project** page, click **Delete**, then **Confirm Delete**.

The project is deleted from the system.

## Undelete a project

To bring a project back into your TeamForge Lab Management site, you can undelete the project.

Only TeamForge Lab Management Domain Admins can delete and undelete projects. When a project has been deleted, only TeamForge Lab Management Domain Admins can access it, and then only to undelete it.

1. On the **Administration** tab, click **Manage Projects** in the navigation bar.  
In the list of all the projects in the system, deleted projects are in italics and are marked "Currently deleted."
2. Click the project you want to undelete.
3. On the **Edit Project** page for the deleted project, click **Undelete** button, then **Confirm Undelete**.


The project is restored to being a fully active TeamForge Lab Management project.

## Expand your domain with virtual guests

By using virtual guests, you can provide your users with many more machines than you actually have.

### *Create a virtual guest*

To add a virtual guest, go to the configuration page for the host where the guest will live.

 **Note:** You can only create virtual guests on machines that are allocated to you, are not guests themselves, and are not marked immutable.

1. On the **Manage Hosts** page, (or any other page where the host is visible), click the down arrow next to the host where your virtual guest will live and select **Manage Virtual Guests**.
2. On the host's **Manage Virtual Guests** page, specify the details for your new virtual guest under **Add New Virtual Guest**.

### *Migrate a virtual guest*


When a virtual guest consumes too many system resources, or you wish to rebuild or perform other more drastic maintenance on the virtual guest's current virtual host's use, migrate the guest to a different host.

Before moving a virtual guest, make sure the following is true:

- You are a TeamForge Lab Management domain administrator.


- The current virtual host and guest, and the desired new virtual host, are all in the same project.
- The target virtual host has enough RAM, disk, and CPU to accommodate the virtual guest you wish to migrate to it. (TeamForge Lab Management displays a drop-down list of hosts that fulfill all these criteria.)
- The virtual guest being moved and the target virtual host are on the same subnet.

1. Go to the **Host|Admin** page for the virtual guest that you wish to migrate.
2. In the **Move to Virtual Host** drop-down box, select the host you wish to migrate this virtual guest to, and click **Update Allocation Details**.

 **Note:** If you do not see this drop-down box, but instead see the text "No hosts are eligible for re-parenting," there are no other virtual hosts in this project that could accommodate this virtual guest. You must delete or reduce the hardware resources assigned to one or more existing guests.


3. Click **Confirm Update**.

The virtual guest immediately shuts down and begins migrating. The owner of the virtual guest (not necessarily the user that requested the migration) receives an e-mail when the migration is complete.

 **Note:** There is no exact time frame or estimate for how long a guest migration takes. A typical "base" system without much data (e.g., under 10GB) usually migrates in 5-15 minutes. The time required is mostly dependent on the size of the virtual machine, coupled with the load on the source and destination virtual hosts.

### *Delete a virtual guest*

You delete a virtual guest the same way you delete a physical host.

 **Important:** Before you delete your guest, be aware of one very important difference: As soon as you delete a virtual guest, the virtual guest system and all its data are wiped out forever! This is in contrast to deleting a physical host, in which case the machine gets removed from TeamForge Lab Management, but the contents of its hard drive are not wiped out, and could still be rescued by an administrator with physical access to the system.

1. On the **Host|Admin** screen, click **Delete Host**.
2. Click **Confirm Delete**.

The virtual guest is permanently deleted.

### **Manage profiles**

Like hosts and projects, profiles must be added at the domain level before they can be used by individual projects. Likewise, they can be administratively disabled from use at the domain level.

#### *Add a profile to the Lab Management domain*


In Lab Management, a "profile" is a reproducible collection of software packages that comprise a runnable operating system and accompanying set of applications. Every system inside the Lab Management environment has an operating system profile installed on it.

There are three ways to add a new profile to Lab Management :

- Directly add it to the profile repository using a Subversion client, then import it into Lab Management through the **Add Profile** page.
- Create the Profile XML file directly, and upload the file through the **Add Profile** page.
- Make a copy of an existing profile on that profile's Admin page, and edit it to suit your needs.

Among the three, the third is the easiest, and used most often by our own profile creators.

1. Log into TeamForge Lab Management as a user with TeamForge Lab Management Domain Admin privileges.
2. Click the **Administration** tab.
3. On the **Administration** menu, click **Manage Profiles** and then the **Add New Profile** link.
4. Enter the following information.

 **Note:** All fields are mandatory.

#### **Profile name**

The name of the profile. Must already exist in the site's Subversion repository before it can be added into

TeamForge Lab Management. This name can never be changed, so be sure to get the profile named properly when you initially create it.

You can also upload a profile from your local system into Lab Management's Subversion repository. The profile must be in valid Lab Management profile XML format.

<b>Summary</b>	A brief summary of the profile. The Summary can be changed later on using the Profile Admin page.
<b>Project</b>	What project the profile belongs to. If you leave it blank, the profile will belong to the domain, and either be accessible to <i>all</i> projects (if the profile is public) or <i>no</i> projects (if the profile is private). This property can be changed later on using the Profile Admin page.
<b>Public</b>	Whether the profile is public (can be used by any project) or private (can only be used by the project which owns the profile). This property can be changed later on using the Profile Admin page.
<b>Prebuilt</b>	Whether or not the profile is a prebuilt image file. <i>This profile property cannot be altered once the profile is created.</i> If your profile is an image file, this will be Yes. If your profile installs via a network install method like Kickstart(tm) or Jumpstart(tm), this will be No.
<b>Size (in GB)</b>	The minimum disk size required to build and run this profile. The minimum disk size should be at least 10 GB.

5. Click **Add New** to add the new profile.

#### *Modify the XML definition for a profile*

Rather than using a Subversion client to check out and edit a profile, you can edit the profile XML definition in a text box in the Lab Management user interface. This allows you the greatest flexibility in changing parts of the profile.

1. Log into TeamForge Lab Management as a user with TeamForge Lab Management Domain Admin privileges.
2. Click the **Administration** tab.
3. On the **Administration** menu, click **Manage Profiles**.
4. Right-click a profile and select **Admin** from its context menu.
5. In the Modify Profile: Profile XML Definition section, click **View/Edit XML Definition of Profile**.  
This brings up the content of the profile XML definition in an editable text box.
6. To edit the content, we recommend that you copy it to a text editor, edit and save it. Then copy it back into the text area.
7. To double-check your changes your changes before checking in the new version, click **Show Diff**.
8. When you're done, click **Check In New Version**.

If, by some chance, you made a mistake and the profile is invalid, don't worry -- since profiles are versioned, you can always mark that version as unbuildable.

#### *Stop using a profile*

In Lab Management, profiles can be permanently deleted or disabled.

#### Delete a profile


If a profile in Lab Management has never been used, it can be deleted. This can be helpful, for example, if you wish to clean up some profiles that were created with incorrect names.

1. Log into TeamForge Lab Management with TeamForge Lab Management Domain Admin privileges.
2. Click the **Administration** tab.
3. On the Administration menu, click **Manage Profiles**.
4. In the list of all profiles currently in the system, click the profile you wish to disable.
5. Click the **Admin** tab.

#### Disable a profile

When you disable a profile in Lab Management, it will prevent the users from building or rebuilding a host with that profile. It is possible to make disabled Lab Management profiles accessible again.

A disabled profile gets removed from all projects that currently allow the profile. If you disable a profile and then enable it later, project managers who previously allowed that profile will have to go back into their project's Profile Library and add the profile again.

 **Important:** Use caution before disabling a profile. Contact the owners of the systems using that profile, and owners of projects that include that profile, before disabling the profile.

1. Log into TeamForge Lab Management with TeamForge Lab Management Domain Admin privileges.
2. Click the **Administration** tab.
3. On the Administration menu, click **Manage Profiles**.
4. In the list of all profiles currently in the system, click the profile you wish to disable.
5. Click the **Admin** tab.
6. Click **Disable Profile**.

The profile is disabled and removed from all projects that currently allow it.

#### *Enable a profile*

To build a machine with a profile that has been disabled, you must re-enable the profile.

1. Log into TeamForge Lab Management with TeamForge Lab Management Domain Admin privileges.
2. Click the **Administration** tab.
3. On the **Administration** menu, click **Manage Profiles**.
4. In the list of all profiles currently in the system, click the profile you wish to re-enable.
5. Click the **Admin** tab.
6. On the **Admin** page, click **Enable Profile** to re-enable the profile.

#### **Administer clouds**

As an administrator, you can set up and manage the clouds in your Lab Management site.

#### *Add a new cloud*

When you pool the physical resources into a cloud, they can be partitioned in to different virtual machines and used in any project on the domain that needs them.

1. Select **Administration > Clouds**
2. In the **Manage Clouds** page, click **New Cloud**.
3. Enter a name and brief summary for the cloud.
4. Enter any additional details in the **Description**.
5. Select an owning project for the cloud.

The owning project allows the physical source systems in the cloud to be from that project. If you don't specify an owning project, the sources for the cloud can only be systems that are unallocated and don't belong to any project.

6. Specify whether this is a public or private cloud.
  - A public cloud appears in any project on the site that allows clouds, and is available for other project administrators to include in their list.
  - A private cloud appears only in the project that owns it, so that only the users in that project can allocate systems from it.
7. Specify whether hosts allocated from this cloud should be deleted when they are deallocated.

- Specify the maximum time (in days) for which a host can be allocated from this cloud.

The default value is 0 and indicates that there is no limit.

- Click **New Cloud**.

When your cloud gets created, you can proceed to add hosts to it.

#### *Add a cloud host*

When you've created a cloud, you can add hosts from physical sources or Amazon EC2 sources.

- On the **Manage Clouds** page, click a cloud name and select **Admin**.
- Add a cloud source using either of these options:
  - To add a physical host source, select one from the drop-down and click **Add Host**.  
Physical hosts can be partitioned into virtual machines that can be allocated into different projects on the site.
  - To add an Amazon EC2 source, click **Add EC2 Source**.

 **Note:** You will need information from your *Amazon Web Services* account to complete this task.

- Add a description for the source.
  - Specify the maximum number of hosts that may exist at any given time from the set of EC2 credentials specified on this page.
  - Fill in the EC2 credentials from the information in your Amazon Web Services account. See *Set up Amazon EC2 credentials* on page 376 for details on how to get this information - you'll need to provide the API Key, API Secret, and the Keypair Name and Keypair Value (your RSA private key).
- Click **Save**.

When you add an EC2 source, make sure that one or both EC2 profiles - Amazon\_EC2\_FC4\_base\_i386, Amazon\_EC2\_FC8\_base\_i386 - are enabled for your Lab Management projects.

#### *Define a host type*

To add a host type for a cloud, specify its attributes such as CPU, disk size and memory, and price for usage in an XML configuration file.

- On the cloud's **Admin** page, click **Modify Cloud**.
- Edit the XML file.

Here's an example:

```
<host_type value="Sun">
  <contains>
    <size value="extra_small">
      <price base="0.1" />
      <provides>
        <memsize value="256" />
        <disksize value="10" />
        <ncpu value="1" />
      </provides>
    </size>
    <size value="small">
      <price base="0.2" />
      <provides>
        <memsize value="512" />
        <disksize value="10" />
        <ncpu value="1" />
      </provides>
    </size>
    <size value="medium">
      <price base="0.3" />
      <provides>
        <memsize value="1024" />
        <disksize value="20" />
        <ncpu value="1" />
      </provides>
    </size>
  </contains>
</host_type>
```

```


        </provides>
    </size>
    <size value="large">
        <price base="0.4" />
        <provides>
            <memsize value="1024" />
            <disksize value="40" />
            <ncpu value="1" />
        </provides>
    </size>
    <size value="extra_large">
        <price base="0.5" />
        <provides>
            <memsize value="2048" />
            <disksize value="80" />
            <ncpu value="1" />
        </provides>
    </size>
</contains>
<requires category="cpuarch" op="in">
    sun4u
    sun4v
</requires>
<requires category="type" op="eq">
    physical
</requires>
</host_type>

```

### 3. Click **Update**.

#### *Delete a cloud*

When a cloud, and its hosts, are no longer in use, you can delete it.

 **Note:** You can delete a cloud only when all of its physical sources have been removed, and there are no hosts that were created from the cloud.


1. On the **Clouds** page, click on a cloud name and select **Admin**.
2. Click **Delete Cloud**.

#### *Set up Amazon EC2 credentials*

To be able to create an EC2 cloud source in Lab Management, you need to set up an Amazon Web Services account and supply the EC2 credentials to Lab Management.


To get started, you need to create an *Amazon Web Services* (AWS) account (or sign in with an existing Amazon account). You will be required to provide a credit card as part of sign up; this card will be charged for all your running EC2 instances. You will also be required to provide a callback number to which you have ready access. Amazon will generate an automated callback to this number to complete registration, and ask you to enter a generated PIN into your phone. Finally you'll receive a verification email letting you know that your account is set up.

1. Go to <https://console.aws.amazon.com/ec2/> to view the AWS console.

 **Note:** You should default to the US East region. This is the default that Lab Management assumes for all AMIs, instances, and key pairs.

2. Get the API credentials you'll give Lab Management to use your AWS account.
  - a) In **My Account > Security Credentials**, scroll down to your Access Keys, and get the **Access Key ID** (known as API Key in Lab Management).
  - b) Show the **Secret Access Key** (known as API Secret in Lab Management).
3. Generate or upload a **Key Pair**.

This is a named public/private key pair that will be used to access instances as root.

 **Note:** This key pair is associated with this AWS account, and therefore the credit card and the user that creates a cloud in Lab Management - take the appropriate precautions!



Amazon stores only the public key for passwordless SSH access to the instances. If you generate a key pair, you'll download the private key, which you'll be responsible for keeping secure. If you upload your own public key, you're still responsible for keeping your private key secure. For Lab Management-controlled EC2 instances, this key is used for passwordless root SSH access to the instance.

When you *create an EC2 cloud source in Lab Management*, you'll need to provide the API Key and API Secret viewed above, as well as the Keypair Name and Keypair Value (your RSA private key) generated above.

## Reference information about CollabNet TeamForge 17.4

---

This is some information about the scripts, configuration files and log files you can use to work with CollabNet TeamForge.


### Reference information about your CollabNet TeamForge site

Here is some detailed information to help you make the most of TeamForge.

#### Wiki syntax


You can use special wiki text markup to do a wide variety of cool and useful things on TeamForge wiki pages.

For a quick introduction, see the [documentation for JSPwiki](#), which is the wiki engine that drives the TeamForge wiki tool.

 **Note:** These tools are for use only when you are editing a wiki page in text mode. If you try to use them in WYSIWYG mode, they are displayed just the way you typed them in, which is not what you want.

#### Text effects on wiki pages

Wiki markup is great for making text look the way you need it to look.


 **Note:** These tools are for use only when you are editing a wiki page in text mode. If you try to use them in WYSIWYG mode, they are displayed just the way you typed them in, which is not what you want.

Syntax	Effect	Details
----	Creates a horizontal rule.	
\\	Creates a line break.	
!!!text	Creates a level 1 (large) header.	
!!text	Creates a level 2 (medium) header.	
!text	Creates a level 3 (small) header.	
' 'text'	Creates italic text. (That's two single quotes on each side.)	
__text__	Creates bold text. (That's two underscores on each side.)	
{{text}}	Creates monospaced text.	
*text	Creates a bulleted list item.	
#text	Creates a numbered list item.	
;term:ex	Creates a definition for the word "term" with the explanation "ex."	

Syntax	Effect	Details
<code>{{text}}</code>	Creates pre-formatted text.	
<code>%( &lt;css-style&gt; )&lt;your text&gt;%</code>	Defines a CSS style command.	<code>%( font-size: 150%; color: red; ) Hello, world!%</code>
Blank line	Starts a new paragraph.	

### Bring TeamForge data into wiki pages


Use this markup format to bring in data from elsewhere on your TeamForge site.

-  **Note:** These are for use only when you are editing a wiki page in text mode. If you try to use them in WYSIWYG mode, they are displayed just the way you typed them in, which is not what you want.

Syntax	Effect	Details
<pre>[{INSERT  ExcelToHTMLPlugin  src='c:\somesheet.xls'}] or [{INSERT  ExcelToHTMLPlugin  border='1'  src='\\the_server  \somesheet.xls'}]</pre>	Reads a Microsoft Excel file and displays it as an HTML table.	Parameters: <ul style="list-style-type: none"> <li>src: URL / Attachment file name</li> <li>srcsheet: Sheet name</li> <li>height: height attribute for the html table</li> <li>width: width attribute for the html table</li> <li>border: border attribute for the html table</li> </ul> More at <a href="http://www.ecyrd.com/JSPWiki/wiki/ExcelToHTMLPlugin">http://www.ecyrd.com/JSPWiki/wiki/ExcelToHTMLPlugin</a>

### Text navigation tools for wiki pages

You can use wiki syntax to help readers get around.

-  **Note:** To use these tools, copy and paste the sample syntax into your Wiki page in "Plain Editor" mode, then customize it appropriately.


Syntax	Effect	Details
<code>%%insert-toc %%</code>	Creates a table of contents consisting of the header text on the page.	
<code>[link]</code>	Creates a link to a new Wiki page called "link."	If the link is a complete URL, a link to the URL is created. If the link points to a .gif, .jpg, or .png image, the image is rendered directly in the page.
<code>[title link]</code>	Creates a link to a new Wiki page called "link" with the text "title" displayed for the URL.	If the link is a complete URL, a link to the URL is created. If the link points to a .gif, .jpg, or .png image, the image is rendered directly on the page with "title" as ALT text.
<code>~TestText</code>	Disables link creation for a CamelCase word.	CamelCase words are two or more uppercase words with no spaces.


Syntax	Effect	Details
		By default, a CamelCase word automatically creates a link to a new Wiki page.
<code>[[link]]</code>	Creates the text "[link]."	
<pre>{IFramePlugin url='http:// open.collab.net/' width='100%' height='500' border='1' scrolling='yes' align='center'}}</pre>	Embeds an iframe into a wiki page.	<ul style="list-style-type: none"> <li>attachment: Attachment path, e.g. 'IFramePlugin.jar(info)'</li> <li>url: A URL, e.g 'http://www.google.com'</li> <li>align: Align the iFrame to left/center/right</li> <li>border: Whether there is a border or not</li> <li>width: Width of the iFrame</li> <li>height: Height of the iFrame</li> <li>marginwidth: Margin width of the iFrame</li> <li>marginheight: Margin height of the iFrame</li> <li>scrolling: Whether the iFrame can be scrolled or not</li> </ul> <p>See <a href="http://www.ecyrd.com/JSPWiki/wiki/iFramePlugin">http://www.ecyrd.com/JSPWiki/wiki/iFramePlugin</a></p>
Tab Completion <Keyword+Tab>	You can type a keyword and hit the Tab key (under the Tab Completion mode). The editor will fill in with a sample template for the specific markup represented by the keyword.	<ul style="list-style-type: none"> <li>link: Inserts a sample link.</li> <li>h1: Inserts level 1 heading sample.</li> <li>h2: Inserts level 2 heading sample.</li> <li>h3: Inserts level 3 heading sample.</li> <li>bold: Inserts a bold text sample.</li> <li>italic: Inserts an italics text sample.</li> <li>mono: Inserts a mono text sample.</li> <li>sup: Inserts a superscript sample.</li> <li>sub: Inserts a subscript sample.</li> <li>strike: Inserts a strike through text sample.</li> <li>br: Inserts a line break.</li> <li>hr: Inserts a horizontal line.</li> <li>pre: Inserts a pre-formatted text sample.</li> <li>code: Inserts a code block sample.</li> <li>dl: Inserts a definition list block sample.</li> <li>toc: Inserts the Table of Contents plugin syntax.</li> <li>tab: Inserts a sample tabbed section block syntax.</li> <li>table: Inserts a sample table syntax.</li> </ul>

Syntax	Effect	Details
		<ul style="list-style-type: none"> <li>img: Inserts a sample image plugin syntax.</li> <li>quote: Inserts a sample quoted text block.</li> <li>sign: Inserts the user's signature.</li> </ul>

### Attachments for wiki pages

You can use wiki markup to bring in information from external sources.

 **Note:** These tools are for use only when you are editing a wiki page in text mode. If you try to use them in WYSIWYG mode, they are displayed just the way you typed them in, which is not what you want.


Syntax	Effect	Details
<code>[WikiPageName/ attachmentName]</code>	Embeds an attachment in the page.	If the attachment is a .gif, .jpg, or .png image file, the attachment will be embedded in the page; otherwise, the name of the attachment will display as a downloadable link. After adding attachments, the exact syntax for including the current page's attachments is shown next to each attachment's name in the Attachments section of the Edit Wiki page. You can use the same syntax to embed attachments from other wiki pages in the same project.
<code>[[InsertAttachment page='WikiPage/ attachment']]]</code>	Inserts the contents of an attachment (text file) into a page.	If the attachment is a text file, the content of the text file is inserted into the page. For more information, <a href="#">click here</a> .   <b>Note:</b> <ul style="list-style-type: none"> <li>This markup is for inserting content of a text file.</li> <li>Files with the following extensions are allowed as attachment: .txt, .html, .xml, .cpp and .java.</li> <li>In addition to image files, inserting files with the following extensions is not supported by this markup: .doc, .xls, .pdf, .zip, .jar.</li> </ul>
<code>[[MediaPlayer src='fileName.wmv']]]</code>	Embeds a Windows Media Player or Quicktime Player on a wiki page.	<ul style="list-style-type: none"> <li>src: Media URL / Attachment file name</li> <li>playertype: "mediaplayer" / "quicktime"</li> <li>width, height: Dimension of the embedded media displayed</li> </ul>

Syntax	Effect	Details
		<ul style="list-style-type: none"> <li>• movieheight, moviewidth: Dimension of the display screen</li> <li>• caption: Caption to be displayed below the media player</li> <li>• control: Displays Control bar. mediaplayer: 1 (Show) / 0 (Hide); quicktime: true (Show) / false (Hide)</li> <li>• autostart: Play automatically. mediaplayer: 1 (Auto) / 0 (Manual, Click to play); quicktime: true (Auto) / false (Manual, Click to play)</li> <li>• autorewind: Automatically rewinds when play ends. mediaplayer: 1 (Auto Rewind) / 0 (Play once); quicktime: true (Auto Rewind) / false (Play once)</li> <li>• playcount: Number of times the movie will play. 0 represents always play.</li> </ul> <p>See <a href="http://www.ecyrd.com/JSPWiki/wiki/MediaPlayerPlugin">http://www.ecyrd.com/JSPWiki/wiki/MediaPlayerPlugin</a></p>
<pre> [{{INSERT  ExcelToHTMLPlugin  src='WikiPage  \somesheet.xls'}}] or [{{INSERT  ExcelToHTMLPlugin  border='1'  src='\\the_server  \somesheet.xls'}}] </pre>	<p>Reads a Microsoft Excel file and displays it as an HTML table.</p>	<p>Parameters:</p> <ul style="list-style-type: none"> <li>• src: URL / Attachment file name</li> <li>• srcsheet: Sheet name</li> <li>• height: height attribute for the html table</li> <li>• width: width attribute for the html table</li> </ul> <p>More at <a href="http://www.ecyrd.com/JSPWiki/wiki/ExcelToHTMLPlugin">http://www.ecyrd.com/JSPWiki/wiki/ExcelToHTMLPlugin</a></p>
<pre> [{{arnaud.Flash  src='yourAttachedFlash.swf </pre>	<p>Embeds a Flash Player on your wiki page.</p>	<ul style="list-style-type: none"> <li>• width='n'</li> <li>• height='n'</li> <li>• controls='true false'</li> <li>• play='true false'</li> <li>• loop='true false'</li> <li>• parameters='param1=value1, &amp;param2=value2'</li> </ul> <p>See <a href="http://www.ecyrd.com/JSPWiki/wiki/AFlashPlugin">http://www.ecyrd.com/JSPWiki/wiki/AFlashPlugin</a></p>
<pre> [{{IFramePlugin  url='http://  open.collab.net/'  width='100%'  height='500'  border='1'  scrolling='yes' </pre>	<p>Embeds an iframe into a wiki page.</p>	<ul style="list-style-type: none"> <li>• attachment: Attachment path, e.g. 'IFramePlugin.jar(info)'</li> <li>• url: A URL, e.g 'http://www.google.com'</li> <li>• align: Align the iFrame to left/center/right</li> </ul>

Syntax	Effect	Details
<code>align='center']}]</code>		<ul style="list-style-type: none"> <li>border: Whether there is a border or not</li> <li>width: Width of the iFrame</li> <li>height: Height of the iFrame</li> <li>marginwidth: Margin width of the iFrame</li> <li>marginheight: Margin height of the iFrame</li> <li>scrolling: Whether the iFrame can be scrolled or not</li> </ul> <p>See <a href="http://www.ecyrd.com/JSPWiki/wiki/iFramePlugin">http://www.ecyrd.com/JSPWiki/wiki/iFramePlugin</a></p>

### Tables on wiki pages

You can use wiki markup to organize information in tables on a wiki page.

 **Note:** These tools are for use only when you are editing a wiki page in text mode. If you try to use them in WYSIWYG mode, they are displayed just the way you typed them in, which is not what you want.

Syntax	Effect	Details
<pre> [{{Table &lt;table- parameters&gt;    Table Header Example    More...   Table Data Example   More... }} </pre>	Inserts a table on the wiki page. See this page for the markup for table elements.	<ul style="list-style-type: none"> <li>rowNumber: &lt;integer&gt; Row number starts counting at this value; default = 0 (used in conjunction with '#' syntax)</li> <li>style: &lt;css-style&gt; Add formatting to the table, e.g. style='border=2px solid black;'</li> <li>dataStyle: &lt;css-style&gt; Format all data cells (prefixed by single pipe signs ' ')</li> <li>headerStyle: &lt;css-style&gt; Format all header cells (prefixed by double pipe signs '  ')</li> <li>evenRowStyle: &lt;css-style&gt; Format the even rows, e.g. evenRowStyle='background: #ffff00;'</li> <li>oddRowStyle: &lt;css-style&gt; Format the odd rows, e.g. oddRowStyle='color: red;'</li> </ul>
<code>  head1  head2</code>	Creates a table column with header text "head1" in the first cell and "head2" in the second cell.	
<code> col1 col2</code>	Creates a table row containing the text "col1" in the first cell and "col2" in the second cell.	
<code> &lt;</code>	Collapses a cell with the previous cell so it spans multiple columns.	<a href="http://www.ecyrd.com/JSPWiki/wiki/TablePlugin">http://www.ecyrd.com/JSPWiki/wiki/TablePlugin</a>

Syntax	Effect	Details
<code>  &lt;</code>	Collapses a header cell with the previous header cell so it spans multiple columns.	<a href="http://www.ecyrd.com/JSPWiki/wiki/TablePlugin">http://www.ecyrd.com/JSPWiki/wiki/TablePlugin</a>
<code> ^</code>	Collapses a cell with the cell above so that it spans multiple rows.	<a href="http://www.ecyrd.com/JSPWiki/wiki/TablePlugin">http://www.ecyrd.com/JSPWiki/wiki/TablePlugin</a>
<code>  ^</code>	Collapses a header cell with the header cell above so that it spans multiple rows.	<a href="http://www.ecyrd.com/JSPWiki/wiki/TablePlugin">http://www.ecyrd.com/JSPWiki/wiki/TablePlugin</a>
<code>  ( &lt;css-style&gt; )</code>	Inside a table cell, adds specific formatting to a cell.	<a href="http://www.ecyrd.com/JSPWiki/wiki/TablePlugin">http://www.ecyrd.com/JSPWiki/wiki/TablePlugin</a>
<code>#</code>	Inside a table cell, displays the current row number with auto row numbering.	<a href="http://www.ecyrd.com/JSPWiki/wiki/TablePlugin">http://www.ecyrd.com/JSPWiki/wiki/TablePlugin</a>
<pre>[{INSERT   ExcelToHTMLPlugin   src='WikiPage   \somesheet.xls' }] or [{INSERT   ExcelToHTMLPlugin   border='1'   src='\\the_server   \somesheet.xls' }]</pre>	Reads a Microsoft Excel file that is attached to a wiki page, and displays it as an HTML table.	<p>Parameters:</p> <ul style="list-style-type: none"> <li>• <code>src</code>: If the Excel file is attached to the current wiki page, this is the attachment file name. If it is attached to some other wiki page, this is the URL of the attachment.</li> <li>• <code>sheetsheet</code>: Sheet name</li> <li>• <code>height</code>: height attribute for the html table</li> <li>• <code>width</code>: width attribute for the html table</li> <li>• <code>border</code>: border attribute for the html table</li> <li>• <code>cellpadding</code>: cellpadding attribute for the html table</li> <li>• <code>cellspacing</code>: cellspacing attribute for the html table</li> <li>• <code>background</code>: background attribute for the html table. Attachment file name is accepted as value</li> <li>• <code>backgroundcolor</code>: backgroundcolor attribute for the html table</li> <li>• <code>keepformat</code>: Formatting specified in the excel sheet is applied for the html table. Note: The complete formatting from excel sheet is not applied to the html table, for example, font, font size, etc are not applied and background color, foreground color, etc are applied. yes/no is accepted as value</li> <li>• <code>headercolor</code>: Foreground color for the header (eg: #dcdcdc)</li> </ul>

Syntax	Effect	Details
		<ul style="list-style-type: none"> <li>• headerbackgroundcolor: Background color for the header (eg: #adadad)</li> <li>• evenrowcolor: Foreground color for the even rows (eg: #adadad)</li> <li>• evenrowbackgroundcolor: Background color for the even rows (eg: #adadad)</li> <li>• oddrowcolor: Foreground color for the odd rows (eg: #adadad)</li> <li>• oddrowbackgroundcolor: Background color for the odd rows (eg: #adadad)</li> <li>• tableclass: Style class name for the HTML table</li> <li>• headerclass: Style class name for header</li> <li>• evenrowclass: Style class name for the even rows</li> <li>• oddrowclass: Style class name for the even rows</li> <li>• stylesheet: Stylesheet for the table. Attachment file name is accepted as value</li> </ul> <p>See <a href="http://www.ecyrd.com/JSPWiki/wiki/ExcelToHTMLPlugin">http://www.ecyrd.com/JSPWiki/wiki/ExcelToHTMLPlugin</a></p>
<code>[[TableOfContents ]]\</code>	Creates a table of contents consisting of the header text on the page.	

### Wiki plugins

You can use pre-defined wiki plugins to format your current page.

Plugin Name	Details
Section Headings Template	Inserts a page template that includes a table of contents and several section headings. Typically, this can be used in a new or blank page.
Sortable Table	Inserts a new table that can be sorted when you click on the column headers.
Zebra Table	Inserts a new table that has alternating background colors for each row.
Table of Contents Plugin	This plugin automatically generates table of contents that provides links to all the headings on your page.
Insert Page Plugin	This plugin will insert a copy of another page into the current page. You must specify the name of the page to insert.
Current Time Plugin	This plugin displays the current date and local time of the server when the page is viewed.



Plugin Name	Details
Insert Attachment	The insert attachment plugin allows you to insert the contents of an attachment into a page.
Media Player	This is an embedded player in your wiki page, supporting Windows Media Player and Apple QuickTime.
Insert Table	Additional table support including multi-line table editing, cell merging, and automatic row numbering.
Flash Player	An embedded flash player for the wiki page.
Insert Excel	Allows you to insert a Microsoft Excel file as a table.
Insert iFrame	Allows embedding attachments, external urls, and files (relative to the docbase).
Code to HTML	Allows source code syntax to be rendered as a HTML output and supports syntax from 130 different programming languages.
Index Plugin	Displays all the pages in wiki in a alphabetical order.
Recent Changes Plugin	Inserts the latest changes in order.
Referred Pages Plugin	Finds and lists all the pages that are referred to by the current page. The depth parameter allows to display a recursive tree of referred pages.
Referring Pages Plugin	Finds and lists all the pages that refer to the current page.
Undefined Pages Plugin	Lists all the pages that are referred to, but not yet created.
Unused Pages Plugin	Lists all the pages that are not currently referred to by any other page.
ExcelToHTMLPlugin	Provides a HTML view for the Excel spread sheet files.
PDFPlugin	Provides the PDF output for the files.
WikiContentToHTML Plugin	Exports a specific page to HTML.

### The Document Details page

Each document in TeamForge has a **Document Details** page where you can find information about the document.

### Contents

The **Document Details** page has these components:

#### Document Details

View details about the document such as the name, description, file type, and size.

#### Status

View the document status, including whether it is locked for editing.

#### Created By

Identify and contact the author of each document version.

#### Versions

View the version history of the document.

#### Associations

View any items associated with the document.

**Version Comment**

Explanation of changes from one version of the document to another.

**Monitor**

Begin or stop monitoring the document.

**Filters for automatically sorting emails from your TeamForge site**

X-headers are special instructions that are added to many email messages. Messages that come from your TeamForge site have X-headers describing the purpose of the email, the event that triggered it, and lots of other information.

You can use these X-headers to sort and filter emails sent by TeamForge.

<b>X-header</b>	<b>Example</b>	<b>Description</b>
X-TeamForge-Application	Tracker	Application, if applicable. Values include Documents, Wiki, Source Code, TeamForge, Tasks, Tracker, Discussion, File Releases.
X-TeamForge-Artifact-Monitoring	true	For artifacts, true if the user is monitoring the artifact directly. If value is false, you are monitoring the folder, not the artifact.
X-TeamForge-AssignedTo	admin	User assigned to artifact or task.
X-TeamForge-CommitPath	/path/to/my/branch	For Subversion commit emails: reports the greatest common path of all files in the commit. For example, if you modify the paths /a/b/file1 and /a/c/file2, the value of the X-TeamForge-CommitPath header is /a. If you modify /a/b/file1 and /file2, the value is / (the root path of the repository).
X-TeamForge-FolderId	tracker1005	If the object is a folder, the object's ID, or the ID of the object's container.
X-TeamForge-Forum	New user forum	For posts, this is the forum title.
X-TeamForge-PlanningFolder	/path/to/my/folder	The path to a planning folder.
X-TeamForge-ProjectId	proj1006	Project ID, if applicable.
X-TeamForge-ServerName	collab.net	Hostname portion of TeamForge URL.
X-TeamForge-Status	Open	Status of artifacts, tasks, documents, and document reviews. Values include: <ul style="list-style-type: none"> <li>• For artifacts: Open, Closed, Pending, or user-defined.</li> <li>• For tasks: Not Started, OK, Complete, Warning, or Alert.</li> <li>• For documents: Draft, Review, or Final.</li> </ul>

X-header	Example	Description
		<ul style="list-style-type: none"> <li>For document reviews: Open or Closed.</li> </ul>
X-TeamForge-Type	Artifact	The type of the object, if applicable. Values include Tracker, Artifact, Document Folder, Document, Document Version, Wiki Page, Wiki Page Version, Task Folder, Task, FRS Package, FRS Release, Discussion Forum, Discussion Topic, SCM Repository, SCM Commit.

## Reference information for TeamForge project administrators

Here are some sample files and other resources to help you figure out what is what in this CollabNet product.

### The Project Dashboard page

The **Project Dashboard** offers a centralized view into all development projects managed in TeamForge .

### Overview

The **Project Dashboard** provides managers with an at-a-glance overview of the status of each of their projects. It provides summary information on the number and status of the tasks and tracker artifacts in each project, and calculates project overrun and underrun statistics.

The **Project Dashboard** also provides overview information such as project start and end dates and project ranking.

You can see the **Project Dashboard** if you have both the View Tracker and View Task permissions for one or more projects. Only those projects for which you have both the View Tracker and View Task permissions appear on your **Project Dashboard**.

### Getting there

In the TeamForge navigation bar, click **My page > Dashboard**.

You can view the **Project Dashboard** if you have both the View Tracker and View Task permissions for one or more projects. Only those projects for which you have both the View Tracker and View Task permissions are displayed on your **Project Dashboard**.

### Contents

For each project, the **Project Dashboard** displays the following information:

- **Project Activity ranking** - The activity of the project in relation to all other TeamForge projects.
- **Start Date or End Date** - The start and end date of the project, based on the start and end dates of all project tasks.
- **Task Status** - The status of the project, based on the "rolled-up" status of all project tasks and task folders. You can configure the "roll up" criteria for each project from the project's **Task Manager Settings** page.
- **Status History** - The history of the project's "rolled-up" status color. These figures are calculated in real time, but do not calculate time that the project's status was **Not Started** or **Completed**.

Click the status bar to go to the project's **Task Summary** page.

- **Task and Tracker Effort** - The project's current overrun or underrun, based on the difference between estimated and actual effort spent on project tasks and tracker artifacts.

Only completed and closed tasks and tracker artifacts, with values in the estimated and actual effort fields, contribute to the overrun or underrun calculations.

- **Tracker Status** - The number of open tracker artifacts in the project, per priority value. The number of open tracker artifacts is indicated in parentheses.

Click the status bar to go to the project's **Tracker Summary** page.

### User permissions in CEE and CTF 17.4 projects

Each permission a user can have in a CollabNet Enterprise Edition project has an equivalent permission in a CollabNet TeamForge project, with a few exceptions.

#### TeamForge 17.4 Permissions for Audit log

Permission Mapping Details

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/Gain Details	Comments
Audit Log - View	None	User loses permission	The user loses the ability in absence of "Domain - Administer" permission.	Conditional mapping. This is handled in the mapping for "Domain - Administrator" permissions.

#### TeamForge 17.4 Permissions for Announcements

The news application in TeamForge does not have its own permissions. Users' ability to view project news depends on their level of access to the project itself. The permissions for the CollabNet Enterprise Edition Announcements tool are mapped to the closest possible project permissions in the TeamForge.

Permission Mapping Details

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/Gain Details	Comments
Project News - Approve	None	User loses the permission	The user will not be able to create the news if the role does not have the "Project Admin" permissions	These tasks are project administrator only tasks in TeamForge. It is not correct to grant the "Project Admin" permissions to map this permission as that will grant much higher privilege. Hence, this permission is dropped.
Project News - Delete	None	User loses the permission	The user will not be able to delete the news if the role does not have the "Project Admin" permissions	These tasks are project administrator only tasks in TeamForge. It is not correct to grant the "Project Admin" permissions to map this permission as that will grant much higher privilege. Hence,

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/Gain Details	Comments
				this permission is dropped.
Project News - Edit	None	User loses the permission	The user will not be able to edit the news if the role does not have the "Project Admin" permissions	These tasks are project administrator only tasks in TeamForge. It is not correct to grant the "Project Admin" permissions to map this permission as that will grant much higher privilege. Hence, this permission is dropped.
Project News - Suggest	None	Feature is not available in TeamForge		TeamForge 17.4 doesn't have any such permission.
Project News - View	None	CollabNet Enterprise Edition permission is retained		In TeamForge 17.4, if the user has a role, then the "Project - view" permission is automatically granted.

### TeamForge 17.4 Permissions for Discussions

Permissions are mapped to TeamForge's Discussions application.

#### Permission Mapping Details

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/Gain Details	Comments
Discussions - Discussion Add	Forums Admin	User loses permission	The user loses the permission if the role does not have "Discussion Edit" permission and "Discussion - Use Private" in specific cases.	Conditional Mapping. For Global/Site-wide role, projects having private discussion, and parent projects having child projects with private forum: Grant only if the role has "Discussion - Discussion Edit" and "Discussions - Discussion Use Private" permissions also. If a project-level role and the project/child projects do not have any

<b>CollabNet Enterprise Edition Permission</b>	<b>TeamForge 17.4 Permission</b>	<b>Change in Behavior</b>	<b>Functionality Loss/ Gain Details</b>	<b>Comments</b>
				private forum, do not consider "Use Private".
Discussions - Discussion Delete	Delete/View	CollabNet Enterprise Edition permission is retained		All public forums are mapped. Additionally, all private forums are mapped if the user has "Discussions - Discussion Use Private" permission.
Discussions - Discussion Edit	Forums admin	User loses permission	The user loses the permission if the role does not have "Discussion Add" permission and "Discussion - Use Private" in specific cases.	Conditional Mapping. For Global/ Site-wide role, projects having private discussion, and parent projects having child projects with private forum: Grant only if the role has "Discussion - Discussion Add" and "Discussions - Discussion Use Private" permissions also. If a project-level role and the project/child projects do not have any private forum, do not consider "Use Private".
Discussions - Discussion Trusted	ALTERNATE MEANS	CollabNet Enterprise Edition permission is retained		The users having this permission are made the trusted users of all the forums in the project subject to "Discussion – Use Private" permission.
Discussions - Message Delete		This feature is not available in TeamForge 17.4		TeamForge 17.4 does not have any such feature.
Discussions - Message Edit		This feature is not available in TeamForge 17.4		TeamForge 17.4 does not have any such feature.
Discussions - Message Post	Post/View	CollabNet Enterprise Edition permission is retained.		All public forums are mapped. Additionally, all private forums are

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/ Gain Details	Comments
				mapped if the user has "Discussions - Discussion Use Private" permission.
Discussions - Discussion View	View only	CollabNet Enterprise Edition permission is retained.		All public forums are mapped. Additionally, all private forums are mapped if the user has "Discussions - Discussion Use Private" permission.
Discussions - Discussion Use Private		CollabNet Enterprise Edition permission is retained.		Other permissions are specified at the folder level. If the role has this permission, then permission can be given to all public and all private forums, otherwise, access is limited to only public forums.

### TeamForge 17.4 Permissions for Documents and Files

Permissions are mapped to TeamForge 17.4's Documents application.

#### Permission Mapping Details

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/ Gain Details	Comments
Project Document - Approve	Create/View	User loses permission	No suggest/approve feature in TeamForge 17.4.	Conditional Mapping. This permission is granted only if the user has Project Document – Suggest also.
Project Document - Delete	Delete/View	CollabNet Enterprise Edition permission is retained		This permission is applicable to all root document folders.
Project Document - Edit	Edit/View	CollabNet Enterprise Edition permission is retained		This permission is applicable to all root document folders.
Project Document - Edit Locked	Document Admin	User loses permission	The user is likely to lose the rights to edit a locked document in certain cases.	Conditional Mapping. This permission is granted only if the user has "Project Document - Approve", "Project

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/ Gain Details	Comments
				Document - Edit", "Project Document - Reserve" and "Project Document – Suggest" permissions also.
Project Document - Reserve		User loses permission	The user will lose the ability to reserve documents in absence of "Project Document - Edit" permission	
Project Document - Reserve for Others		User loses permission	Users without "Project Document – Edit" permission will lose the rights	
Project Document - Reserve Mine		User loses permission	Users without "Project Document – Edit" permission will lose the rights	
Project Document - Suggest		This feature is not available in TeamForge 17.4		TeamForge 17.4 does not have any such feature.
Project Document - View	View Only	CollabNet Enterprise Edition permission is retained		This permission is applicable to all root document folders.

### TeamForge 17.4 Permissions for Issue Tracker/Project Tracker

CollabNet Enterprise Edition uses the same set of permissions for Issue Tracker and Project Tracker. There are some permissions for individual tools too. The CollabNet Enterprise Edition project tracker permissions are appropriately transferred to Project Tracker (integrated application in TeamForge 17.4) or Issue Tracker, based on the tool that the project uses or the inherited tool permissions in child projects or the permissions that are a part of a global or site-wide role.

#### Permission Mapping Details

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/ Gain Details	Comments
Project Issue Tracking - Assignable	Pluggable: Assignable	CollabNet Enterprise Edition permission is retained		This permission is assigned to the user for Project Tracker if the project uses Project Tracker (or) if permissions are inherited to child projects using Project Tracker (or) if the role is a global/site-wide role.



<b>CollabNet Enterprise Edition Permission</b>	<b>TeamForge 17.4 Permission</b>	<b>Change in Behavior</b>	<b>Functionality Loss/Gain Details</b>	<b>Comments</b>
Project Issue Tracking - Assignable	TeamForge 17.4: conditionally dropped	User loses permission	The user loses permissions in the case of global/site-wide role, if the role does not have "Edit/View" permission	User becomes assignable through the "Edit/View" permission. There is no separate permission in TeamForge 17.4.
Project Issue Tracking - Add Comment	TeamForge 17.4: dropped	User loses permission	The user loses rights if the role does not have "Attach File", "Change" and "Change Status" permissions.	
Project Issue Tracking - Attach File	TeamForge 17.4: dropped	User loses permission	The user loses rights if the user does not have "Change", "Add Comment" and "Change Status" permissions	
Project Issue Tracking - Change	Pluggable: Edit Artifact	CollabNet Enterprise Edition permission is retained		This permission is assigned to the user for Project Tracker if the project uses Project Tracker (or) if permissions are inherited to child projects using Project Tracker (or) if the role is a global/site-wide role.
Project Issue Tracking - Change	TeamForge 17.4: Edit/View	CollabNet Enterprise Edition permission is retained	Granted only if the role also has "Attach File", "Add Comment" and "Change Status" permissions.	Conditional mapping.
Project Issue Tracking - Change Status	TeamForge 17.4: dropped	User loses permissions	The user loses rights if the user does not have "Change", "Add Comment", and "Attach File" permissions	
Project Issue Tracking - Configure	Pluggable: Configure project	CollabNet Enterprise Edition permission is retained		This permission is assigned to the user for Project Tracker if the project uses Project Tracker (or) if permissions are inherited to child

<b>CollabNet Enterprise Edition Permission</b>	<b>TeamForge 17.4 Permission</b>	<b>Change in Behavior</b>	<b>Functionality Loss/Gain Details</b>	<b>Comments</b>
				projects using Project Tracker (or) if the role is a global/site-wide role.
Project Issue Tracking - Configure	TeamForge 17.4: Tracker Admin	User loses permission	The user can lose the rights if the role does not have "Submit", "Change" and "Add Comment" permissions. Also, the user may gain "Attach", "Change status" permissions	Conditional mapping. Granted only if the role also has "Submit", "Change" and "Add Comment" permissions.
Project Issue Tracking - Move Issue	Pluggable: Move artifact	CollabNet Enterprise Edition permission is retained		This permission is assigned to the user for Project Tracker if the project uses Project Tracker (or) if permissions are inherited to child projects using Project Tracker (or) if the role is a global/site-wide role.
Project Issue Tracking - Move Issue	TeamForge 17.4: Delete/View	Permissions are gained	Granted only if the user gets "Submit/View" and "Edit/View" permissions.	Conditional mapping. User may get additional permissions to delete artifacts and also the additional permissions gained via the "Submit/View" and "Edit/View" permissions. Drop if Project role.
Project Issue Tracking - Query	Pluggable: Query	CollabNet Enterprise Edition permission is retained		This permission is assigned to the user if the role is a project role (or)if the project uses Project Tracker (or) if permissions are inherited to child projects using Project Tracker (or) if the role is a global/site-wide role.
Project Issue Tracking - Query	TeamForge 17.4: View Only	CollabNet Enterprise Edition permission is retained		There is no exact parity, users having "View" permission will be able to use

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/Gain Details	Comments
				Filter and Search. Also, Issue Tracker uses this permission for "View" access.
Project Issue Tracking - Submit	Pluggable: Enter artifact	CollabNet Enterprise Edition permission is retained		This permission is assigned to the user if the role is a project role (or)if the project uses Project Tracker (or) if permissions are inherited to child projects using Project Tracker (or) if the role is a global/site-wide role.
Project Issue Tracking - Submit	TeamForge 17.4: Submit/View	CollabNet Enterprise Edition permission is retained		
Project Issue Tracking - View	Pluggable: View artifact	CollabNet Enterprise Edition permission is retained		This permission is assigned to the user if the role is a project role (or)if the project uses Project Tracker (or) if permissions are inherited to child projects using Project Tracker (or) if the role is a global/site-wide role.
Project Issue Tracking - View	TeamForge 17.4: dropped	CollabNet Enterprise Edition permission is retained		This permission is redundant, Issue Tracker uses "Query" permission to determine view access.
Domain Issue Tracking - Configure	Pluggable: Configure site	CollabNet Enterprise Edition permission is retained		Conditional mapping. This permission has an impact only if it is part of a site-wide role, otherwise it can be dropped.
Domain Issue Tracking - Configure	TeamForge 17.4: dropped	This feature is not available in TeamForge 17.4		TeamForge 17.4 does not have any such feature.

### TeamForge 17.4 Permissions for Project Dashboard and Tasks

Permissions are mapped to TeamForge 17.4's Tasks application. As there are several gaps between CollabNet Enterprise Edition and TeamForge 17.4's Task tracking application, it is difficult to come up with an exact role mapping. The role played by the user having the permission is used as the baseline for the mapping.

## Permission Mapping Details

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/Gain Details	Comments
Project Dashboard - View		This feature is not available in TeamForge 17.4		TeamForge 17.4 does not have any such feature.
Project Dashboard - View Manager Reports		This feature is not available in TeamForge 17.4		TeamForge 17.4 does not have any such feature.
Tasks - Configure	Task Admin	CollabNet Enterprise Edition permission is retained		There are differences between CollabNet Enterprise Edition and TeamForge 17.4 features, so there isn't any exact comparison.
Tasks - Edit Others	Edit/View	CollabNet Enterprise Edition permission is retained		
Tasks - Edit Self		User loses permission	The user loses rights to edit other's tasks assigned to self.	In TeamForge 17.4 users with "Edit" permission can also edit other tasks in the specified folder(s).
Tasks - Remove Plan	Delete/View	CollabNet Enterprise Edition permission is retained		
Tasks - View	View only	User gains permission	The user gains the ability to view tasks assigned to others	
Tasks - View Others	View only	CollabNet Enterprise Edition permission is retained		"View Only" in TeamForge 17.4 means the user can view any task in the specified folder(s).

**TeamForge 17.4 Permissions for Projects**

TeamForge 17.4 does not have fine level controls with regard to project functions as a project-level role. However, project operations are a part of a site-wide role. Hence the mapping is done to site-wide roles, except for project administrator.

## Permission Mapping Details

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/Gain Details	Comments
Project - Associate Project		User loses permission	The user loses the ability to map project	TeamForge 17.4 allows only project administrators to

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/Gain Details	Comments
			to a category/project group	do so, which is way above this permission.
Project - Associate With Project		User loses permission	The user loses the ability to map project to a category/project group	
Project - Delete	Projects: Delete/View, Categories: Delete/View	User loses permission	The user loses rights to delete project if the role is project-level/Global	Conditional mapping. This permission is granted only if the role is a site-wide role.
Project - Designate Subproject		User loses permission	The user loses the ability to set parent-child relationships between projects	TeamForge 17.4 allows project administrators to set parent-child relationships.
Project - Designate Superproject		User loses permission	The user loses the ability to set parent-child relationship between projects	TeamForge 17.4 allows project administrators to set parent-child relationships.
Project - Edit	Project Admin: Project Admin Permissions, Project Pages: Project Pages Admin	User gains permission	The user gains other project administrator capabilities like setting parent-child relationship etc.	Mapping is done to PCE also since it is a project administrator's responsibility to configure PCE pages.
Project - Invite		This feature is not available in TeamForge 17.4		TeamForge 17.4 does not have any such feature.
Project - Lock		The user loses the ability to lock/unlock projects	Project administrators can do this in TeamForge 17.4.	
Project - View	Projects: View Only	CollabNet Enterprise Edition permission is retained		Conditional mapping. The proposed mapping is done when the role is a domain role. For project/global project roles, user will have the view access just by the role assignment. In the case of a role project/global project role with just this permission, the users

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/ Gain Details	Comments
				belonging to the role are added as a member of the project in case of private projects.
Project Content - View		User gains permission	Without having such a permission, users will gain access to remote published contents.	TeamForge 17.4's remote publishing implementation needs only "View" access to the project.
Project Member - List		This feature is not available in TeamForge 17.4		TeamForge 17.4 does not have any such feature.
Project Page - View	Projects: View Only	User gains permission	Exact changes not clear	Conditional mapping. The same approach as that of "Project – View" permission is followed. There is no clear separation of responsibilities between "Project – View" and "Project Pages – View". In TeamForge 17.4, these are merged.
Private Project Document - View		CollabNet Enterprise Edition permission is retained		
Project - Site-specific		User gains permission	The user will lose the ability to control access to site specific contents. Users will see all project content if the user gets permissions to the project.	TeamForge 17.4 does not have any such feature.

### TeamForge 17.4 Permissions for Role

TeamForge 17.4 does not have fine level controls with regard to roles functions at project-level. However, role operations are a part of a site-wide role. Hence the mapping is done only in the case of domain roles. For project level roles, project administrator has these capabilities.

#### Permission Mapping Details

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/ Gain Details	Comments
Role - Add	Roles: Create/Edit/View	User loses permission	The user loses the ability if it is a	Conditional mapping. The permission is


CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/Gain Details	Comments
			project-level role and the user does not get project administrator permissions.	granted if the role is a domain role and the role also has "Role – Grant" and "Role – Edit" permissions.
Role - Delete	Roles: Delete/View	User loses permission	The user loses the ability if it is a project-level role and the user does not get project administrator permissions.	Conditional mapping. The permission is granted if the role is a domain role.
Role - Edit	Roles: Create/Edit/View	User loses permission	The user loses the ability if it is a project-level role and the user does not get project administrator permissions.	Conditional mapping. The permission is granted if the role is a domain role and the role also has "Role – Add" and "Role – Grant" permissions.
Role - Grant	Roles: Create/Edit/View	User loses permission	The user loses the ability if it is a project-level role and the user does not get project administrator permissions.	Conditional mapping. The permission is granted if the role is a domain role and the role also has "Role – Add" and "Role – Edit" permissions.
Role - Request	Alternate Means	CollabNet Enterprise Edition permission is retained		The user may get additional permissions if the user does not have "Project – View" permission.
Unapproved Project Role - Grant		This feature is not available in TeamForge 17.4		TeamForge 17.4 does not have any such feature.

### TeamForge 17.4 Permissions for Resource Patterns


#### Permission Mapping Details

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/Gain Details	Comments
Resource - Add		User loses permission	The user will lose the ability unless it is a project role and the user gets project administrator permissions.	Path-based permissions (PBP) is the equivalent and is configurable only by project administrators.

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/ Gain Details	Comments
				No such feature at domain level.
Resource - Delete		User loses permission	The user will lose the ability unless it is a project role and the user gets project administrator permissions.	Path-based permissions (PBP) is the equivalent and is configurable only by project administrators. No such feature at domain level.
Resource - Edit		User loses permission	The user will lose the ability unless it is a project role and the user gets project administrator permissions.	Path-based permission (PBP) is the equivalent and is configurable only by project administrators. No such feature at domain level.

 **Tip:** Some examples given below may give you an idea of how PBP can be used to replace the regular patterns that were being used.

Example of regex that can be transformed completely using PBP:	<p>/tags(\$/.**) - This pattern implies anything that starts with /tags/ or exact '/tags'.</p> <p>As the regex is like a super path pattern, here a PBP like [/tags] can suffice.</p>
Example of regex that can be transformed only partially:	<p>/branches(\$/)\$ - This pattern implies only '/branches' and it does not match '/branches/xyz'.</p> <p>This can not be achieved using PBP as [/branches] would match everything under '/branches'. However, as a complex approach, you may need to grant access to [/branches] and disable all its child path access with explicit path components.</p>

 **Note:** It is recommended that you try this on stage before doing it in production.

## TeamForge 17.4 Permissions for Search

### Permission Mapping Details

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/ Gain Details	Comments
Search - Subversion Revision Properties		User gains permissions		TeamForge 17.4 doesn't have any such permission and it is controlled by version control permissions.



### TeamForge 17.4 Permissions for Version Control

Permissions are mapped to TeamForge 17.4's Source Code application. Resource pattern association with the permissions are not handled and will appear in the exception report. The "Change in Behavior" column ignores the regex aspect.

#### Permission Mapping Details

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/ Gain Details	Comments
VersionControl - Add	Commit/View	User gains permission		The user will get permission to modify existing files unless the role also has VersionControl – Modify.
VersionControl - Copy	Commit/View	User loses permission		Conditional mapping. This is mapped only if the user also has the "VersionControl – Add" permission.
VersionControl - Delete	Delete/View	CollabNet Enterprise Edition permission is retained		
VersionControl - Modify	Commit/View	User gains permission		The user will get permission to add new files unless the role also has the "VersionControl – Add" permission.
VersionControl - Read	View only	CollabNet Enterprise Edition permission is retained		
VersionControl CC - Admin				This Clear Case permission is dropped in TeamForge 17.4.
VersionControl CC - Change Event				This Clear Case permission is dropped in TeamForge 17.4.
VersionControl CC - Change Type				This Clear Case permission is dropped in TeamForge 17.4.
VersionControl CC - Checkout				This Clear Case permission is dropped in TeamForge 17.4.
VersionControl CC - Protect				This Clear Case permission

<b>CollabNet Enterprise Edition Permission</b>	<b>TeamForge 17.4 Permission</b>	<b>Change in Behavior</b>	<b>Functionality Loss/ Gain Details</b>	<b>Comments</b>
				is dropped in TeamForge 17.4.
VersionControl CC - Reserve				This Clear Case permission is dropped in TeamForge 17.4.
VersionControl CC - Unreserve				This Clear Case permission is dropped in TeamForge 17.4.
VersionControl CC Attribute - Create				This Clear Case permission is dropped in TeamForge 17.4.
VersionControl CC Branch - Create				This Clear Case permission is dropped in TeamForge 17.4.
VersionControl CC BranchType - Create				This Clear Case permission is dropped in TeamForge 17.4.
VersionControl CC Hyperlink - Create				This Clear Case permission is dropped in TeamForge 17.4.
VersionControl CC Label - Create				This Clear Case permission is dropped in TeamForge 17.4.
VersionControl CC Trigger - Create				This Clear Case permission is dropped in TeamForge 17.4.
VersionControl CVS - Admin	Source Code Admin	User gains permission		This mapping is done only if project role and the project uses CVS (or) if the permissions are inherited to child projects using CVS (or) the permission is part of a global/ site-wide role. As there is no one-to-one mapping between CollabNet Enterprise

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/ Gain Details	Comments
				Edition actions and TeamForge 17.4 clusters, all read-only operations are mapped to "View Only" permission and write operations are mapped to the "Commit/View" permissions. The user may get additional permissions in some cases.
VersionControl CVS - Edit	Commit/View	User gains permission		This mapping is done only if project role and the project uses CVS (or) if the permissions are inherited to child projects using CVS (or) the permission is part of a global/site-wide role. As there is no one-to-one mapping between CollabNet Enterprise Edition actions and TeamForge 17.4 clusters, all read-only operations are mapped to "View Only" permission and write operations are mapped to the "Commit/View" permissions. The user may get additional permissions in some cases.
VersionControl CVS - Module	View only	User gains permission		This mapping is done only if project role and the project uses CVS (or) if the permissions are inherited to child projects using CVS (or) the permission is part of a global/site-wide role. As there is no one-to-one mapping between

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/ Gain Details	Comments
				CollabNet Enterprise Edition actions and TeamForge 17.4 clusters, all read-only operations are mapped to "View Only" permission and write operations are mapped to the "Commit/View" permissions. The user may get additional permissions in some cases.
VersionControl CVS - Release	View only	User gains permission		This mapping is done only if project role and the project uses CVS (or) if the permissions are inherited to child projects using CVS (or) the permission is part of a global/site-wide role. As there is no one-to-one mapping between CollabNet Enterprise Edition actions and TeamForge 17.4 clusters, all read-only operations are mapped to "View Only" permission and write operations are mapped to the "Commit/View" permissions. The user may get additional permissions in some cases.
VersionControl CVS - Tag	Commit/View	User gains permission		This mapping is done only if project role and the project uses CVS (or) if the permissions are inherited to child projects using CVS (or) the permission is part of a global/site-wide role. As there is no one-to-

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/ Gain Details	Comments
				one mapping between CollabNet Enterprise Edition actions and TeamForge 17.4 clusters, all read-only operations are mapped to "View Only" permission and write operations are mapped to the "Commit/View" permissions. The user may get additional permissions in some cases.
VersionControl CVS - Watch	Commit/View	User gains permission		This mapping is done only if project role and the project uses CVS (or) if the permissions are inherited to child projects using CVS (or) the permission is part of a global/ site-wide role. As there is no one-to-one mapping between CollabNet Enterprise Edition actions and TeamForge 17.4 clusters, all read-only operations are mapped to "View Only" permission and write operations are mapped to the "Commit/View" permissions. The user may get additional permissions in some cases.
VersionControl CVS Module - Add	Commit/View	User gains permission		This mapping is done only if project role and the project uses CVS (or) if the permissions are inherited to child projects using CVS (or) the permission is part of a global/ site-wide role. As

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/ Gain Details	Comments
				there is no one-to-one mapping between CollabNet Enterprise Edition actions and TeamForge 17.4 clusters, all read-only operations are mapped to "View Only" permission and write operations are mapped to the "Commit/View" permissions. The user may get additional permissions in some cases.
VersionControl CVS Module - Delete	Commit/View	User gains permission		This mapping is done only if project role and the project uses CVS (or) if the permissions are inherited to child projects using CVS (or) the permission is part of a global/site-wide role. As there is no one-to-one mapping between CollabNet Enterprise Edition actions and TeamForge 17.4 clusters, all read-only operations are mapped to "View Only" permission and write operations are mapped to the "Commit/View" permissions. The user may get additional permissions in some cases.
VersionControl CVS Module - Edit	Commit/View	User gains permission		This mapping is done only if project role and the project uses CVS (or) if the permissions are inherited to child projects using CVS (or) the permission is part of a global/

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/ Gain Details	Comments
				site-wide role. As there is no one-to-one mapping between CollabNet Enterprise Edition actions and TeamForge 17.4 clusters, all read-only operations are mapped to "View Only" permission and write operations are mapped to the "Commit/View" permissions. The user may get additional permissions in some cases.
VersionControl CVS Module - View	View only	User gains permission		This mapping is done only if project role and the project uses CVS (or) if the permissions are inherited to child projects using CVS (or) the permission is part of a global/site-wide role. As there is no one-to-one mapping between CollabNet Enterprise Edition actions and TeamForge 17.4 clusters, all read-only operations are mapped to "View Only" permission and write operations are mapped to the "Commit/View" permissions. The user may get additional permissions in some cases.
VersionControl CVS Repository - Delete	Delete/View	User gains permission		This mapping is done only if project role and the project uses CVS (or) if the permissions are inherited to child projects using CVS (or) the permission

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/ Gain Details	Comments
				is part of a global/ site-wide role. As there is no one-to-one mapping between CollabNet Enterprise Edition actions and TeamForge 17.4 clusters, all read-only operations are mapped to "View Only" permission and write operations are mapped to the "Commit/View" permissions. The user may get additional permissions in some cases.
VersionControl CVS Repository - Rename	Source Code Admin	User gains permission		This mapping is done only if project role and the project uses CVS (or) if the permissions are inherited to child projects using CVS (or) the permission is part of a global/ site-wide role. As there is no one-to-one mapping between CollabNet Enterprise Edition actions and TeamForge 17.4 clusters, all read-only operations are mapped to "View Only" permission and write operations are mapped to the "Commit/View" permissions. The user may get additional permissions in some cases.
VersionControl CVS Repository - ToggleExecutable	Commit/View	User gains permission		This mapping is done only if project role and the project uses CVS (or) if the permissions are inherited to child projects using CVS



<b>CollabNet Enterprise Edition Permission</b>	<b>TeamForge 17.4 Permission</b>	<b>Change in Behavior</b>	<b>Functionality Loss/ Gain Details</b>	<b>Comments</b>
				(or) the permission is part of a global/ site-wide role. As there is no one-to-one mapping between CollabNet Enterprise Edition actions and TeamForge 17.4 clusters, all read-only operations are mapped to "View Only" permission and write operations are mapped to the "Commit/View" permissions. The user may get additional permissions in some cases.

#### **TeamForge 17.4 Permissions for Wiki**

Wiki will come in as an integrated application post TeamForge 17.4. At that time, these permissions will be transferred to the integrated application. As of now, the permissions are mapped to TeamForge 17.4 wiki permissions and hence users will get the same permission in TeamForge 17.4's wiki.

#### Permission Mapping Details

<b>CollabNet Enterprise Edition Permission</b>	<b>TeamForge 17.4 Permission</b>	<b>Change in Behavior</b>	<b>Functionality Loss/ Gain Details</b>	<b>Comments</b>
Project Wiki - Administer	Wiki Admin, Delete/View	CollabNet Enterprise Edition permission is retained		CollabNet Enterprise Edition wiki allows replacing CollabNet Enterprise Edition actions with ACL in wiki pages. This mapping ignores that setting since TeamForge 17.4 does not have that feature.
Project Wiki - Contribute	Create/Edit/View	CollabNet Enterprise Edition permission is retained		CollabNet Enterprise Edition wiki allows replacing CollabNet Enterprise Edition actions with ACL in wiki pages. This mapping ignores that setting since TeamForge 17.4 does not have that feature.

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/Gain Details	Comments
Project Wiki - View	View only	CollabNet Enterprise Edition permission is retained		CollabNet Enterprise Edition wiki allows replacing CollabNet Enterprise Edition actions with ACL in wiki pages. This mapping ignores that setting since TeamForge 17.4 does not have that feature.

## Reference information about CollabNet TeamForge site administration




Here is some detailed reference information to help you figure out what is what in this CollabNet product.

### Replica status icons



The icons represent the health of a replica server, the state of a replicated repository, or the status of a replication command.




You can get more information related to a specific status by clicking on its icon. However, not all icons are links.

### Replica server





Icon	Description
	There are no errors. The replica server has polled the master within 3 times the <b>Command polling interval</b> . For example, if the polling frequency is 10 seconds, TeamForge has a record of contact within the last 30 seconds.
	This could mean one of two things: <ul style="list-style-type: none"> <li>The server has an error in the last 30 commands, although the most recent command is successful.</li> <li>The server has not polled the master within 3 times the <b>Command polling interval</b>, although it has within 30 times the polling interval.</li> </ul>
	This could mean one of two things: <ul style="list-style-type: none"> <li>The most recent command is a failure.</li> <li>The server has not polled the master within 30 times the <b>Command polling interval</b>.</li> </ul>

### Replicated repository

Icon	Description
	Replicated repository is in sync with the repository on the master.
	Replicated repository is being added.

Icon	Description
	Replicated repository is being deleted.
	Replicated repository is being synchronized using <code>svnsync</code> .
	Synchronization events are queued on the TeamForge server and waiting for the replica to process them.

### Replication command

Icon	Description
	The command is queued in the TeamForge server for the replica.
	Depending on the replica server queue, the command is being processed or scheduled to be processed.
	The execution of the command failed in the replica server. You will see a link to the <i>server log</i> that contains details about the error.
	The execution of the command succeeded in the replica server.

### Migrating tools from Enterprise Edition to TeamForge 17.4

As you move onto CTF 17.4, you may want to know how the CollabNet Enterprise Edition (CEE) project information is brought over to CollabNet TeamForge 17.4 (CTF 17.4).

### Migrating tools from CollabNet Enterprise Edition to CollabNet TeamForge 17.4

As you move onto CTF 17.4, you may want to know how the CollabNet Enterprise Edition (CEE) project information is brought over to CollabNet TeamForge 17.4 (CTF 17.4).

### CEE and CTF 17.4 Tool Mapping Details

Most of the tools and their functionality are similar between CEE and CTF 17.4. However, within each tool, some work-arounds might have been done to bring over the CEE project information and resources.

This table shows how tools in CEE map to CTF 17.4 tools:

CEE Tool	CTF 17.4 Tool
Announcements	Project News
Discussions	Discussions
Documents & Files	Documents
Wiki	<Data is not migrated>
CVS	CVS
Subversion	Subversion
Project Dashboard/Tasks	Tasks
Issue Tracker	Tracker
Project Tracker	Project Tracker (Integrated Application)

CEE Tool	CTF 17.4 Tool
Project metrics	<Equivalent tool is not available.>
Project Pages	Project Pages
Remote Publishing	Remote Publishing

This table shows how objects in CEE map to CTF 17.4 objects:

CEE Objects	CTF 17.4 Objects
Project	Project
Project Group	Project
Category	Category + Project

### How to find what you are looking for?

There are several tools that are migrated from CollabNet Enterprise Edition to CollabNet TeamForge 17.4.

You can look through the list for the required tool details. Within each tool, the information is divided into two sections - **Migration Overview** and **Notes**. Migration Overview section provides the mapping information whereas the Notes section lists the work-arounds or dissimilarities, if any.

### Overview

- Post migration, CEE's audit log entry table is available in CTF 17.4's data folder as "cee\_backedup\_audit\_log" table.
- Post migration, project tracker permissions are not migrated to CTF 17.4's tracker and vice versa. Project Tracker permissions are transferred to Project Tracker which is brought in to CTF 17.4 as an integrated application.
  - 👉 **Note:** Project hierarchy and flow of permissions to child projects is a known issue as of now. Suppose before migration, the parent project is Issue Tracker (CEE) based, child project is Project Tracker (PT) based and some roles from the parent projects are in use in the child project. In all probability, the role inheritance will not be migrated.
- The look project contains a single consolidated exception report - Documents>Root Folder>cee\_migration\_exception\_report> GLOBAL-REPORT-CONSOLIDATED.txt.
- The Documents>Root Folder>cee\_migration\_exception\_report><project name>REPORT.txt folder holds all per project exception reports.
- The CEE files that were committed in a CVS repository, display junk characters when directly viewed in CTF 17.4. It is recommended to download the file and view it.
- The "www" content from CEE's www project is moved to the "html" directory in the branding repository of "look" project in CTF 17.4.

The assumptions made while migrating CEE to CTF 17.4 are given here: [Migration Assumptions](#) on page 412

### List of the Tools Migrated from CEE to CTF 17.4

#### Migration Assumptions

As you move onto CTF 17.4, you may want to know the assumptions that were made while migrating CollabNet Enterprise Edition (CEE) project information to CollabNet TeamForge 17.4 (CTF 17.4).

- Global Data:
  - Default usergroups '\*Domain Users' and '\*\*Host Users' are not migrated.
  - System users, if part of any user group, are dropped.
  - System users like guest, webmaster, postmaster and mailer-domain are not migrated and their references are mapped to 'Nobody' user.

- Host Superuser (root) is not migrated in favor of administrator account in CTF 17.4 and its references are mapped to the administrator user.
- Users who are associated with Host Admin role are migrated as Site administrators.
- Domain roles; 'Anonymous Guest' and 'Registered User' are dropped and their permissions are transferred to each public project's Default Access Permissions user class.
- Host roles are not migrated.
- Project announcements of www project are moved to the Look project as news.
- CEE's HAC\_AUDIT\_LOG\_ENTRY table is moved to cee\_backedup\_audit\_log.
- Projects Data:
  - Unapproved and deleted projects are not migrated.
  - Default projects like testproject, www, project-templates are not migrated.
  - Categories and project groups are migrated as normal projects.
  - License detail of projects is not migrated.
  - CVS repositories are accessible with a new URL only.
  - CVS working copies must be recreated.
  - SVN working copies continue to work as is.
  - The Alias URL of the project is not migrated.
  - When the project creation date is not found for a project in the audit log, the default fallback value (2000-01-01) is used.
  - Project owners message is moved to the news component as a news item.
- Discussion services tool:
  - Soft deleted discussions are not migrated.
  - Trusted discussions are migrated as moderated discussion and trusted subscribers are moved to the trusted users list.
  - The following discussion settings are dropped as CTF 17.4 does not support these settings:
    - Web UI View
    - Web Posting
    - Web Post CAPTCHA
    - Display Email Address
    - Remove Deleted Messages
    - Remove Unmoderated Messages
    - Content in Email
    - Message Size
    - Attachments
    - Advanced Email Filters
  - The following discussion topic value is dropped as CTF 17.4 does not support it:
    - Message headers
  - In CTF 17.4, discussion email address must begin with an alphabet and it can not contain special characters except an underscore. So, if a CEE's discussion address contains hypens or periods, they are replaced with underscores. Also, if it begins with a number, it is replaced with the equivalent number in word followed by an underscore.
  - Lengthy discussion topic titles, if any, are truncated to 250 characters.
  - Soft deleted approved and unapproved topics including replies are not migrated.
- Documents and files tool:
  - Unapproved folders and documents are not migrated.
  - In CTF 17.4, a document's status can be either draft or final. So, CEE's documents in 'Stable' status are migrated as final and documents in any other status are migrated as draft.
- Announcements tool:
  - Announcements are migrated to the news component in CTF 17.4.

- CEE features like Image url, publishing and event dates are not supported in CTF 17.4. Those details are appended as part of announcements itself.
- Project pages tool:
  - The documents existing in CEE's "Document" component are moved to a sub-folder in the CTF 17.4's "Documents" tool's "cee-migrated-docs" folder.
  - The page components, HTML and Activity table components are migrated as text/html component.
  - The page components Metrics report, Milestones table and Project overview are not migrated.
  - The availability property of a page 'All users' and 'Users with project-edit' is mapped to 'Visible'; and 'None' will be mapped to 'Hidden' in CTF 17.4.
  - Backward compatibility of free-form page references in html components is not supported.
  - Sub-page icons are not migrated.
- Issue Tracker tool:
  - Issue tracker in CEE is migrated to a tracker called 'Issue tracker' in CTF 17.4.
  - Issues are migrated as artifacts.
  - The email notification preferences are not migrated.
  - The associations between Component and Subcomponent, Component and Version and, Component and Target Milestone are dropped as CTF 17.4 does not support these associations.
  - Issue tracking configuration parameters not migrated.
  - Internal users of Issue tracker are dropped.
  - Dependencies are migrated as associations.

#### Mapping functionality between Enterprise Edition and TeamForge 17.4

Most of the tools and their functionality are similar between CEE and CTF 17.4. However, within each tool, some work-arounds might have been done to bring over the CEE project information and resources.

There are several tools that are migrated from CollabNet Enterprise Edition to CollabNet TeamForge 17.4.

You can look through this section for the required tool details and work-arounds or dissimilarities, if any.

This table shows how the tools in CEE map to CTF 17.4 tools:

CEE Tool	CTF 17.4 Tool
Announcements	Project News
Discussions	Discussions
Documents & Files	Documents
Wiki	<Data is migrated post-migration, as required>
CVS	CVS
Subversion	Subversion
Project Dashboard/Tasks	Tasks
Issue Tracker	Tracker
Project Tracker	Project Tracker (Integrated Application)
Project metrics	<Equivalent tool is not available.>
Project Pages	Project Pages
Remote Publishing	Remote Publishing

This table shows how the objects in CEE map to CTF 17.4 objects:

CEE Objects	CTF 17.4 Objects
Project	Project
Project Group	Project
Category	Category + Project

### Migrating projects and categories

Here are the migration details and migration notes for projects and categories.

#### Migration Overview

- The public and private projects are retained appropriately.
- The project members in CTF 17.4 consist of users assigned a role in the current project.
  - 👉 **Note:** In CEE, a project member must have a role whereas, in CTF 17.4, a user can be a project member without being assigned a role.
- Locked projects in CEE are migrated as locked projects to CTF 17.4.
  - 👉 **Note:** The locked projects seem to allow certain actions even after being locked. For example, forums can still be monitored by administrators.
- CEE projects that are imported into CTF 17.4 can be selectively deleted, if required.
- CEE's Project groups project is migrated as a regular project in CTF 17.4. Members of the Project groups project are migrated as it's child projects in CTF 17.4.
- In CEE, categories could behave as projects, having tools like tracker, documents, discussion and so on. However in CTF 17.4, categories are plain containers which have projects in them, enabling any project to be placed in a category. CTF 17.4 categories do not have project tools like tracker, documents, discussions etc.
  - 👉 **Note:** Categories from CEE to CTF 17.4 are migrated in two steps:
    - A CTF category (with the same CEE category name) is created
    - The tool data (eg. artifacts, documents, discussions etc) are migrated into a CTF 17.4 project with the CEE category name.

Consider this example. Assuming there exists a CEE category (named - "cat") having projects (proj1, proj2) with tool data of its own (document folders: folder1, folder2). When this category is migrated to CTF 17.4, a CTF 17.4 category called "cat" is created, and the projects (proj1 and proj2) are placed in this category. Also, a CTF 17.4 project "cat" is created which holds the tool data that the category had in CEE (document folders folder1, folder2).

- The project-subproject, category-subcategory association is maintained.

This table shows how Project fields in CEE map to Project fields in CTF 17.4:

CEE Projects	CTF 17.4 Projects
Project Name	Project Name
Summary	Description
Parent Project	Parent Project
Project Categories	None
Public Project	None
Project Language	None
Versioning Component	None
Tracking Component	None

CEE Projects	CTF 17.4 Projects
Artifact Code	None
Project Owner	None
Project Proposer	Created by user name
Date Last Modified	Date Created/Date Last Modified
Project Approver	Last modified by (user)
Owner's Message	None
CVS Projects	None
SVN Projects	None

This table shows how Category fields in CEE map to Category fields in CTF 17.4:

CEE Categories	CTF 17.4 Categories
Category Name	Category Name
Parent Category	None
Public Category	None
Use Category Index.html	None
Category Language	None
Initial Projects	None
Use Category Home Page	None

## Notes

- In CTF 17.4, unapproved projects (requested project) are just project-creation requests, with the project name and description. However, in CEE, unapproved projects can be actively used by the project creators. As the unapproved projects are conceptually different between CEE and CTF 17.4, these projects are not migrated.
- The IsDeleted field in CTF 17.4 corresponds to whether the project is deleted or not in CEE.
- If the project creation date is not found in the audit log, the fall-back value specified in the `/opt/collabnet/cee_dump_load/conf/c6migrate.conf` file is used.
- The Date Created field in CEE does not have an equivalent field in CTF 17.4.
- There is no default project language in CTF 17.4.
- The monitoring feature in CTF 17.4 supports desired user notification.
- In CTF 17.4 a project is created "Private" by default.
- In CTF 17.4, the "Artifact Code" is constant across the site. It is not configurable project-wise.
- In CTF 17.4, the project creator becomes default project owner.
- CEE categories have project tools while the CTF 17.4 categories do not have tools.
- The name of the category can be edited in CTF 17.4.
- Only the site administrator can create project category in CTF 17.4.

## Migrating roles and permissions

Here are the migration details and migration notes for roles and permissions data.

### Migration Overview

- Project roles are migrated appropriately.



- CEE's equivalent of Project roles with global/all projects visibility are global project roles in CTF 17.4 and these roles are migrated appropriately.
- CEE's equivalent of Domain Roles are Site-wide roles in CTF 17.4 and these are migrated appropriately.
- There is no equivalent of CEE's Host roles, users belonging to "Host Admin" role are migrated as Site Administrators in CTF 17.4. Users belonging to other Host roles, if any, will lose such roles after migration and the details are exception reported.
- CEE's Host SuperUser (root) is mapped to CTF 17.4's TeamForge Administrator (admin).
- For details of the mapped permissions, see [User permissions on a CEE and CTF 17.4 site](#) on page 430 or [User permissions in CEE and CTF 17.4 projects](#) on page 388
- User and group associations with the roles are retained.
- Role associations with internal user groups - Host Users and Domain Users, are not migrated since the feature is not available in CTF. These mappings will appear in the exception report.
- Project role requests from CEE projects are migrated as same to CTF 17.4.
- CTF 17.4 administrators can configure permission mapping using the configuration file. The `permissionmapping.conf` configuration file is available here: `/opt/collabnet/cee_dump_load/conf/`.

This table shows how fields in CEE map to fields in CTF 17.4:

CEE	CTF 17.4
Host administrators	Site-administrators
Domain roles	Site-wide roles
Project roles	Project roles
Project roles with All projects visibility	Global project roles
Name	Role Name
Description	Description
Requestable	Requestable
Ownership Role	None
Grant role on subproject creation	None
Block recursion into private projects	Prevent Access (Site-wide roles) or Prevent Inheritance (Project/Global roles)

### Notes

- CEE's Host roles are not migrated, multi-tenancy is implemented differently in CTF 17.4.
- Domain Roles - Registered User role is dropped in favor of "Default Access Permissions" at individual project level. In CTF 17.4, there is no global control like a domain role.
- Domain Roles - Anonymous Guest role is dropped in favor of "Default Access Permissions" at individual project level. In CTF 17.4, there is no global control like a domain role.
- The Ownership role and Grant role on sub-project creation functionalities are dropped.
- If the role creation/modification date is not found in the audit log, either the project creation/modification or the fall-back date is used. The default fall-back value is "2000-01-01 00:00:00", as specified in the `/opt/collabnet/cee_dump_load/conf/c6migrate.conf` file.

### Migrating SCM integrations

Here are the migration details and migration notes for Source Code Manager integrations.

## Migration Overview

- Two SCM integrations can be created during migration. One for Subversion and another for CVS. CVS creation is optional.
- The version control's view page in CTF 17.4 shows extra information in migrated folders in the **Age**, **Author** and **Last Log Entry** columns.

## Notes

- In the migrated repository (in CTF 17.4), the **Commits this week** column displays "0", even if there had been commits in CEE during that week. This happens because the historical commit details are not captured for migration.
- In CEE, there can be only one version control component for a project whereas in CTF 17.4, there can be multiple version control components associated to a project.
- Post-migration, new commits are not archived in commit discussions. However, during migration, the existing archived commits are migrated.
- Global search does not work for source code. The ability to search source code is not supported as the revisions are not converted to commit objects during migration.

## Migrating users and user groups

Here are the migration details and migration notes for users and user groups.

## Migration Overview

- CEE's Host Superuser "root - root@domainname" is migrated as CTF 17.4's TeamForge administrator "admin - root@collab.net".
- CEE's Host administrators are migrated as CTF 17.4's site administrators.
- User group mapping and project-role association is retained.

This table shows how fields in CEE map to fields in CTF 17.4:

CEE	CTF 17.4
Host Superuser	Administrator
Host Administrator	Site Administrator
User - User Group association	User - User Group association
User Group - Project - Role association	User Group - Project - Role association
Group Name	Full Name
Short Description	Description
Initial users	None
Allow new user to group	Add
User Location	None
Username	User Name
Email	Email Address
Full Name	Full Name
Organization	Organization
Password	Password
Confirm Password	Confirm Password
Email Language	None

CEE	CTF 17.4
Creation Options	Send Welcome Message?
None	Restricted User
None	Site Admin
Secondary Email	Alternate emails
Has seen click-through	None
Has seen welcome page?	None
Include user in subversion reporting	None
Project Tracker Self notification	None
Project Tracker maximum report time	None
Group memberships	None
Times logged in	None
Domain-wide role	
Host-wide roles	Site Admin
None	Profile Picture
None	Status
None	Date Time Format
None	Detail

## Notes

- Post migration, CEE users must raise a "Forgot password" request to login into CTF 17.4. Using the response e-mail, the users can reset their passwords.
- CEE's default groups – "\*\*\*Host Users" and "\*Domain Users" are not migrated to CTF 17.4. The default user group associations (role mapping) will not be retained post-migration.
- The following are the default users in CEE: "guest", "webmaster", "postmaster", "mailer-deamon" and "root". Post-migration, "root" is mapped to "admin" and other default users are mapped to "nobody" in CTF 17.4.
- Unix system user accounts are not migrated from CEE to CTF 17.4. They are either dropped or renamed.
- Issue Tracker's internal users (for example, issues@ert) are mapped to "Nobody".
  - 👉 **Note:** Users of the form issues@projectname are not migrated. All assignments or concepts to such users will be set to "nobody/none" after migration.
- Special Users of CEE are not migrated to CTF 17.4.
- CEE allows users to add up to four secondary email addresses to their accounts. In CTF the supported alternate (secondary) email addresses are three. The first three secondary emails are migrated to CTF; the fourth email address is not migrated.
- Email language feature is not available in CTF 17.4.
- In CTF 17.4, notification setting is same for all the tools.

## Migrating wiki

Here are the migration details and migration notes for wiki pages.

## Migration Overview

- Only the latest revision of wiki pages is migrated from CEE to CTF 17.4.

- Wiki migration tool compares each of the wiki pages taken for migration with a static set of hard-coded page names given below. If the page names are found to be same, the wiki pages are not migrated. The list of reserved wiki pages is here: [Reserved Wiki pages](#)
- Source wiki page "FrontPage" is migrated as "HomePage" wiki page in CTF 17.4.
- Source wiki page "HomePage" is migrated as "HomePageMig" wiki page in CTF 17.4.
- If the source wiki page contains a URL without enclosing square brackets, it is migrated as text (a broken link).

### Notes

- Wiki pages which do not have CamelCase page names are not migrated.
- If an attachment file is missing at the source, then it is migrated as a broken link.
- Subpages links specified in source wiki page are not retained post-migration.
- Macros are not supported in CTF 17.4, hence macros are not migrated.

### Migrating project pages

Here are the migration details and migration notes for project pages.

### Migration Overview

- For PCE-enabled subversion projects, CTF 17.4 Home Page is replaced with CEE Home Page.
- For CVS and non-PCE enabled subversion projects, CTF 17.4 Home Page is replaced with remote publishing content (index.html).
- Top level pages and sub-pages are retained.
- The HTML, Activity table, Subprojects, freeform pages and documents components of CEE are migrated to CTF 17.4.
  - 👉 **Note:** The documents PCE content in CEE is migrated to the Documents tools in CTF 17.4 and the functionality to display the same via CTF 17.4 Docs PCE is provided. Post migration, the documents relevant to PCE are stored under the `Documents from CEE Project Pages` folder. Appropriate permissions on the Documents tool in CTF 17.4 are required.
- CEE's HTML and Activity table components are migrated to the Text component in CTF 17.4.

This table shows how fields in CEE map to fields in CTF 17.4:

CEE field	CTF 17.4 field
Home Page	CEE Home Page - for PCE enabled Subversion projects
Home Page	Index.html - for CVS/Non-PCE enabled Subversion projects
Top level - Sub pages associations	Top level - Sub pages associations
Html, Activity table, Subprojects, Free form pages, Document components	Text, Subprojects, Free form pages, Document components

### Notes

- The free-form page links mentioned in the sub-pages do not work in CTF 17.4 as URL compatibility doesn't work for free-form links referred in HTML components. The CTF 17.4 site/project admin would need to point the links to correct pages.
- The icons added in sub-page are not migrated to CTF 17.4.
- The CEE components that are not migrated to CTF 17.4 include: Metrics report, milestones and project overview.
- The project home pages that displayed sub-projects in CEE will not show the sub-projects list in CTF 17.4 until the PCE is enabled.
- In CTF 17.4, you can enable the PCE for CVS projects, if desired.

- The project page content is not version controlled in CTF 17.4. During migration, CEE's latest version is moved over to CTF 17.4. The old information (older revisions) can be accessed through the current project's subversion repository from `trunk/www/cn-project-pages/`.

## Migrating Issue Tracker

Here are the migration details and migration notes for Issue Tracker.

### Migration Overview

- Issue Tracker in CEE is migrated to a tracker called "Issue tracker" in CTF.
- Each CEE issue becomes an artifact in the "Issue tracker" tracker.
- Each CEE issue can be identified with its corresponding artifact, as the **Issue Id** appears as a prefix in the corresponding artifact's **title**. For example, after migration the title of an artifact might be like - "[issue 1] - Issue Title", where "issue 1" would be the original artifact Id in CEE.
- The association's comment is of the form: "[issueX] depends on [issueY]", where X and Y are the source and target issue ids.
  - 👉 **Note:** The date and time stamps are retained in all instances. Only the last date and time stamps are displayed.
- All the issue activities except dependency activities in CEE are displayed under the **Change Log** tab in CTF. The issue Id of dependent and/or blocker issues are migrated as "Associations".
  - 👉 **Note:** CTF does not have corresponding fields for "dependent" or "blocker" issues. As the change log tab displays only the changes that are made to the tracker fields, dependency and blocker changes are not displayed there.
- The comments and status changes are displayed under the **Status/Comments** tab.
- All the attachments in an issue are made available under the **Attachments** tab of the respective artifact.
- The **data**, **file name**, **file size**, **creator name** and **creation time** details of the attachments are retained.
- Each user on the **CC** list of an issue is added to the **Users Monitoring** list of the corresponding artifact.

This table shows how fields in CEE map to fields in CTF 17.4:

CEE field	CTF field
Summary	Title
First Description	Description
Priority	Priority
Status - Unconfirmed/New/Started/Reopened/Resolved/Verified/Closed	Status - Unconfirmed/New/Started/Reopened/Resolved/Verified/Closed
State - Open/Close	State - Open/Close
Issue Depends On	Associations
Issue Blocking	Associations
Issue Closed Date	Artifact Closed Date
Issue Activities	Change Log tab
Issue Comments	Status/Comments tab
Attachments	Attachments tab
Attachment - Description	Status/Comments tab
Cc Users	Monitoring Users
None	Group

CEE field	CTF field
None	Category
None	Customer
None	Reported in Release
None	Fixed in Release
Zero	Fixed in Release
Zero	Actual Hours

New fields created in CollabNet TeamForge:

- URL
- Issue Type
- OS
- Component
- Subcomponent
- Platform
- Version
- Resolution
- Target Milestone
- QA Contact
- Status Whiteboard
- Keywords

## Notes

- In CTF 17.4, you can monitor an artifact even after you are no longer assigned to it. You must explicitly stop monitoring, if desired.
  - 👉 **Note:** In CEE, you do not have to be a project member to be in the "Assigned To" list. However, in CTF 17.4, you must be a project member to be in the "Assigned To" list.
- In CTF 17.4, only the project members can be part of monitoring and "Assigned To" lists.
- Internal users, for example, "issues@ert" in CEE are mapped to "None" in CTF 17.4.
- CEE Projects that have set up the "Unconfirmed" state to be the initial status of an artifact are migrated with a change. The Unconfirmed issues in CEE are migrated as Unconfirmed in CTF 17.4, but the "Unconfirmed" configuration setting is not migrated to CTF 17.4. So, that means, no confirmation is required in CTF 17.4.
  - 👉 **Note:** In CTF 17.4, an artifact can always be created only in "New" state (and not in "Unconfirmed" state as possible in CEE); as the user must have the "Submit" permission to create an artifact in CTF 17.4.
- "QA Contact" is a single-select field in CEE, but is made as multi-select field in CTF 17.4.
- History is not retained for "QA Contact", "Issue Owner", "Cc" and "Dependencies" (dependent/blocker issues) fields.
- Votes are not retained. CTF 17.4 does not support this feature.
- If the mime types in CEE attachments are not mapped in correspondence with the attachment type; post-migration, the attachments may not open appropriately. However, the attachments get migrated and can be saved locally and viewed.
- The email notification options set in CEE are not retained while migrating to CTF 17.4.
- Associations in CEE between Component-Subcomponent, Component-Versions and Component-Target Milestones are not retained post-migration.

## Migrating discussions

Here are the migration details and migration notes for discussions.

## Migration Overview

- All the active discussion forums, topics, posts, user subscribers and attachments are migrated to CTF 17.4.
- Attachment formats are retained.
- All the discussion user subscribers in CEE are added to the **Monitoring Users** list in CTF 17.4.
- Un-moderated and moderated discussions are migrated appropriately. However, the trusted discussions are migrated as moderated discussions and trusted subscribers as trusted users.
- Approved and waiting-for-approval topics/posts are migrated as is.
- Only those users who had the "Discussions - Discussion Use Private" permission in CEE can view the migrated private discussions in CTF 17.4.

This table shows how fields in CEE map to fields in CTF 17.4:

CEE	CTF 17.4
Title	Forum Name
Email address	Email Address
Description	Description
Restrictions	Enable Moderation/Type
Discussion User Subscribers	Monitoring Users
Un-moderated/Moderated Discussions	Un-moderated/Moderated Discussions
Trusted discussions/trusted subscribers	Moderated discussions/trusted users
Discussion Description	Forum Description
Prefix	Subject Prefix
Email Posting	Enable Mailing List
Email Address	Email Address
Footer	Footer Text

While posting attachments via the user interface in CTF 17.4, attachments can only be posted one at a time. However, while posting using email, you can post multiple attachments.

This table lists the message posting mapping details:

CEE	CTF 17.4
Subject	Subject
Body	Message
Attachments	Attachments
Other Recipients	Nil

## Notes

- The discussion mailing address format in CEE is `forumalias@project.domain` whereas in CTF 17.4 it is `forumalias-project@domain`.



**Note:** Post-migration, the response emails to the existing CEE mailing address will be re-directed to the new CTF 17.4 mailing address and captured into appropriate mailing lists. However, it's recommended to use the new CTF 17.4 mailing address for any new posts. The future releases of CTF 17.4 may not continue supporting the old mailing address compatibility.

- 👉 **Tip:** A CEE forum list name can start with a number or "\_" and can also have "." and "-" characters, whereas, CTF 17.4 imposes certain naming conventions. So, if the toAddressUsername (list name) in CEE contains "." or "-", it is replaced with "\_" in CTF 17.4. If it starts with "\_" in CEE, the underscore is removed in CTF 17.4. If it starts with a number, post-migration, the number is replaced with the number in word, suffixed by an underscore. For example, "1team" mailing list address becomes "one\_team" post-migration.
- The posts made by guest users in CEE are migrated with their actual mail-ids being prefixed in their messages.
- The **Default Discussions** feature is not available in CTF 17.4. The different default discussions (Issues/Commits) serve as archives for any change to issues/version control commits. These are migrated to CTF 17.4 but post-migration, these forums do not continue as archives.
- Any form (approved/un-approved) of soft-deleted posts/topics/forums is not migrated to CTF 17.4 as CTF 17.4 does not have soft deletion feature.
- The CJK file names are not preserved.
- Subject Prefix and Footer text fields available in CEE's create page are mapped in CTF 17.4's edit page.
- Email messages can not be blocked based on attachments in CTF 17.4.
- In CEE, "Content in email" has two options: Send full text of posted messages and send email notification with link only. In CTF 17.4, send full text is the default email content behavior
- In CEE, "Email notifications" has three options: Message-by-message, digest format and start page only. In CTF 17.4, email notifications are set to "message-per-change" by default.
- In CEE, there is no character length restriction on "Topic Title", whereas, in CTF 17.4, it is restricted to 250 characters.
- In CEE, subscription is allowed at the discussion level whereas in CTF 17.4, it is allowed even at the thread level.
- CEE features that are not available in CTF 17.4 include: Default Discussions, Message size, Message Header, Soft deletion, Guest Posting, User Groups, Anonymous and Sub-list subscribers, Email Subscription, Other Recipients, Web Post CAPTCHA, Remove deleted messages and Remove unmoderated messages.
- In CEE, "Reply Behavior" has four options: Reply to list, To sender only, specific header and do not reply to. In CTF 17.4, Reply to post is the only reply behavior available.
- CTF 17.4 does not support CEE's advanced email filter options like mime types, headers, reject content.
- In CTF 17.4, based on the user's permissions, web UI and/or web posting can be accessed.
- By default, forum mailing list is displayed in CTF 17.4 as the display email address.

### Migrating documents

Here are the migration details and migration notes for documents and files.

### Migration Overview

- In CEE the text field supports three formats: "Default", "Exactly as entered" and "HTML source". In CTF 17.4, there are two types of document mime types: "Text/Plain" and "Text/Html".
- The file format selected as "html source" in CEE is migrated to CTF 17.4 as an attachment.
- The CEE document files with "Default" and "Exactly as entered" format type are displayed as "Text" format type in CTF 17.4.

👉 **Note:** These files can be edited by selecting the **Create online** option under **Update Document**.

- The CEE document files with "Html source" format type are displayed as "Html" format type in CTF 17.4.


👉 **Note:** These files can be edited by selecting the **Upload from file** option under **Update Document**.

- In CEE, a document can be in either of the following states:
  - Draft
  - Reviewed
  - Baselined
  - Archival
  - Obsolete
  - Stable




. In CTF 17.4, a document can only be in either "Draft" or "Final" state.

- All CEE documents except those in "Stable" state are moved over to CTF 17.4 as "Draft" documents. The "Stable" documents are moved over as "Final" file status.

 **Note:** The "Obsolete" state may have a prefix identifier in the "Draft" state.

- The folder hierarchy is retained.
- The file version of moved files is set as version 1.
- Locked documents are retained as locked.

 **Note:** If a document in CEE is set with a future date for locking, when migrated, it will be migrated as an un-locked document only. Locked documents can be viewed but not edited.

This table shows how fields in CEE map to fields in CTF 17.4:

CEE	CTF 17.4
File formats: <ul style="list-style-type: none"> <li>• Default</li> <li>• Exactly as entered</li> <li>• Html source</li> </ul>	File formats: <ul style="list-style-type: none"> <li>• Text/Plain</li> <li>• Text/Html</li> </ul>
File Status: <ul style="list-style-type: none"> <li>• Draft</li> <li>• Reviewed</li> <li>• Baselined</li> <li>• Archival</li> <li>• Obsolete</li> <li>• Stable</li> </ul>	File Status: <ul style="list-style-type: none"> <li>• Draft</li> <li>• Final</li> </ul>
Pre-migration document states: Draft, Reviewed, Baselined, Archival or Obsolete	Post migration document state: Draft
Pre-migration document state: Stable	Post migration document state: Final
First document created date	Folder creation date
Folder name	Document Folder Name
Name	Document Name
Status	Status
Description	Description
Text	Create Online
Attachments	Upload from File
Link	Enter Document URL

### Notes

- The CEE features like infinite lock and a specific period are not mapped to CTF 17.4.
- While creating documents, version comments can be entered in CTF 17.4.
- In CEE, the file format can not be changed while editing the file, whereas, in CTF 17.4, the file format is editable as well.
- In CTF 17.4, version comments also can be specified.

- The CEE features that are not available in CTF 17.4 include: Advisory/Strict lock, Future date setting and Folder description.

### Migrating announcements

Here are the migration details and migration notes for announcements.

#### Migration Overview

- Announcements in CEE are migrated as News in CTF 17.4.
- Image URLs, publishing and event dates in CEE announcements are migrated as news content in CTF 17.4.
- If the Project Home page is configured via remote publishing (using index.html), the News component would not appear in the Home Page. However, the news component would appear in My Workspace in CTF 17.4.
- CEE project owner's message is migrated to the CTF 17.4 news component as project news.

This table shows how fields in CEE map to fields in CTF 17.4:

CEE	CTF 17.4
Announcements	News
Headline	Title – character limit is 250
Body	Body
Announcements Modified Date	News Creation Date
Image URLs, Event Dates	News - Body


#### Notes

- In CEE, Announcements is available as an independent tool, whereas, in CTF 17.4, News is available as a project page component.
- For public projects in CTF 17.4, news is displayed in Login page even after logout.
- Advanced options for announcements like publishing and event dates are not migrated to CTF 17.4. as CTF 17.4 doesn't support these features.
- News post headlines are migrated only up to the first 250 characters.
- Deleted announcements are not migrated from CEE to CTF 17.4.

### Migrating project dashboard

Here are the migration details and migration notes for project dashboard.

#### Migration Overview


- CEE's PD/Task is not migrated to CTF 17.4. However, the customers can access all the versions of imported and synchronized MPP files which they had in CEE. Post migration, any version of the file can be chosen and imported manually into the corresponding project in CTF 17.4.
- The MPP files of a project can be found under the "plans" directory of the respective project's CVS or SVN repository. The version number in the sync file's name corresponds to a repository revision of the plan file.
- The files will have the following naming convention:
  - Plan files: <domain name>\_<project name>\_plan.mpp
  - Sync files: <domain name>\_<project name>\_sync\_<version no.>.mpp
-  **Tip:** As an example, "mydomain\_myproject\_sync\_1.18.mpp" will be the sync file corresponding to the revision "1.18" of plan file, "mydomain\_myproject\_plan.mpp".
- Any revision of the plan file or a sync file can be retrieved from the repository.
- You can take a backup of the exception report from the following location: <source folder for cee dump and load directory>/backup/<domain name>/projects/<projectname>/projectname-REPORT.txt

- The global exception report can be accessed from: <source folder for cee dump and load directory>/backup/<domain name>/global/GLOBAL-REPORT-CONSOLIDATED.txt

### Migrating remote publishing tool

Here are the migration details and migration notes for remote publishing tool.

#### Migration Overview

- For PCE-enabled subversion projects, CTF Home Page is replaced with CEE Home Page.
- For CVS and non-PCE enabled subversion projects, CTF Home Page is replaced with remote publishing content (index.html).
-  **Note:** If PCE is enabled in CEE, by default the "Project Page" option appears selected in CTF's **Project Admin > Project Settings** page. If PCE is not enabled in CEE, the remote publishing content is displayed in CTF.
- The remote publishing content of CEE is migrated to the "Publishing" repository of CTF.

#### Notes

- In CTF, there is no equivalent functionality for governing the access to remote published content. Hence, the corresponding CEE permissions (resource patterns) are dropped and anyone with project view page permission is allowed to view the remote published content in CTF.

### Mapping CollabNet Enterprise Edition's Host or Domain Configure settings to CollabNet TeamForge 17.4

The mapping details are given in respective tables.

#### Note about SAS-70

To meet SAS-70 certification requirement, the following criteria need to be met:

- After 20 bad logons, account must lock
- Reset passwords must be changed on next logon
- Cannot reuse previous three passwords
- Passwords expire every 90 days
- Password must be at least 8 alphanumeric digits long
- Passwords must contain at least one of each of these:
  - At least one lower case letter
  - At least one upper case letter
  - At least one number (0-9)
- At least one special character (e.g. !@#%&\*()+) )

While migrating the CEE password settings to CTF, if the CEE setting is more liberal than the SAS-70 requirement, the SAS-70 minimum requirement setting is chosen; otherwise the CEE setting is taken.

Here's how the CTF site's password settings are selected:

- If 'Authentication failures before lockout' domain setting was set to  $\leq 0$  or  $> 20$ , then set CTF's LOGIN\_ATTEMPT\_LOCK token to 20. Same value otherwise.
- If 'Size of user's password history' was set to  $< 3$ , then set CTF's PASSWORD\_HISTORY\_AGE token to 3.
- If 'Size of user's password history' was set to  $> 10$ , then set CTF's PASSWORD\_HISTORY\_AGE token to 10.
- If 'Size of user's password history' was set between 3 and 10, then set CTF's PASSWORD\_HISTORY\_AGE to the CEE value.
- If 'Minimum password length' was set to  $< 8$ , then set MINIMUM\_PASSWORD\_LENGTH to 8. Same value otherwise.
- Set PASSWORD\_REQUIRES\_MIXED\_CASE, PASSWORD\_REQUIRES\_NON\_ALPHANUM and PASSWORD\_REQUIRES\_NUMBER to true regardless of CEE setting.
- Set ALLOW\_PASSWORD\_DICTIONARY\_WORD to false regardless of CEE setting.

- If 'Password expiration period for non-admins (days)' was set to  $\leq 0$  or  $> 90$ , then set history.age to 90 in ctf-pass.conf. Same value otherwise.

This table shows the mapping of CEE's Host or Domain Configure settings to CTF 17.4's "runtime-options.conf" file.

CEE's Host or Domain Configure Settings	CTF 17.4 Tokens	Comments
Size of user's password history	PASSWORD_HISTORY_AGE	The CEE value is transferred to CTF 17.4.
Authentication failures before lockout	LOGIN_ATTEMPT_LOCK	When the CEE value is "0", the CTF 17.4 value is set to "5", else the same CEE value is transferred.
Minimum password length	MINIMUM_PASSWORD_LENGTH	The CEE value is transferred to CTF 17.4.
Allow user join	DISABLE_USER_SELF_CREATION	When the CEE value is "true", the CTF 17.4 value is set to "false"; else it is set to "true".
Project path separator string	PROJECT_PATH_SEPARATOR	The CTF 17.4 value is hard coded to ">".
Suppress gathering of full name and organization	INCLUDE_ORGANIZATION_USER	When the CEE value is "true", the CTF 17.4 value is set to "false"; else it is set to "true".
Users can edit own organizations	ORGANIZATION_EDITABLE	The CEE value is transferred to CTF 17.4.
Number of required character classes	PASSWORD_REQUIRES_MIXED_CASE	When the CEE value is " $\leq 1$ " or "2", the CTF 17.4 value is set to "false"; else it is set to "true".
Number of required character classes	PASSWORD_REQUIRES_NON_ALPHANUM	When the CEE value is " $\leq 1$ " or "2" or "3", the CTF 17.4 value is set to "false"; else it is set to "true".
Number of required character classes	PASSWORD_REQUIRES_NUMBER	When the CEE value is " $\leq 1$ ", the CTF 17.4 value is set to "false"; else it is set to "true".
Reject passwords based on dictionary words	ALLOW_PASSWORD_DICTIONARY_WORDS	When the CEE value is "true", the CTF 17.4 value is set to "false"; else it is set to "true".
Password ticket expiration period for non-admins (days)	ACTIVATION_CODE_TIMEOUT	When the CEE value is " $\leq 0$ ", then set the CTF 17.4 value as 336. CEE value is multiplied by 24 (converting days to hours); else the same CEE value is transferred.
Password ticket expiration period for non-admins (days)	INITIAL_PASSWORD_CHANGE_ACTIVATION_CODE_TIMEOUT	When the CEE value is " $\leq 0$ ", then set the CTF 17.4 value as 336. CEE value is multiplied by 24 (converting days to hours); else the same CEE value is transferred.

CEE's Host or Domain Configure Settings	CTF 17.4 Tokens	Comments
Allow guest access/Allow anonymous access to domain home	ALLOW_ANONYMOUS_ACCESS	When the CEE value is "true", the CTF 17.4 value is set to "true"; else it is set to "false".

This table shows the mapping of CEE's Host or Domain Configure settings to CTF 17.4's "/etc/ctf-pass.conf" file:

CEE's Host or Domain Configure Settings	CTF 17.4 Properties	Comments
Password expiration period for non-admins (days)	history.age	The CEE value is transferred to CTF 17.4.
Users with non-expiring password	special.users	Same CEE value must be added to the existing users of CTF 17.4.
Password expiry warning period (days)	history.warn	The CEE value is transferred to CTF 17.4.

### How is global data handled during migration?

Global (or "framework") data is validated for integrity, correctness and completeness after migration from Enterprise Edition (CEE) to TeamForge 17.4 (CTF 17.4). Here are the detailed assumptions behind the validation process.

- This includes; category, project, project group, global announcements, global SCM settings, global project roles, project tracker's global data, users and user groups.
- As part of data validation, any mismatch of global value count between CEE and CTF is reported in "Global.log". Absence of global value reference in CTF tables is reported in "Global Validate.log".
- CEE's HAC\_AUDIT\_LOG\_ENTRY table is moved to cee\_backedup\_audit\_log.

### Migrating Issue Tracker

When Issue Tracker data is migrated from CollabNet Enterprise Edition to TeamForge, you will see these changes.

### Common features

Here's a list of what you had in Issue Tracker and what you now have instead in TeamForge tracker.

- Issue Tracker in CEE is migrated to a tracker called "Issue Tracker" in CTF.
- Each CEE issue becomes an artifact in 'Issue Tracker' tracker.
- Each CEE issue can be identified with its corresponding artifact, as issue Id appears as a prefix in the corresponding artifact's title. For example, [issue 1].

CEE field	CTF field
Summary	Title
First Description	Description

### Dissimilarities

Some behavior in TeamForge tracker is different from Issue Tracker functionality.

- In CTF, you can monitor an artifact even after you are no longer assigned to it. You must explicitly stop monitoring, if desired. This behavior is different from CEE.
- In CTF, only the project members can be part of monitoring and "Assigned To" lists.

CEE field	CTF behavior
QA Contact	Single-select field in CEE, but is made as multi-select field in CTF.
Cc and Dependencies (dependent/blocker issues) fields	History is not retained for both these fields.

### User permissions on a CEE and CTF 17.4 site

Each permission a user can have on a CollabNet Enterprise Edition site has an equivalent permission on a CollabNet TeamForge site, with a few exceptions.

### Notes

- For integrations on TeamForge, users with the "Domain - Administer" permission on CEE get the "SCM Admin" and "SWLA Admin" permissions.
- CEE's Guest user is not migrated.
- CEE's Registered User and Anonymous Guest domain roles are dropped in favor of Default Access Permissions.
- Permissions of CEE's Registered User role are transferred to CTF 17.4's project's Default Access Permissions. Project is created as public (User class: "All logged in users") if the guest access was turned ON in CEE, or as gated (User class: "All project members and all non-restricted users") otherwise.
- Permissions of CEE's Anonymous Guest role are transferred to CTF 17.4's public project's Default Access Permissions (User class: "All users") if the guest access was turned ON in CEE.
- CEE's Host users with Host Admin role are migrated as CTF 17.4's site Admins.
- CEE's other host users are migrated as non-administrators in CTF 17.4.
- The CEE's host roles that are dropped are captured in the exception report.
- CEE and CTF 17.4 permission mapping details are configurable.
- There are no permissions (actions) for CUBiT, they are just roles without any permissions.

### TeamForge 17.4 Permissions for Authentication Realms

#### Permission Mapping Details

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/Gain Details	Comments
Authentication Realm - Add		This feature is not available in TeamForge 17.4		There isn't any corresponding UI option. This is a configurable option in TeamForge 17.4.
Authentication Realm - Delete		This feature is not available in TeamForge 17.4		There isn't any corresponding UI option. This is a configurable option in TeamForge 17.4.
Authentication Realm - Edit		This feature is not available in TeamForge 17.4		There isn't any corresponding UI option. This is a configurable option in TeamForge 17.4.
Authentication Realm - Map User		This feature is not available in TeamForge 17.4		There isn't any corresponding UI option. This is a

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/Gain Details	Comments
				configurable option in TeamForge 17.4.
Authentication Realm - View		This feature is not available in TeamForge 17.4		There isn't any corresponding UI option. This is a configurable option in TeamForge 17.4.

#### TeamForge 17.4 Permissions for Category

##### Permission Mapping Details

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/Gain Details	Comments
Category - Add	Create/Edit/View	User gains permissions	The user will get the permissions to edit the name and description.	This is available through site-wide roles.

#### TeamForge 17.4 Permissions for Domain

##### Permission Mapping Details

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/Gain Details	Comments
Domain - Administer	System Tools Permissions	User loses permission	The user may lose some permissions.	Conditional mapping. This permission is granted if the user also has "Audit Log – View" and "Session – List" permissions.
Domain Page - View		User gains permission	Gains the ability to view the community page if the user does not have this permission.	This permission is used to control who can view the Community tab. As this is part of the out-of-the-box "Registered User" role, dropping this is safe in most of the cases.
Domain Reports - View		CollabNet Enterprise Edition permission is retained		This permission is not being used in CollabNet Enterprise Edition.

#### TeamForge 17.4 Permissions for License

##### Permission Mapping Details

<b>CollabNet Enterprise Edition Permission</b>	<b>TeamForge 17.4 Permission</b>	<b>Change in Behavior</b>	<b>Functionality Loss/ Gain Details</b>	<b>Comments</b>
Licenses - Administer		This feature is not available in TeamForge 17.4		TeamForge 17.4 does not have any such feature.

*TeamForge 17.4 Permissions for Project and Project Group*

Permission Mapping Details

<b>CollabNet Enterprise Edition Permission</b>	<b>TeamForge 17.4 Permission</b>	<b>Change in Behavior</b>	<b>Functionality Loss/ Gain Details</b>	<b>Comments</b>
Project - Approve	Projects: Projects Admin	User loses permission	The user loses the permission if the role does not have "Project – Delete" permission.	Conditional mapping. This permission is available through a site-wide role. It is granted only if the role also has the "Project – Delete" permission.
Project - Suggest		User gains permission	All users will be able to suggest projects	Project suggest is not controlled through permissions in TeamForge 17.4.
ProjectGroup - Add		This feature is not available in TeamForge 17.4		TeamForge 17.4 does not have any such feature.

*TeamForge 17.4 Permissions for Session*

Permission Mapping Details

<b>CollabNet Enterprise Edition Permission</b>	<b>TeamForge 17.4 Permission</b>	<b>Change in Behavior</b>	<b>Functionality Loss/ Gain Details</b>	<b>Comments</b>
Session - List		User loses permission	The user loses the ability if the user does not have the "Domain – Administer" permission.	This permission is handled in the mapping for "Domain - Administer".

*TeamForge 17.4 Permissions for Start Page*

Permission Mapping Details



CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/Gain Details	Comments
Start Page - View		User gains permission	The user will have access to <b>My Workspace</b> , but is unlikely since this is part of the "Registered User" role.	TeamForge does not use permissions for this page.

### TeamForge 17.4 Permissions for User

#### Permission Mapping Details

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/Gain Details	Comments
User - Add	Create/Edit/View	User loses permission	The user loses the ability if the role does not have the "User – Edit" permission.	Conditional mapping. This permission is granted if the user also has the "User – Edit" permission.
User - Delete	Delete/View	CollabNet Enterprise Edition permission is retained		
User - Edit		User loses permission	The user loses the ability if the role does not have the "User – Add" permission.	This is handled in the mapping for the "User – Add" permission.
User - Edit - Self	Pluggable: Edit user preferences	CollabNet Enterprise Edition permission is retained		
User - Edit - Self	TeamForge 17.4: Drop	User gains permission	Roles not having this permission get the ability	TeamForge 17.4 does not have any such feature.
User - List		User loses permission	The user loses the ability unless the user gets "Create/Edit/View" or "Delete/View" permissions.	In TeamForge 17.4, there isn't any separate cluster for the view permission.
User - View		User gains permission	All users will be able to see any other user in the system	TeamForge 17.4 doesn't have such permission control.
User - View - Self		User gains permission	All users will be able to see their profiles	TeamForge 17.4 doesn't have such permission control.

### TeamForge 17.4 Permissions for User Group

#### Permission Mapping Details

CollabNet Enterprise Edition Permission	TeamForge 17.4 Permission	Change in Behavior	Functionality Loss/ Gain Details	Comments
UserGroup - Add	Create/Edit/View	User loses permission	The user loses the ability if the role does not have "UserGroup – Edit" and "Add/Delete Element" permissions	Conditional mapping. This permission is granted if the role also has "UserGroup – Edit", "UserGroup Element - Add" and "UserGroup Element – Delete" permissions.
UserGroup - Delete	Delete/View	CollabNet Enterprise Edition permission is retained		
UserGroup - Edit	Create/Edit/View	User loses permission	The user loses the ability if the role does not have "UserGroup – Add", "UserGroup Element - Add" and "UserGroup Element – Delete" permissions.	Conditional mapping. This permission is granted if the role also has "UserGroup – Add", "UserGroup Element - Add" and "UserGroup Element – Delete" permissions.
UserGroup - List	View Only	CollabNet Enterprise Edition permission is retained		
UserGroup - View	View Only	User gains permission	The user gains additional access for Groups list.	
UserGroup Element - Add		User loses permission	The user loses the ability unless the role gets "Create/Edit/View" permissions.	
UserGroup Element - Delete		User loses permission	The user loses the ability unless the role gets "Create/Edit/View" permissions.	

## CollabNet TeamForge full-text search

Here is some detailed information to help you make the most of TeamForge's full-text search capabilities powered by the Apache Lucene full-text search engine.



**Note:** For more information about Apache Lucene, refer to [Apache Lucene Documentation](#).

### An overview of the TeamForge full-text search

Here is some detailed information about the TeamForge's full-text search.

### Searchable items in TeamForge

The following items are search-able in TeamForge.

- Discussion forums, posts and topics
- Documents and document folders
- File releases, packages, and FRS files
- News posts
- Visible project pages
- Projects
- Source code files and commits
- Tasks and task folders
- Trackers and tracker artifacts
- Users
- Wiki pages
- Integrated applications<sup>#</sup>

#### #Integrated applications

Integrated applications may or may not have search capabilities. Refer to the integrated application's documentation to know more about its search features and to appreciate how the search features of the integrated application and TeamForge differ.

#### Attached files

You can search the contents of attached files. For more information about supported document formats, [click here](#).

#### Notes

You may witness some delay for the TeamForge objects to appear in the search results and the extent of delay depends on the load on the server.

When you search, the contents of all the search-able fields (of an object) are collectively searched for matches. For instance, when you search for an artifact, the contents of the title, description and comment fields, put together, are searched. As an example, the search entry `CollabNet AND TeamForge` returns an artifact if the content of its title, description and comment, put together, has both the words "CollabNet" and "TeamForge". In other words, the word "CollabNet" could be in the artifact title and the word "TeamForge" could be in the artifact description (and not necessarily present on the same field).

If you search with multiple words, items containing any of the words in the search string are returned. For more information, refer to [multiple-terms search](#). On the other hand, if you want to find items where the words, say `CollabNet` and `TeamForge`, both appear, type `CollabNet AND TeamForge`. For more information, refer to [Boolean operators](#).

The TeamForge searches for full words. Use [wildcards](#) for partial word searches.

Search terms are case-insensitive. For example, if you search using the keyword "collabnet", pages that contain `COLLABNET`, `CollabNet` and `collabnet` are all returned.

#### TeamForge full-text search guidelines

Here is some guidelines to help you create effective searches.

#### Single term, multiple terms and phrase

Single-term search looks for all search results that match the search text. For example, a search entry of <code>doc</code> only returns search results of "doc".
---

Multiple-terms search looks for all search results that match any of the words in the search text. For example, a search entry of <code>document plan</code> returns search results of "document", "plan", and "document plan".
---

A group of words surrounded by double quotes, such as "product requirements", return only search results containing the entire phrase.

## Boolean operators

Terms and phrases can be combined with Boolean operators for more complex searches. Boolean operators must be in upper case. Use:

- OR between two terms returns search results containing either of the terms. This is the default operator used if no other operator is specified.
- AND between two terms returns only search results containing both of the terms.
- The + operator before a term makes the term required. Only search results containing the terms are returned.
- The - or NOT operator before a term returns only search results that do not contain the term. The character "-" represents the Boolean operator AND NOT.

### Tip:

You can group Boolean searches using parentheses. For example, (doc OR test) AND plan returns search results containing "doc plan" and "test plan".

## Wildcard searches

To look for search results with a single character replaced, use the ? symbol. For example, to look for search results with "text" or "test", enter te?t.

To look for search results with more than one character replaced, use the \* symbol. For example, to look for search results such as "content" or "contest" or "continuous" or "control", enter cont\*.

### Note:

You can use wildcard symbols in the middle or at the end of a search, but not as the first character of a search keyword.

## Fuzzy searches

To look for search results with spelling similar to the search term entered, use the ~ symbol as the last character of the search keyword. For example, to look for search results with spelling similar to "roam", enter roam~. This returns search results such as "roam" and "rooms".

## Special characters

If you have any of the following special characters in your search text, you must escape them by enclosing the entire phrase in double quotes. + - & | ! ( ) { } [ ] ^ " ~ \* ? : \

For example, to look for search results containing the hyphenated term "product-development", enter product-development.

The special character "+" represents the Boolean operator AND. The special character "-" represents the Boolean operator AND NOT.

## Regular expression search with forward slashes

Lucene 4 supports regular expression searches matching a pattern between forward slashes "/". For example, to look for search results containing the words "moat" or "boat", use the search string /[mb]oat/.

If you are specifically looking for search results containing a forward slash "/" character, you must backslash-escape or quote-escape the forward slash character. For example, to look for search results containing "<opt/collabnet>", use the search string <opt\/collabnet>.

### Excluded words

The following words are considered stop words and are not search-able on their own: a, an, and, are, as, at, be, but, by, for, if, in, into, is, it, no, not, of, on, or, s, such, that, the, their, then, there, these, they, this, to, was, will, with

### Range searches

You can do a range-bound search using the TO operator. For example, the search entry, [artf1100 TO artf1200], returns items containing values between artf1100 and artf1200, including artf1100 and artf1200. To exclude the upper and lower bounds from the search results, use curly brackets {} instead of square brackets [].

# Install / Upgrade

---

## Plan your installation/upgrade

---

Plan your installation or upgrade setup, hardware and software requirements and so on before you begin..

### TeamForge Services

Before you plan your installation or upgrade, let us understand TeamForge and its services.

TeamForge site consists of a core TeamForge application and several tightly integrated services that support it. In addition, you can integrate TeamForge with other third party applications such as Nexus, Jenkins, Jira and so on. Some of the TeamForge services are mandatory and some are optional. You can install the services, all in one single server, or distribute them across two or more servers.

- The core TeamForge application provides the Web interface that users see, and the API that other applications can interact with. It also includes the file system where some user content is stored, such as wiki pages.
- The site database is where most of the user-created content is stored and accessed. Documents, discussion posts, tracker artifacts, project administration settings: all that sort of thing lives in the database.
- The source control server ties any number of Subversion, Git/Gerrit or CVS repositories into the TeamForge site.
- The Extract Transform and Load (ETL) server pulls data from the site database and populates the datamart to generate charts and graphs about how people are using the site. The datamart (Reports DB) is an abstraction of the site database, optimized to support the reporting functionality.
- EventQ is a TeamForge capability that provides traceability for product life cycle activities such as work items, SCM commits, continuous integration (CI) builds, and code reviews.

### Service Name and site-options.conf syntax changes

The names of some of the TeamForge services have been renamed in TeamForge 17.1. In addition, the `site-options.conf` syntaxes for defining the TeamForge services running on a host (**host:SERVICES** token) and TeamForge domain name (**host:PUBLIC\_FQDN** token) have been changed in TeamForge 17.1.

Unless otherwise notified, to ensure backward compatibility, TeamForge 17.1 and later versions support both old and new syntaxes for defining your host and domain tokens. Similarly, TeamForge 17.4 supports both old and new service names. However, it is recommended to adjust your `site-options.conf` in line with these changes as support for older syntax and naming convention would be dropped in due course.

Here's a list of available TeamForge services.

Services	Mandatory/Optional	Old Name	Description
ctfcore	Mandatory	app	Main TeamForge application server
search	Mandatory	indexer	Indexing and Searching
mail	Mandatory	NA	Email server
ctfcore-database	Mandatory	database	Operational database
codesearch	Mandatory	codesearch	Code Search
etl	Optional	etl	ETL for Datamart
ctfcore-datamart	Mandatory if and only if you install <b>etl</b>	datamart	Datamart database
subversion	Optional	subversion	SVN Version Control
cvs	Optional	cvs	CVS Version Control

Services	Mandatory/Optional	Old Name	Description
gerrit	Optional	gerrit	Git/Gerrit Version Control
gerrit-database	Mandatory if and only if you install <b>gerrit</b>	NA	Database for Git/Gerrit. In a distributed setup, add this identifier on the server where you want to run Gerrit database.
binary	Optional	Optional	Artifact repository integration
binary-database	Mandatory if and only if you install <b>binary</b>	binary	Database for artifact repository integration. Binary app (binary) and database (binary-database) have to be installed on the same server.
reviewboard	Optional	reviewboard	Review Board code review tool
reviewboard-database	Mandatory if and only if you install <b>reviewboard</b>	NA	Review Board code review tool
reviewboard-adapter	Mandatory if and only if you install <b>reviewboard</b>	NA	Database for Review Board. In a distributed setup, reviewboard-adapter must always be installed on the TeamForge Application Server.
eventq	Mandatory	NA	EventQ application server. In a distributed setup, add this identifier on the server where you want to run EventQ application.
redis	Mandatory	NA	EventQ in-memory database/data structure server. In a distributed setup, add this identifier on the server where you want to run EventQ application.
mongodb	Mandatory	NA	EventQ database server. In a distributed setup, add this identifier on the server where you want to run EventQ database.
rabbitmq	Mandatory	NA	EventQ AMQP message server. In a distributed setup, add this identifier on the server where you want to run EventQ message queue.

These service identifiers are used in the `site-options.conf` file's **host:SERVICES** token.

### Single Server or Distributed Setup?

If you are installing TeamForge, are you planning to install on a single server or distribute TeamForge services across two or more servers? How are you going to distribute the services?

In the default setup, all services run on the same server as the main TeamForge application. But in practice, only the TeamForge application needs to run on the TeamForge application server. The other services can share that server or run on other servers, in almost any combination.

Assess your own site's particular use patterns and resources to decide how to distribute your services, if at all. For example, if you anticipate heavy use of your site, you will want to consider running the site database, the source control service, or the reporting engine on separate hardware to help balance the load.

When you distribute your services on multiple servers, you must do some configuration to handle communication among the services. Verify your basic networking setup. See [Set up Networking for TeamForge](#).

## PostgreSQL or Oracle?

PostgreSQL 9.6.1 is installed automatically when you install TeamForge 17.4. If you intend to use Oracle, CollabNet recommends that you let the installer run its course, make sure things work normally, and then set up your Oracle database and switch over to it.

If you want to use Oracle as your database, consider the following points:

- TeamForge 17.4 supports Oracle server 12c and Oracle client .
- Oracle express edition is not supported for both client and server.
- Review Board 2.5.6.1 was tested with PostgreSQL 9.6.1 only. Review Board with Oracle was not tested. Note that Review Board (application) must be installed on the TeamForge application server and therefore uses the PostgreSQL database that is installed by default during TeamForge installation. However, Review Board database can be installed on a separate server.
- GIT integration works only with PostgreSQL. The Git integration will use PostgreSQL even if your TeamForge site uses Oracle.

The efficiency of your database can have an impact on your users' perception of the site's usability. If your site uses a PostgreSQL database (which is the default), you may want to consider tuning it to fit your specific circumstances. The default settings are intended for a small-to-medium site running on a single server. See [What are the right PostgreSQL settings for my site?](#) on page 943 for recommendations from CollabNet's performance team on optimizing PostgreSQL for different conditions.

## Integrations

TeamForge supports integration with a wide array of third party applications such as Review Board, Git, Nexus, Jira and so on. As a customer you may or may not always want (or have) all of TeamForge's supported integrated applications. It's also quite possible that some of the integrated applications may not always run on all the platforms supported by TeamForge.

To accommodate a wider audience, by default, TeamForge install and upgrade instructions include steps to integrate such third party applications with TeamForge. However, use your discretion to ignore and skip such steps if they are not relevant to your site. See [TeamForge installation requirements](#) on page 441 to understand what it takes to run TeamForge 17.4 with integrations.

## One-hop upgrade compatibility

Though the TeamForge 17.4 installer supports one-hop upgrade from TeamForge 16.3 or later versions, TeamForge 17.4 upgrade instructions, in general, are for upgrading from TeamForge 17.1 (including update releases, if any) to TeamForge 17.4.

During an upgrade, the TeamForge 17.4 migration script detects the TeamForge version you run, checks if it's TeamForge 16.3 or later, and if 'true', proceeds with the data migration.

The migration script aborts with an error message if it detects TeamForge 7.2 or earlier versions. You must upgrade your site to TeamForge 16.3 or later and then upgrade to TeamForge 17.4.

## Hardware and Backup

If you aren't the person who first installed your current TeamForge site (or maybe, even if you are), it's essential to catalog the hosts where your services are running and to know what configuration has been applied to them.

While upgrading to a latest TeamForge version, you can choose to upgrade on the same hardware or on new hardware. In general, it is good to have a backup plan in place. Same hardware upgrades need no backup. However, it's recommended to take a back up as a measure of caution. See [Backup and Restore TeamForge and EventQ](#) for more information.



## Other Dos and Don'ts

Here's a list of dos, don'ts and points to remember when you install or upgrade TeamForge.

- Understand [TeamForge installation requirements](#) on page 441.
- Get your [TeamForge license](#) key and keep it handy.
- Installing or upgrading TeamForge needs root privileges. You must log on as root or use a root shell to install or upgrade TeamForge.
- Verify your basic networking setup before installing or upgrading TeamForge. See [Set up Networking for TeamForge](#).
- If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.
- SSL is enabled by default and a self-signed certificate is auto-generated. However, you can use a few `site-options.conf` tokens to adjust this behavior. To generate the SSL certificates, see [Generate SSL certificates](#) on page 718.
- Do not customize your operating system installation. Select only the default packages list.
- You need to add the SSL certificate to the Java keystore if the site uses a self-signed certificate. For instructions on adding the self-signed certificate to the Java keystore, see [Protect Integrations with SSL](#).
- For the ETL service to run as expected in a distributed TeamForge installation, all servers must have the same time zone.
- If you have Git integration on a separate server, both TeamForge and Git servers must have their time and date synchronized.
- While you can run both EventQ and TeamForge on the same server, CollabNet recommends such an approach only for testing purposes. It's always recommended to run EventQ on a separate server for optimal scalability.
- Look for new or modified `site-options.conf` tokens and update your `site-options.conf` file as required during the upgrade process. See [Site Options Change Log](#).
- Do not manually modify TeamForge-managed site option tokens such as the `AUTO_DATA` token. See [.](#)  for more information.
- No backup is required for same hardware upgrades. However, you can create a backup as a precaution. See [Backup and Restore TeamForge and EventQ](#) for more information.
- Always use compatible JDBC drivers meant for specific database versions. See [JDBC Drivers Reference](#) for more information. Also see: [Why do ETL jobs fail post TeamForge upgrade?](#) on page 962
- You can run the initial load job any time after the installation of TeamForge. We recommend that you run it before you hand over the site to the users. For more information, see [When do I run the initial load job?](#)
- Set up a TeamForge Stage Server before you upgrade your Production Server.
- Stop TeamForge services on all servers in a distributed setup while upgrading to TeamForge 17.4.
- Uninstall hot fixes and add-ons, if any, before you start the TeamForge 17.4 upgrade procedure.
- SOAP50 APIs and event handlers are no longer supported in TeamForge 16.10 and later. Use the latest TeamForge SOAP/REST APIs.

## TeamForge installation requirements

---

Here's what it takes to install and run TeamForge.

### Hardware requirements


The following table lists the CPU, RAM and JVM Heap Size recommendations for Small, Medium, Large and Extra-large sites.


	Small	Medium	Large	X-Large
Users	100	500	1000-5000	10000+
CPU	Octa-core	12-core	> 12-core	> 16-core

	Small	Medium	Large	X-Large
RAM	16GB	24GB	32GB	32GB
Jboss JVM Heap Size	1.5GB	3GB	6GB	> 8GB
Elasticsearch JVM Heap Size*	2GB	2GB	2GB	2GB
200 GB (or more) hard drive #				

The following table highlights the factors that can impact TeamForge performance. Numbers are indicative. Anything more than the prescribed numbers may impact the performance.

	Small	Medium	Large	X-Large
Artifacts	15000	70000	100000	100000
Flex Fields	25	50	100	100
Projects	20	80	500	500
Integrations	0-1	0-2	0-2	0-2
Integrated Applications	0-2	0-3	3+	3+

 **Important:** On Medium, Large and X-Large sites, it is highly recommended that you install the TeamForge Application, Database and SCM services on separate 64-bit servers based on the usage pattern.

 **Important:** Do not customize your operating system installation. Select only the default packages list.

**Operating system** RHEL 6.8 or 7.3 <sup>1</sup> (64 bit) | CentOS 6.8 or 7.3 (64 bit)

**Browser** Google Chrome 58  
Mozilla Firefox 53  
Microsoft IE 11  
Safari 10.1.

## Software requirements

\* You must have adequate RAM to accommodate the JVM heap requirements of Elasticsearch in addition to the JVM heap requirements of other components such as Jboss, integrated applications, and so on.

# The required hard drive capacity depends on the estimated amount of document and file release uploads.

<sup>1</sup> Red Hat Enterprise Linux servers must have access to the [Red Hat Network](#) or equivalent (satellite server, spacewalk, or RHN proxy).

TeamForge user interfaces are best viewed at screen resolution of at least 1280 x 800 (or more) pixels.

**Database** PostgreSQL 9.6.1 is installed by default when you install TeamForge 17.4. However, you can use Oracle if you want to. See [PostgreSQL or Oracle](#) for more information.

Oracle Server 12c

Oracle Client 12c

**Java**<sup>2</sup> Jboss/Wildfly 10.1.0.Final

JRE 1.8.0\_112

Tomcat 8.0.39

**Elasticsearch** 5.1.1

**Yum package manager** 3.2.29 or earlier versions

**TeamForge 17.4 supports the following software versions for integration**

**Review Board** 2.5.6.1<sup>3</sup>

**Git/Gerrit** 2.13.x

**Subversion** 1.8.17

**Subversion Edge** 5.0

**CVS** 1.11.x

**ViewVC** 1.1.24

**Nexus** 2.9.0 or later

**Artifactory Pro** 4.7

**Jira** 6.3-7.1

**TestLink** 1.9.15 and 1.9.16

### Microsoft applications

The following Microsoft applications have been tested with TeamForge:

- Microsoft Office XP (with Service Pack 3) on Windows XP Service Pack 2 and Windows 2000 Service Pack 4.
- Microsoft Office 2003 (with Service Pack 1) on Windows XP Service Pack 2 and Windows 2000 Service Pack 4.

<sup>2</sup> For larger installs, the max JVM heap size should be increased to 1024MB (or larger), depending on available resources on the server.

<sup>3</sup> It is recommended to install Review Board on the TeamForge application server.

## PostgreSQL or Oracle?

PostgreSQL 9.6.1 is installed automatically when you install TeamForge 17.4. If you intend to use Oracle, CollabNet recommends that you let the installer run its course, make sure things work normally, and then set up your Oracle database and switch over to it.

If you want to use Oracle as your database, consider the following points:


- TeamForge 17.4 supports Oracle server 12c and Oracle client 12c.
- Oracle express edition is not supported for both client and server.
- Review Board 2.5.6.1 was tested with PostgreSQL 9.6.1 only. Review Board with Oracle was not tested. Note that Review Board must be installed on the TeamForge application server and therefore uses the PostgreSQL database that is installed by default during TeamForge installation.
- GIT integration works only with PostgreSQL. The Git integration will use PostgreSQL even if your TeamForge site uses Oracle.

The efficiency of your database can have an impact on your users' perception of the site's usability. If your site uses a PostgreSQL database (which is the default), you may want to consider tuning it to fit your specific circumstances. The default settings are intended for a small-to-medium site running on a single server. See [What are the right PostgreSQL settings for my site?](#) on page 943 for recommendations from CollabNet's performance team on optimizing PostgreSQL for different conditions.

## Set up networking for your TeamForge server

---

After installing the operating system, prepare the networking connections and configuration that your TeamForge 17.4 site will require.


 **Note:** You must have root access to all the hosts you will be setting up for your site.

1. Use the NetworkManager to list the DNS servers you want to use for resolving Internet addresses.
2. Open the appropriate ports, and close all other ports.
3. Use the `hostname` command to verify that the machine name is resolvable on the network.

```
hostname
bigbox.supervillain.org
```

4. Use the `nslookup` command to verify that your hostname maps to the right IP address.

```
nslookup bigbox.supervillain.org
Server: 204.16.107.137
Address: 204.16.107.137#53
```

 **Tip:** If there is any doubt about what the system's real IP address is, use the `/sbin/ifconfig` command.

5. If you are installing behind a proxy, specify your proxy settings.

```
export http_proxy=http://
<PROXY_USERNAME>:<PROXY_PASSWD>@<PROXY_HOST>:<PROXY_PORT>
export no_proxy=localhost,127.0.0.0/8,<hostname>
```

6. Use a tool such as Nessus to scan your server for potential vulnerabilities.

(See [Which ports are required to be open for a TeamForge site?](#) on page 444 for detailed security recommendations.)

## Which ports are required to be open for a TeamForge site?

---

The components of a CollabNet TeamForge installation listen on a number of operating system ports.

A small subset must be exposed externally to enable users to access TeamForge services. Any port that is not absolutely needed must be closed.

You can select your open ports in one of these ways:

- Use the firewall configuration GUI tool that comes with your operating system. It's usually launched with a command like `system-config-selinux`.
- Open the `/etc/sysconfig/iptables` file and specify your open ports by hand.

## Ports open to the Internet

Open the following operating system level ports. All other ports must be firewalled off to maintain security.

- 👉 **Important:** Do not open port 7080 or port 8080 to the Internet. These ports are only for communications between the TeamForge application and the source code integration service, when those two site components are running on separate boxes.

### 22 (SSH)

Port 22 is the default port for the secure shell (SSH). This is required for basic SSH administrative functionality and for CVS, as all CVS transactions occur over SSH. If all Teamforge repositories are in SVN (the default for Teamforge), then this port should be closed to the public and only accessible to the system administrators.

If you have to expose SSH to the Internet, the best way to protect it is to require SSH keys and not allow password authentication, and do not permit root logins over SSH. If you must use local authentication for SSH, enforce regular password changes and password complexity.

- 👉 **Note:** If you have to expose SSH internally, limit access to the port to a bastion host if you can; otherwise limit it to specific trusted hosts or subnets.

### 25 (SMTP)

Port 25 is the default port for SMTP (email). CollabNet TeamForge discussion forums include mailing list functionality that allows users to send email to the TeamForge server. The James mail server included with TeamForge listens on port 25 to accept this mail for processing.

### 80 (HTTP)

Port 80 is the default port for Web data transfer. We strongly recommend that you set up SSL and use port 80 only to redirect to port 443.

### 443 (https)

Port 443 is the default port for encrypted Web data transfer (HTTPS). The Apache web server should be configured to encrypt all data so that it cannot be compromised by a third party with malicious intent. Apache can be configured to force all traffic to be sent over HTTPS, even when a request is sent via port 80 (HTTP).

TeamForge can help you take care of this, if you tell it to. See [Set up SSL for your TeamForge site](#) on page 717 for details.

### 29418 (Gerrit SSH)


Port 29418 is the default port which should be open for Gerrit SSH.

### Ports for internal use only

Open the `REPORTS_DATABASE_PORT` if you are granting direct access to the datamart from specific IPs using the `REPORTS_DB_ACCESS_HOSTS` `site-options.conf` token.

### Ports to be open in the firewall environment for TeamForge 17.4

Name	Source Box	Target Box	Port	Notes
Apache	ALL	TeamForge App	80 or 443	443 for ssl
TeamForge Database	TeamForge App	TeamForge Database	5432	
SVN Integration	TeamForge App	SVN	80 or 443	443 for ssl
GIT Integration	TeamForge App	GIT	80 or 443	443 for ssl
GIT ssh	ALL	GIT	29418	
Search	TeamForge App	Search	2099	
Binaries	TeamForge App	Binaries	8500	
Reports DB	TeamForge App	Reports DB	5432 or 5632	5432 is used by default as Reports DB is co-hosted with TeamForge database. 5632 can be used if you want Reports DB on a separate port.
Reports ETL	TeamForge App	Reports ETL	7010	
Code Search (Elasticsearch)	ALL	Code Search (Elasticsearch)	9200	

 **Note:** No manual port configuration is required for the following services if they are all installed on the TeamForge App server:

- TeamForge Database
- SVN Integration
- Search
- Binaries
- Reports DB
- Reports ETL

### Ports used by TeamForge EventQ services

Port	service	Host
8844	HTTP / HTTPS	App server
6379	Redis	App server
27017	MongoDB	DB server
28017	MongoDB HTTP console	DB server
5672	RabbitMQ	MQ server
15672	RabbitMQ management console	MQ server

## Ports to be open in a firewall environment for EventQ

The following use cases detail TeamForge EventQ's firewall/routing requirements. By default, end-user web access is proxied through the primary TeamForge web server. TeamForge EventQ adapters supply data using the MQ layer and therefore need access to the MQ server (default port 5672). There are also private access requirements between the various installed services as detailed below.

Port	From	To	Description
443/80	App server	TeamForge server	App communication with TeamForge server
8844	TeamForge server	App server	TeamForge communication with App server
5672	TeamForge EventQ Adapters	MQ server	Message communication between Adapters and MQ server
5672	App server	MQ server	App communication with MQ server
5672	TeamForge server	MQ server	TeamForge communication with MQ server
15672	ALL	MQ server	App administration of MQ server
27017	App server	DB server	App server communication with DB server
22	App server	MQ server	App ssh to MQ server, installation only
22	App server	DB server	App ssh to DB server, installation only

## TeamForge DevOps Lifecycle Manager port requirements

ActionHub uses the following ports by default.

- 5672: To communicate with RabbitMQ. (Daemon and Web Server both require)
- 5432: To communicate with PG Database (Daemon and Web Server both require)
- 8187: Port used by Tomcat for Web Server front end
- 8191: Port used by the Daemon

ActionFlow uses the following ports by default.

- 80: HTTP not exposed
- 443: HTTPS needs to be exposed to the client
- 5672: RabbitMQ not exposed

ActionDesigner uses the following port by default: 8001.

## Install or upgrade TeamForge

A system administrator provides the infrastructure that lets site administrators, project managers and members collaborate effectively.

## Install TeamForge

When you finish the sequence of tasks given below, you will have a working TeamForge 17.4 site tailored to the specific requirements of your user base and environment.

Installing TeamForge is not difficult, but it can be complex. To succeed, you should be familiar with the essential Linux commands and terminology. Most of the commands are given explicitly, but sometimes you have to substitute the values corresponding to your own specific environment.

You can install TeamForge in a dedicated mode. In dedicated installation mode, the server is fully dedicated to the TeamForge application. This makes the installation process very simple as the TeamForge installer takes care of configuring your Apache, Postgres, SELinux (if enabled) automatically without any manual intervention.

In a "Dedicated" install, it is advised not to install additional software applications other than TeamForge. It is also recommended not to run any unwanted services that are not used by TeamForge.

### Install TeamForge 17.4: All services on a single RHEL/CentOS 7.3 server

The easiest way to install TeamForge is to install it on a single server, dedicated to TeamForge taking the default configuration settings.

#### Before you begin:

*Plan your installation/upgrade* on page 438

#### Remember:

- TeamForge 17.4 supports both RHEL/CentOS 6.8 and 7.3. See *TeamForge installation requirements* on page 441
- For the ETL service to run as expected in a distributed TeamForge installation, all servers must have the same time zone.
- While you can run both EventQ and TeamForge on the same server, CollabNet recommends such an approach only for testing purposes. It's always recommended to run EventQ on a separate server for optimal scalability. See *EventQ installation requirements*.
- Installing or upgrading TeamForge needs root privileges. You must log on as root or use a root shell to install or upgrade TeamForge.

All the following services run on a single RHEL/CentOS 7.3 server (we call this server-01).

- TeamForge Application Server (ctfcore)
- Database Server (ctfcore-database and ctfcore-datamart)
- Mail Server (mail)
- Codesearch Server (codesearch)
- ETL Server (etl)
- Git Integration Server (gerrit and gerrit-database)
- SCM Integration Server (subversion and cvs)
- Search Server (search).
- TeamForge EventQ Server (eventq, mongodb, redis and rabbitmq)

1. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.

See *the Red Hat installation guide* for help.

2. Check your basic networking setup. See *Set up Networking* for more information.
3. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge.

If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.



```
getenforce
```

- b) If the output of the `getenforce` command is "permissive", continue with the next step. If not, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

#### 4. Configure your TeamForge installation repository.

##### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.  
 b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```


- c) Refresh your repository cache.

```
yum clean all
```

##### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit RPM package: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**


 **Note:** In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d)  **Note:** If the Red Hat Enterprise Linux/CentOS installation DVD is mounted already, skip the following instructions. If not, mount the DVD.

Mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
```

```
yum list apr
```


5. Install the TeamForge application packages.

```
yum install teamforge
```

6. Set up your site's master configuration file.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

a) Configure the services and domain name tokens.

 **Note:** cliserver and reviewboard-adapter are the new services added in TeamForge 17.4.

```
localhost:SERVICES=ctfcore ctfcore-database mail etl ctfcore-datamart
search subversion
cvs codesearch eventq mongodb redis rabbitmq cliserver
```

```
localhost:PUBLIC_FQDN=my.app.domain.com
```

b) Add Gerrit identifiers if you are installing Git.


```
localhost:SERVICES=ctfcore ctfcore-database mail etl ctfcore-datamart
search subversion
cvs codesearch eventq mongodb redis rabbitmq cliserver gerrit gerrit-
database
```

c) Add Binary identifiers if you are installing Nexus.

```
localhost:SERVICES=ctfcore ctfcore-database mail etl ctfcore-datamart
search subversion
cvs codesearch eventq mongodb redis rabbitmq cliserver gerrit gerrit-
database binary binary-database
```

d) Add Review Board identifiers if you are installing Review Board.

```
localhost:SERVICES=ctfcore ctfcore-database mail etl ctfcore-datamart
search subversion
cvs codesearch cliserver eventq mongodb redis rabbitmq gerrit gerrit-
database binary binary-database
reviewboard reviewboard-database reviewboard-adapter
```


e)  **Attention:** *SSL* is enabled by default and a self-signed certificate is auto-generated. Use the following tokens to adjust this behavior. To generate the SSL certificates, see [Generate SSL certificates](#) on page 718.

Have the custom SSL certificate and private key for custom SSL certificate in place and provide their absolute paths in these tokens. *SSL\_CHAIN\_FILE* (intermediate certificate) is optional.

```
SSL_CERT_FILE=
```

```
SSL_KEY_FILE=
```

```
SSL_CHAIN_FILE=
```

 **Important:** All SSL certificates including self-signed certificates are added automatically.

f) TeamForge 7.1 and later support automatic password creation. See [AUTO\\_DATA](#) on page 736 for more information.

g) If the token *REQUIRE\_PASSWORD\_SECURITY* is enabled, then set a value for the token, *PASSWORD\_CONTROL\_EFFECTIVE\_DATE*.



**Caution:** The Password Control Kit (PCK) disables, deletes or expires user accounts that don't meet the password security requirements starting from the date set for the `PASSWORD_CONTROL_EFFECTIVE_DATE` token. If a date is not set, the PCK disables, deletes or expires user accounts immediately. See [PASSWORD\\_CONTROL\\_EFFECTIVE\\_DATE](#) for more information.

- h) If the token `REQUIRE_RANDOM_ADMIN_PASSWORD` is already set to true, then set the token `ADMIN_EMAIL` with a valid email address.

```
ADMIN_EMAIL=root@{__APPLICATION_HOST__}
```

- i) If you have LDAP set up for external authentication, you must set the `REQUIRE_USER_PASSWORD_CHANGE` site options token to false.
- j) Make sure the PostgreSQL tokens in the `site-options.conf` file are set as recommended in the following topic: [What are the right PostgreSQL settings for my site?](#) on page 943
- k) Save the `site-options.conf` file.

## 7. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

## 8. Run the TeamForge post installation script. For more information, see [post-install.py](#).



**Important:** If the token `REQUIRE_USER_PASSWORD_CHANGE` is set to true, login to TeamForge user interface, change the admin password and then run the `post-install.py` script.

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

## 9. Restart TeamForge.

```
setenforce 1
/opt/collabnet/teamforge/bin/teamforge restart
```

## 10. Verify TeamForge installation.

- Reboot the server and make sure all services come up automatically at startup.
  - Log on to the TeamForge web application using the default Admin credentials.
    - Username: "admin"
    - Password: "admin"
  - Create a sample project. See [Create a TeamForge project](#)
  - Write a welcome message to your site's users. See [Create a site-wide broadcast](#).
- [Install Review Board](#)
  - [Run TeamForge in SELinux enabled mode](#)
  - [Install Memcached](#) on page 694
  - [Install EventQ Adapters](#)

## Install TeamForge 17.4: All services on a single RHEL/CentOS 6.8 server

The easiest way to install TeamForge is to install it on a single server, dedicated to TeamForge taking the default configuration settings.

### Before you begin:

[Plan your installation/upgrade](#) on page 438

### Remember:

- TeamForge 17.4 supports both RHEL/CentOS 6.8 and 7.3. See [TeamForge installation requirements](#) on page 441
- For the ETL service to run as expected in a distributed TeamForge installation, all servers must have the same time zone.
- While you can run both EventQ and TeamForge on the same server, CollabNet recommends such an approach only for testing purposes. It's always recommended to run EventQ on a separate server for optimal scalability. See [EventQ installation requirements](#).
- Installing or upgrading TeamForge needs root privileges. You must log on as root or use a root shell to install or upgrade TeamForge.

All the following services run on a single RHEL/CentOS 6.8 server (we call this server-01).

- TeamForge Application Server (ctfcore)
- Codesearch Server (codesearch)
- Mail Server (mail)
- Database Server (ctfcore-database and ctfcore-datamart)
- ETL Server (etl)
- Git Integration Server (gerrit and gerrit-database)
- SCM Integration Server (subversion and cvs)
- Search Server (search).
- TeamForge EventQ Server (eventq, mongodb, redis and rabbitmq)

1. Install Red Hat Enterprise Linux/CentOS 6.8 and log in as root.

- Don't customize your installation. Select only the default packages list.
- The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.
- See [the Red Hat installation guide](#) for help.

2. Check your basic networking setup. See [Set up Networking](#) for more information.

3. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge.


If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.

a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

b) If the output of the `getenforce` command is "permissive", continue with the next step. If not, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

4.  **Attention:** Do this if and only if you are upgrading TeamForge on RHEL/CentOS 6.8.

Delete the `python-crypto` package.

```
yum erase python-crypto
```

5. Configure your TeamForge installation repository.

#### **TeamForge installation repository configuration for sites with internet access**

- Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.
- Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

c) Refresh your repository cache.

```
yum clean all
```


#### **TeamForge installation repository configuration for sites without internet access**

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.

- Red Hat Enterprise Linux/CentOS 6.8 64 bit RPM package: **CTF-Disconnected-media-17.4.824-110.rhel6.x86\_64.rpm**

- b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- c)  **Note:** If the Red Hat Enterprise Linux/CentOS installation DVD is mounted already, skip the following instructions. If not, mount the DVD.

Mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- d) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- e) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

## 6. Install the following application packages.


- a) TeamForge: To install the TeamForge application packages run the following command:

```
yum install teamforge
```

## 7. Set up your site's master configuration file.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

- a) Configure the services and domain name tokens.

-  **Note:** cliserver and reviewboard-adapter are the new services added in TeamForge 17.4.

```
localhost:SERVICES=ctfcore ctfcore-database ctfcore-datamart mail etl
search codesearch
subversion cvs eventq mongodb redis rabbitmq cliserver
```

```
localhost:PUBLIC_FQDN = my.app.domain.com
```

- b) Add Gerrit identifiers if you are installing Git.

```
localhost:SERVICES=ctfcore ctfcore-database ctfcore-datamart mail etl
search codesearch
```


```
subversion cvs eventq mongodb redis rabbitmq cliserver gerrit gerrit-
database
```

- c) Add Binary identifiers if you are installing Nexus.

```
localhost:SERVICES=ctfcore ctfcore-database ctfcore-datamart mail etl
search codesearch
subversion cvs eventq mongodb redis rabbitmq cliserver gerrit gerrit-
database binary binary-database
```

- d) Add Review Board identifiers if you are installing Review Board.

```
localhost:SERVICES=ctfcore ctfcore-database ctfcore-datamart mail etl
search codesearch
subversion cvs eventq mongodb redis rabbitmq gerrit gerrit-database
binary binary-database cliserver
reviewboard reviewboard-database reviewboard-adapter
```


- e)  **Attention:** *SSL* is enabled by default and a self-signed certificate is auto-generated. Use the following tokens to adjust this behavior. To generate the SSL certificates, see [Generate SSL certificates](#) on page 718.

Have the custom SSL certificate and private key for custom SSL certificate in place and provide their absolute paths in these tokens. *SSL\_CHAIN\_FILE* (intermediate certificate) is optional.


```
SSL_CERT_FILE=
```

```
SSL_KEY_FILE=
```

```
SSL_CHAIN_FILE=
```

 **Important:** All SSL certificates including self-signed certificates are added automatically.

- f) TeamForge 7.1 and later support automatic password creation. See [AUTO\\_DATA](#) on page 736 for more information.
- g) If the token *REQUIRE\_PASSWORD\_SECURITY* is enabled, then set a value for the token, *PASSWORD\_CONTROL\_EFFECTIVE\_DATE*.

 **Caution:** The Password Control Kit (PCK) disables, deletes or expires user accounts that don't meet the password security requirements starting from the date set for the *PASSWORD\_CONTROL\_EFFECTIVE\_DATE* token. If a date is not set, the PCK disables, deletes or expires user accounts immediately. See [PASSWORD\\_CONTROL\\_EFFECTIVE\\_DATE](#) for more information.

- h) If the token [REQUIRE\\_RANDOM\\_ADMIN\\_PASSWORD](#) is already set to true, then set the token *ADMIN\_EMAIL* with a valid email address.

```
ADMIN_EMAIL=root@{__APPLICATION_HOST__}
```

- i) If you have LDAP set up for external authentication, you must set the "[REQUIRE\\_USER\\_PASSWORD\\_CHANGE](#)" site options token to false.
- j) Make sure the PostgreSQL tokens in the `site-options.conf` file are set as recommended in the following topic: [What are the right PostgreSQL settings for my site?](#) on page 943
- k) Save the `site-options.conf` file.

## 8. Deploy services.


```
/opt/collabnet/teamforge/bin/teamforge provision
```

 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

9. Run the TeamForge post installation script. For more information, see [post-install.py](#).

 **Important:** If the token `REQUIRE_USER_PASSWORD_CHANGE` is set to true, login to TeamForge user interface, change the admin password and then run the `post-install.py` script.

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

10. Restart TeamForge.

```
setenforce 1
/opt/collabnet/teamforge/bin/teamforge restart
```

11. Verify TeamForge installation.

- a) Reboot the server and make sure all services come up automatically at startup.
- b) Log on to the TeamForge web application using the default Admin credentials.
  - Username: "admin"
  - Password: "admin"
- c) Create a sample project. See [Create a TeamForge project](#)
- d) Write a welcome message to your site's users. See [Create a site-wide broadcast](#).

- [Install Review Board](#)
- [Run TeamForge in SELinux enabled mode](#) .
- [Install Memcached](#) on page 694
- [Install EventQ Adapters](#)

### Install TeamForge, EventQ, Database and Datamart on three separate servers

Three-server distributed setup with TeamForge, EventQ, and Database (including Datamart) services installed on three separate servers.

#### Dos and Don'ts

Check [this](#) list before installing or upgrading TeamForge.

#### Three Server Setup

Here's how the services are distributed across three separate RHEL/CentOS 7.3 servers.

Services	TeamForge Application Server (server-01)	EventQ Server (server-02)	Database Server (server-03)
ctfcore	✓		
mail	✓		
search	✓		
codesearch	✓		
etl	✓		
gerrit	✓		
gerrit-database			✓

Services	TeamForge Application Server (server-01)	EventQ Server (server-02)	Database Server (server-03)
reviewboard	✓		
reviewboard-database	✓		
<b>reviewboard-adapter</b>	✓		
subversion	✓		
cvs	✓		
binary	✓		
binary-database			✓
<b>cliserver</b>	✓		
eventq		✓	
mongodb		✓	
redis		✓	
rabbitmq		✓	
ctfcore-database			✓
ctfcore-datamart			✓

The TeamForge installer takes care of the database configurations. In this setup, we use a separate port for the reports database. By default, both the site database and the datamart use port 5432, but when heavy traffic is expected, it is a good idea to use a separate port(5632) for datamart.

### Do this step by step on TeamForge Application Server (server-01)

1. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.

See [the Red Hat installation guide](#) for help.

2. Check your basic networking setup. See [Set up Networking](#) for more information.
3. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge.

If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is "permissive", continue with the next step. If not, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

4. Configure your TeamForge installation repository.



**TeamForge installation repository configuration for sites with internet access**

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to /tmp.
- b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

**TeamForge installation repository configuration for sites without internet access**

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD.

The DVD contains the necessary software and utilities required for installing TeamForge without internet access. In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDROM
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

5. Install the TeamForge application packages.

```
yum install teamforge
```

**Do this step by step on Database Server (server-03)**

6. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.

See [the Red Hat installation guide](#) for help.

7. Check your basic networking setup. See [Set up Networking](#) for more information.
8. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge.

If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is "permissive", continue with the next step. If not, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

9. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.
- b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

#### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD.

The DVD contains the necessary software and utilities required for installing TeamForge without internet access. In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDROM
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
```

```
enabled=1
gpgcheck=0
```

f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```


10. Install the TeamForge database packages.

```
yum install teamforge-database
```


11. Set up your site's master configuration file.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

### host: SERVICES Token

 **Note:** cliserver and reviewboard-adapter are the new services added in TeamForge 17.4.

- ```
server-01:SERVICES=ctfcore mail etl search subversion cvs codesearch
cliserver gerrit
binary reviewboard reviewboard-database reviewboard-adapter
```
- ```
server-02:SERVICES=eventq mongodb redis rabbitmq
```
- ```
server-03:SERVICES=ctfcore-database ctfcore-datamart gerrit-database
binary-database
```

 **Note:** You may remove the identifiers of components you do not want. For example, remove binary and binary-database if you are not planning to install binary repository managers such as Nexus. See TeamForge services for more information.

Save the `site-options.conf` file.

For further customization of your site configuration:

### SSL Tokens

SSL is enabled by default and a self-signed certificate is auto-generated. Use the following tokens to adjust this behavior.

```
SSL_CERT_FILE=
```

```
SSL_KEY_FILE=
```

```
SSL_CHAIN_FILE=
```

- To generate the SSL certificates, see [Generate SSL certificates](#) on page 718.
- Have the custom SSL certificate and private key for custom SSL certificate in place and provide their absolute paths in these tokens. `SSL_CHAIN_FILE` (intermediate certificate) is optional.
- All SSL certificates including self-signed certificates are added automatically.

### Password Tokens

- TeamForge 7.1 and later support automatic password creation. See [AUTO\\_DATA](#) on page 736 for more information.
- If the token `REQUIRE_PASSWORD_SECURITY` is enabled, then set a value for the token, `PASSWORD_CONTROL_EFFECTIVE_DATE`.



**Warning:** The Password Control Kit (PCK) disables, deletes or expires user accounts that don't meet the password security requirements starting from the date set for the `PASSWORD_CONTROL_EFFECTIVE_DATE` token. If a date is not set, the PCK disables, deletes

or expires user accounts immediately. See [PASSWORD\\_CONTROL\\_EFFECTIVE\\_DATE](#) for more information.

- If the token [REQUIRE\\_RANDOM\\_ADMIN\\_PASSWORD](#) is already set to true, then set the token [ADMIN\\_EMAIL](#) with a valid email address.

```
ADMIN_EMAIL=root@{__APPLICATION_HOST__}
```

- If you have LDAP set up for external authentication, you must set the “[REQUIRE\\_USER\\_PASSWORD\\_CHANGE](#)” site options token to false.

### PostgreSQL Tokens and Settings

- Enable the [POSTGRES\\_INTERFACE](#) token in the `site-options.conf` file if the value of "ifconfig -a" is not `eth0/enp0*`.
- Make sure the PostgreSQL tokens in the `site-options.conf` file are set as recommended in the following topic: [What are the right PostgreSQL settings for my site?](#) on page 943

## 12. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

### Do this step by step on TeamForge Application Server (server-01)

13. Copy the `site-options.conf` file from the Database Server (server-03) to the TeamForge Application Server's `/opt/collabnet/teamforge/etc/` directory.

## 14. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

### Do this step by step on EventQ server (server-02)

15. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.

See [the Red Hat installation guide](#) for help.

16. Check your basic networking setup. See [Set up Networking](#) for more information.

17. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge.

If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is "permissive", continue with the next step. If not, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

## 18. Configure your TeamForge installation repository.

### TeamForge installation repository configuration for sites with internet access

- Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to /tmp.
- Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- Refresh your repository cache.

```
yum clean all
```

### TeamForge installation repository configuration for sites without internet access

- Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD.

The DVD contains the necessary software and utilities required for installing TeamForge without internet access. In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- Verify your yum configuration files.

```
yum list httpd
yum list apr
```

## 19. Install the EventQ packages.

```
yum install teamforge-eventq
```

## 20. Copy the site-options.conf file from the Database Server (server-03) to the EventQ Server's /opt/collabnet/teamforge/etc/ directory.

**21. Deploy services.**

```
/opt/collabnet/teamforge/bin/teamforge provision
```

**Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

**Do this step by step on TeamForge Application Server (server-01)****22. Run the TeamForge post installation script.** For more information, see [post-install.py](#).

**Important:** If the token `REQUIRE_USER_PASSWORD_CHANGE` is set to true, login to TeamForge user interface, change the admin password and then run the `post-install.py` script.

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

**23. Restart TeamForge.**

```
setenforce 1
/opt/collabnet/teamforge/bin/teamforge restart
```

**24. Verify TeamForge installation.**

- a) Reboot the server and make sure all services come up automatically at startup.
- b) Log on to the TeamForge web application using the default Admin credentials.
  - Username: "admin"
  - Password: "admin"
- c) Create a sample project. See [Create a TeamForge project](#)
- d) Write a welcome message to your site's users. See [Create a site-wide broadcast](#).

**Post Installation Tasks**

- [Install Review Board](#)
- [Run TeamForge in SELinux enabled mode](#) .
- [Install Memcached](#) on page 694
- [Install EventQ Adapters](#)

**Install TeamForge, Codesearch, and EventQ on separate servers**

Three-server distributed setup with TeamForge, Codesearch and EventQ services installed on three separate servers.

**Dos and Don'ts**

Check [this](#) list before installing or upgrading TeamForge.

**Three Server Setup**

Here's how the services are distributed across three separate RHEL/CentOS 7.3 servers.

| Services   | TeamForge Application Server (server-01) | EventQ Server (server-02) | Codesearch Server (server-03) |
|------------|------------------------------------------|---------------------------|-------------------------------|
| ctfcore    | ✓                                        |                           |                               |
| mail       | ✓                                        |                           |                               |
| search     | ✓                                        |                           |                               |
| codesearch |                                          |                           | ✓                             |

| Services                   | TeamForge Application Server (server-01) | EventQ Server (server-02) | CodeSearch Server (server-03) |
|----------------------------|------------------------------------------|---------------------------|-------------------------------|
| etl                        | ✓                                        |                           |                               |
| gerrit                     | ✓                                        |                           |                               |
| gerrit-database            | ✓                                        |                           |                               |
| reviewboard                | ✓                                        |                           |                               |
| reviewboard-database       | ✓                                        |                           |                               |
| <b>reviewboard-adapter</b> | ✓                                        |                           |                               |
| subversion                 | ✓                                        |                           |                               |
| cvs                        | ✓                                        |                           |                               |
| binary                     | ✓                                        |                           |                               |
| binary-database            | ✓                                        |                           |                               |
| <b>cliserver</b>           | ✓                                        |                           |                               |
| eventq                     |                                          | ✓                         |                               |
| mongodb                    |                                          | ✓                         |                               |
| redis                      |                                          | ✓                         |                               |
| rabbitmq                   |                                          | ✓                         |                               |
| ctfcore-database           | ✓                                        |                           |                               |
| ctfcore-datamart           | ✓                                        |                           |                               |

### Do this step by step on TeamForge Application Server (server-01)

1. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.

See [the Red Hat installation guide](#) for help.

2. Check your basic networking setup. See [Set up Networking](#) for more information.
3. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge.

If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is "permissive", continue with the next step. If not, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

#### 4. Configure your TeamForge installation repository.

##### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to /tmp.
- b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```


- c) Refresh your repository cache.

```
yum clean all
```

##### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit RPM package: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**

 **Note:** In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDROM
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

#### 5. Install the TeamForge application packages.


```
yum install teamforge
```

#### 6. Set up your site's master configuration file.


```
vi /opt/collabnet/teamforge/etc/site-options.conf
```



**host: SERVICES Token**

 **Note:** cliserver and reviewboard-adapter are the new services added in TeamForge 17.4.

- ```
server-01:SERVICES=ctfcore ctfcore-database ctfcore-datamart mail etl
search subversion cvs
cliserver gerrit gerrit-database binary binary-database reviewboard
reviewboard-database reviewboard-adapter
```
- ```
server-02:SERVICES=eventq mongodb redis rabbitmq
```
- ```
server-03:SERVICES=codesearch
```

 **Note:** You may remove the identifiers of components you do not want. For example, remove binary and binary-database if you are not planning to install binary repository managers such as Nexus. See TeamForge services for more information.

**host: PUBLIC\_FQDN Token**

```
server-01:PUBLIC_FQDN=my.app.domain.com
```

Save the `site-options.conf` file.

For further customization of your site configuration:

**SSL Tokens**

SSL is enabled by default and a self-signed certificate is auto-generated. Use the following tokens to adjust this behavior.

```
SSL_CERT_FILE=
```

```
SSL_KEY_FILE=
```

```
SSL_CHAIN_FILE=
```

- To generate the SSL certificates, see [Generate SSL certificates](#) on page 718.
- Have the custom SSL certificate and private key for custom SSL certificate in place and provide their absolute paths in these tokens. `SSL_CHAIN_FILE` (intermediate certificate) is optional.
- All SSL certificates including self-signed certificates are added automatically.

**Password Tokens**

- TeamForge 7.1 and later support automatic password creation. See [AUTO\\_DATA](#) on page 736 for more information.
- If the token `REQUIRE_PASSWORD_SECURITY` is enabled, then set a value for the token, `PASSWORD_CONTROL_EFFECTIVE_DATE`.



**Warning:** The Password Control Kit (PCK) disables, deletes or expires user accounts that don't meet the password security requirements starting from the date set for the `PASSWORD_CONTROL_EFFECTIVE_DATE` token. If a date is not set, the PCK disables, deletes or expires user accounts immediately. See [PASSWORD\\_CONTROL\\_EFFECTIVE\\_DATE](#) for more information.

- If the token `REQUIRE_RANDOM_ADMIN_PASSWORD` is already set to true, then set the token `ADMIN_EMAIL` with a valid email address.

```
ADMIN_EMAIL=root@{__APPLICATION_HOST__}
```

- If you have LDAP set up for external authentication, you must set the `“REQUIRE_USER_PASSWORD_CHANGE”` site options token to false.

### PostgreSQL Tokens and Settings

- Enable the `POSTGRES_INTERFACE` token in the `site-options.conf` file if the value of "ifconfig -a" is not `eth0/enp0*`.
- Make sure the PostgreSQL tokens in the `site-options.conf` file are set as recommended in the following topic: [What are the right PostgreSQL settings for my site?](#) on page 943

### 7. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

#### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

### Do this step by step on Codesearch Server (server-03)

#### 8. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.

See [the Red Hat installation guide](#) for help.

#### 9. Check your basic networking setup. See [Set up Networking](#) for more information.

#### 10. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.

- Verify if SELinux is running in enforcing mode.

```
getenforce
```

- If the output of the `getenforce` command is either "Disabled" or "Permissive", SELinux is already disabled.
- If not disabled, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

#### 11. Configure your TeamForge installation repository.

##### TeamForge installation repository configuration for sites with internet access

- Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.
- Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- Refresh your repository cache.

```
yum clean all
```

##### TeamForge installation repository configuration for sites without internet access

- Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 6.8 64 bit RPM package: **CTF-Disconnected-media-17.4.824-110.rhel6.x86\_64.rpm**

- Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- d) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDROM
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- e) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

12. Install the Codesearch packages.

```
yum install teamforge-codesearch
```

13. Copy the `site-options.conf` file from the TeamForge Application Server (server-01) to the Codesearch Server's `/opt/collabnet/teamforge/etc/` directory.

14. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



**Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

**Do this step by step on EventQ Server (server-02)**

15. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.

See [the Red Hat installation guide](#) for help.

16. Check your basic networking setup. See [Set up Networking](#) for more information.

17. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is either "Disabled" or "Permissive", SELinux is already disabled.

- c) If not disabled, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

18. Configure your TeamForge installation repository.

**TeamForge installation repository configuration for sites with internet access**

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.  
 b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD.

The DVD contains the necessary software and utilities required for installing TeamForge without internet access. In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

### 19. Install the EventQ packages.

```
yum install teamforge-eventq
```

### 20. Copy the site-options.conf file from the TeamForge Application Server (server-01) to the EventQ Server's /opt/collabnet/teamforge/etc/ directory.

### 21. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```




#### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

### Do this step by step on TeamForge Application Server (server-01)

22. Run the TeamForge post installation script. For more information, see [post-install.py](#).

 **Important:** If the token `REQUIRE_USER_PASSWORD_CHANGE` is set to true, login to TeamForge user interface, change the admin password and then run the `post-install.py` script.

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

23. Restart TeamForge.

```
setenforce 1
/opt/collabnet/teamforge/bin/teamforge restart
```

24. Verify TeamForge installation.

- a) Reboot the server and make sure all services come up automatically at startup.
- b) Log on to the TeamForge web application using the default Admin credentials.
  - Username: "admin"
  - Password: "admin"
- c) Create a sample project. See [Create a TeamForge project](#)
- d) Write a welcome message to your site's users. See [Create a site-wide broadcast](#).

### Post Installation Tasks

- [Install Review Board](#)
- [Run TeamForge in SELinux enabled mode](#) .
- [Install Memcached](#) on page 694
- [Install EventQ Adapters](#)

### Install TeamForge, SCM and Git integration, and EventQ on Three separate servers

Three-server distributed setup with TeamForge (including Database and Datamart), SCM (Subversion and CVS) and Git, and EventQ services installed on three separate servers.

#### Dos and Don'ts

Check [this](#) list before installing or upgrading TeamForge.

#### Three Server Setup

Here's how the services are distributed across three separate RHEL/CentOS 7.3 servers.

Services	TeamForge Application Server (server-01)	EventQ Server (server-02)	SCM and Git Integration Server (server-03)
ctfcore	✓		
mail	✓		
search	✓		
codesearch	✓		
etl	✓		
gerrit			✓
gerrit-database			✓

Services	TeamForge Application Server (server-01)	EventQ Server (server-02)	SCM and Git Integration Server (server-03)
reviewboard	✓		
reviewboard-database	✓		
<b>reviewboard-adapter</b>	✓		
subversion			✓
cvs			✓
binary	✓		
binary-database	✓		
<b>cliserver</b>	✓		
eventq		✓	
mongodb		✓	
redis		✓	
rabbitmq		✓	
ctfcore-database	✓		
ctfcore-datamart	✓		

### Do this step by step on TeamForge Application Server (server-01)

1. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.

See [the Red Hat installation guide](#) for help.

2. Check your basic networking setup. See [Set up Networking](#) for more information.
3. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge.

If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is "permissive", continue with the next step. If not, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

4. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.
- b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```


- c) Refresh your repository cache.

```
yum clean all
```

### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**

 **Note:** In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, umount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```


5. Install the TeamForge application packages.

```
yum install teamforge
```

6. Set up your site's master configuration file.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```


### host: SERVICES Token

 **Note:** cliserver and reviewboard-adapter are the new services added in TeamForge 17.4.

- ```
server-01:SERVICES=ctfcore ctfcore-database ctfcore-datamart mail etl
search codesearch
```

```
cliserver binary binary-database reviewboard reviewboard-database
reviewboard-adapter
```

- `server-02:SERVICES=eventq mongodb redis rabbitmq`
- `server-03:SERVICES=subversion cvs gerrit gerrit-database`

 **Note:** You may remove the identifiers of components you do not want. For example, remove `binary` and `binary-database` if you are not planning to install binary repository managers such as Nexus.

#### host: PUBLIC\_FQDN Token

```
server-01:PUBLIC_FQDN=my.app.domain.com
```

```
server-03:PUBLIC_FQDN=my.scmangit.domain.com
```

Save the `site-options.conf` file.

For further customization of your site configuration:

#### SSL Tokens

SSL is enabled by default and a self-signed certificate is auto-generated. Use the following tokens to adjust this behavior.

```
SSL_CERT_FILE=
```

```
SSL_KEY_FILE=
```

```
SSL_CHAIN_FILE=
```

- To generate the SSL certificates, see [Generate SSL certificates](#) on page 718.
- Have the custom SSL certificate and private key for custom SSL certificate in place and provide their absolute paths in these tokens. `SSL_CHAIN_FILE` (intermediate certificate) is optional.
- All SSL certificates including self-signed certificates are added automatically.

#### Password Tokens

- TeamForge 7.1 and later support automatic password creation. See [AUTO\\_DATA](#) on page 736 for more information.
- If the token `REQUIRE_PASSWORD_SECURITY` is enabled, then set a value for the token, `PASSWORD_CONTROL_EFFECTIVE_DATE`.



**Warning:** The Password Control Kit (PCK) disables, deletes or expires user accounts that don't meet the password security requirements starting from the date set for the `PASSWORD_CONTROL_EFFECTIVE_DATE` token. If a date is not set, the PCK disables, deletes or expires user accounts immediately. See [PASSWORD\\_CONTROL\\_EFFECTIVE\\_DATE](#) for more information.

- If the token `REQUIRE_RANDOM_ADMIN_PASSWORD` is already set to true, then set the token `ADMIN_EMAIL` with a valid email address.

```
ADMIN_EMAIL=root@{__APPLICATION_HOST__}
```

- If you have LDAP set up for external authentication, you must set the `"REQUIRE_USER_PASSWORD_CHANGE"` site options token to false.

#### PostgreSQL Tokens and Settings

- Enable the `POSTGRES_INTERFACE` token in the `site-options.conf` file if the value of `"ifconfig -a"` is not `eth0/enp0*`.



- Make sure the PostgreSQL tokens in the `site-options.conf` file are set as recommended in the following topic: [What are the right PostgreSQL settings for my site?](#) on page 943

## 7. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

### **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

## Do this step by step on SCM and Git Integration Server (server-03)

### 8. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.

See [the Red Hat installation guide](#) for help.

### 9. Check your basic networking setup. See [Set up Networking](#) for more information.

### 10. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.

- Verify if SELinux is running in enforcing mode.

```
getenforce
```

- If the output of the `getenforce` command is either "Disabled" or "Permissive", SELinux is already disabled.
- If not disabled, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

### 11. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

- Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.
- Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```


- Refresh your repository cache.

```
yum clean all
```

#### TeamForge installation repository configuration for sites without internet access

- Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**

 **Note:** In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

12. Install the TeamForge SCM and Git packages.

```
yum install teamforge-scm teamforge-git
```

13. Copy the `site-options.conf` file from the TeamForge Application Server to the SCM and Git Integration Server's `/opt/collabnet/teamforge/etc/` directory.

14. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



**Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

During teamforge provision, the Register SCM integration process fails on sites that use self-signed certificates. Perform these steps in such cases.

- a) Restart JBoss on the TeamForge Application server.

```
/opt/collabnet/teamforge/bin/teamforge restart -s jboss
```

- b) Reinitialize TeamForge on the SCM and Git server.

```
/opt/collabnet/teamforge/bin/teamforge reinitialize
```

**Do this step by step on EventQ Server (server-02)**

15. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.

See [the Red Hat installation guide](#) for help.

16. Check your basic networking setup. See [Set up Networking](#) for more information.

17. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is either "Disabled" or "Permissive", SELinux is already disabled.

- c) If not disabled, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

## 18. Configure your TeamForge installation repository.

### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to /tmp.  
b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD.

The DVD contains the necessary software and utilities required for installing TeamForge without internet access. In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/  
mkdir cdrom  
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]  
name=RHEL CDRom  
baseurl=file:///media/cdrom/Server/  
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release  
enabled=1  
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd  
yum list apr
```

## 19. Install the EventQ packages.

```
yum install teamforge-eventq
```

20. Copy the `site-options.conf` file from the TeamForge Application Server (server-01) to the EventQ Server's `/opt/collabnet/teamforge/etc/` directory.

21. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```


 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

**Do this step by step on TeamForge Application Server (server-01)**

22. Run the TeamForge post installation script. For more information, see [post-install.py](#).

 **Important:** If the token `REQUIRE_USER_PASSWORD_CHANGE` is set to true, login to TeamForge user interface, change the admin password and then run the `post-install.py` script.

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

23. Restart TeamForge.

```
setenforce 1
/opt/collabnet/teamforge/bin/teamforge restart
```

24. Verify TeamForge installation.

- a) Reboot the server and make sure all services come up automatically at startup.
- b) Log on to the TeamForge web application using the default Admin credentials.
  - Username: "admin"
  - Password: "admin"
- c) Create a sample project. See [Create a TeamForge project](#)
- d) Write a welcome message to your site's users. See [Create a site-wide broadcast](#).

### Post Installation Tasks

- [Install Review Board](#)
- [Run TeamForge in SELinux enabled mode](#) .
- [Install Memcached](#) on page 694
- [Install EventQ Adapters](#)

### Install TeamForge, Database, SCM, and EventQ on Four separate servers

Four-server distributed setup with TeamForge, Database (including Datamart), SCM (Subversion CVS and Git) and EventQ on four separate servers.

#### Dos and Don'ts

Check [this](#) list before installing or upgrading TeamForge.

#### Four Server Setup

Here's how the services are distributed across four separate RHEL/CentOS 7.3 servers.

| Services | TeamForge Application Server (server-01) | EventQ Server (server-02) | Database Server (server-03) | SCM and Git Integration Server (server-04) |
|----------|------------------------------------------|---------------------------|-----------------------------|--------------------------------------------|
| ctfcore  | ✓                                        |                           |                             |                                            |
| mail     | ✓                                        |                           |                             |                                            |

| Services                   | TeamForge Application Server (server-01) | EventQ Server (server-02) | Database Server (server-03) | SCM and Git Integration Server (server-04) |
|----------------------------|------------------------------------------|---------------------------|-----------------------------|--------------------------------------------|
| search                     | ✓                                        |                           |                             |                                            |
| codesearch                 | ✓                                        |                           |                             |                                            |
| etl                        | ✓                                        |                           |                             |                                            |
| gerrit                     |                                          |                           |                             | ✓                                          |
| gerrit-database            |                                          |                           |                             | ✓                                          |
| reviewboard                | ✓                                        |                           |                             |                                            |
| reviewboard-database       | ✓                                        |                           |                             |                                            |
| <b>reviewboard-adpater</b> | ✓                                        |                           |                             |                                            |
| subversion                 |                                          |                           |                             | ✓                                          |
| cvs                        |                                          |                           |                             | ✓                                          |
| binary                     | ✓                                        |                           |                             |                                            |
| binary-database            | ✓                                        |                           |                             |                                            |
| <b>cliserver</b>           | ✓                                        |                           |                             |                                            |
| eventq                     |                                          | ✓                         |                             |                                            |
| mongodb                    |                                          | ✓                         |                             |                                            |
| redis                      |                                          | ✓                         |                             |                                            |
| rabbitmq                   |                                          | ✓                         |                             |                                            |
| ctfcore-database           |                                          |                           | ✓                           |                                            |
| ctfcore-datamart           |                                          |                           | ✓                           |                                            |

### Do this step by step on TeamForge Application Server (server-01)

1. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.

See [the Red Hat installation guide](#) for help.

2. Check your basic networking setup. See [Set up Networking](#) for more information.
3. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge.

If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is "permissive", continue with the next step. If not, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

#### 4. Configure your TeamForge installation repository.

##### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.  
b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```


- c) Refresh your repository cache.

```
yum clean all
```

##### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit RPM package: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**

 **Note:** In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/  
mkdir cdrom  
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd  
yum list apr
```

5. Install the following application packages.

- a) TeamForge: To install the TeamForge application packages run the following command:

```
yum install teamforge
```

**Do this step by step on Database Server (server-03)**

6. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.

See [the Red Hat installation guide](#) for help.

7. Check your basic networking setup. See [Set up Networking](#) for more information.

8. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is either "Disabled" or "Permissive", SELinux is already disabled.

- c) If not disabled, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

9. Configure your TeamForge installation repository.

**TeamForge installation repository configuration for sites with internet access**

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to /tmp.  
b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

**TeamForge installation repository configuration for sites without internet access**

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**



**Note:** In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/  
mkdir cdrom  
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```


10. Install the TeamForge database packages.

```
yum install teamforge-database
```


11. Set up your site's master configuration file.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

**host: SERVICES Token**

 **Note:** cliserver and reviewboard-adapter are the new services added in TeamForge 17.4.

- server-01:SERVICES=ctfcore mail etl search codesearch cliserver binary binary-database reviewboard reviewboard-database reviewboard-adapter
- server-02:SERVICES=eventq mongodb redis rabbitmq
- server-03:SERVICES=ctfcore-database ctfcore-datamart
- server-04:SERVICES=subversion cvs gerrit gerrit-database

 **Note:** You may remove the identifiers of components you do not want. For example, remove binary and binary-database if you are not planning to install binary repository managers such as Nexus. See TeamForge services for more information.

**host: PUBLIC\_FQDN Token**

```
server-01:PUBLIC_FQDN=my.app.domain.com
```

```
server-04:PUBLIC_FQDN=my.scmangit.domain.com
```

Save the site-options.conf file.

For further customization of your site configuration:

**SSL Tokens**

SSL is enabled by default and a self-signed certificate is auto-generated. Use the following tokens to adjust this behavior.

```
SSL_CERT_FILE=
```

```
SSL_KEY_FILE=
```

```
SSL_CHAIN_FILE=
```



- To generate the SSL certificates, see [Generate SSL certificates](#) on page 718.
- Have the custom SSL certificate and private key for custom SSL certificate in place and provide their absolute paths in these tokens. `SSL_CHAIN_FILE` (intermediate certificate) is optional.
- All SSL certificates including self-signed certificates are added automatically.

### Password Tokens

- TeamForge 7.1 and later support automatic password creation. See [AUTO\\_DATA](#) on page 736 for more information.
- If the token `REQUIRE_PASSWORD_SECURITY` is enabled, then set a value for the token, `PASSWORD_CONTROL_EFFECTIVE_DATE`.



**Warning:** The Password Control Kit (PCK) disables, deletes or expires user accounts that don't meet the password security requirements starting from the date set for the `PASSWORD_CONTROL_EFFECTIVE_DATE` token. If a date is not set, the PCK disables, deletes or expires user accounts immediately. See [PASSWORD\\_CONTROL\\_EFFECTIVE\\_DATE](#) for more information.

- If the token `REQUIRE_RANDOM_ADMIN_PASSWORD` is already set to true, then set the token `ADMIN_EMAIL` with a valid email address.

```
ADMIN_EMAIL=root@{__APPLICATION_HOST__}
```

- If you have LDAP set up for external authentication, you must set the `“REQUIRE_USER_PASSWORD_CHANGE”` site options token to false.

### PostgreSQL Tokens and Settings

- Enable the `POSTGRES_INTERFACE` token in the `site-options.conf` file if the value of "ifconfig -a" is not `eth0/enp0*`.
- Make sure the PostgreSQL tokens in the `site-options.conf` file are set as recommended in the following topic: [What are the right PostgreSQL settings for my site?](#) on page 943

## 12. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

### Do this step by step on TeamForge Application Server (server-01)

13. Copy the `site-options.conf` file from the Database Server to the TeamForge Application Server's `/opt/collabnet/teamforge/etc/` directory.

## 14. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

### Do this step by step on the SCM and Git Integration Server (server-04)

15. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.

See [the Red Hat installation guide](#) for help.

16. Check your basic networking setup. See [Set up Networking](#) for more information.

17. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.

a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

b) If the output of the `getenforce` command is either "Disabled" or "Permissive", SELinux is already disabled.

c) If not disabled, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

18. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to /tmp.

b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```


c) Refresh your repository cache.

```
yum clean all
```

#### TeamForge installation repository configuration for sites without internet access

a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**

 **Note:** In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDROM
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
```

```
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

19. Install the TeamForge SCM and Git packages.

```
yum install teamforge-scm teamforge-git
```

20. Copy the `site-options.conf` file from the Database Server (server-03) to the SCM and Git Integration Server's `/opt/collabnet/teamforge/etc/` directory.

21. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

During teamforge provision, the Register SCM integration process fails on sites that use self-signed certificates. Perform these steps in such cases.

- a) Restart JBoss on the TeamForge Application server.

```
/opt/collabnet/teamforge/bin/teamforge restart -s jboss
```

- b) Reinitialize TeamForge on the SCM server.

```
/opt/collabnet/teamforge/bin/teamforge reinitialize
```

**Do this step by step on EventQ server (server-02)**

22. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.

See [the Red Hat installation guide](#) for help.

23. Check your basic networking setup. See [Set up Networking](#) for more information.

24. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is either "Disabled" or "Permissive", SELinux is already disabled.

- c) If not disabled, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

25. Configure your TeamForge installation repository.

**TeamForge installation repository configuration for sites with internet access**

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.  
b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

**TeamForge installation repository configuration for sites without internet access**

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.
- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
  - In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.
- b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.centos.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD.

The DVD contains the necessary software and utilities required for installing TeamForge without internet access. In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

## 26. Install the EventQ packages.

```
yum install teamforge-eventq
```

## 27. Copy the site-options.conf file from the Database Server (server-03) to the EventQ Server's /opt/collabnet/teamforge/etc/ directory.

## 28. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```




### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

## Do this step by step on TeamForge Application Server (server-01)

## 29. Run the TeamForge post installation script. For more information, see [post-install.py](#).

-  **Important:** If the token `REQUIRE_USER_PASSWORD_CHANGE` is set to true, login to TeamForge user interface, change the admin password and then run the `post-install.py` script.

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

### 30. Restart TeamForge.

```
setenforce 1
/opt/collabnet/teamforge/bin/teamforge restart
```

### 31. Verify TeamForge installation.

- a) Reboot the server and make sure all services come up automatically at startup.
- b) Log on to the TeamForge web application using the default Admin credentials.
  - Username: "admin"
  - Password: "admin"
- c) Create a sample project. See [Create a TeamForge project](#)
- d) Write a welcome message to your site's users. See [Create a site-wide broadcast](#).

### Post Installation Tasks

- [Install Review Board](#)
- [Run TeamForge in SELinux enabled mode](#) .
- [Install Memcached](#) on page 694
- [Install EventQ Adapters](#)

### Install TeamForge, Oracle database, and EventQ on Three separate servers

Three-server distributed setup with TeamForge, Oracle Database (including Datamart) and EventQ on three separate servers.

#### Dos and Don'ts

Check [this](#) list before installing or upgrading TeamForge.

#### Three Server Setup

Here's how the services are distributed across three separate RHEL/CentOS 7.3 servers.

| Services                   | TeamForge Application Server (server-01) | EventQ Server (server-02) | Database Server (server-03) |
|----------------------------|------------------------------------------|---------------------------|-----------------------------|
| ctfcore                    | ✓                                        |                           |                             |
| mail                       | ✓                                        |                           |                             |
| search                     | ✓                                        |                           |                             |
| codesearch                 | ✓                                        |                           |                             |
| etl                        | ✓                                        |                           |                             |
| gerrit                     | ✓                                        |                           |                             |
| gerrit-database            | ✓                                        |                           |                             |
| reviewboard                | ✓                                        |                           |                             |
| reviewboard-database       | ✓                                        |                           |                             |
| <b>reviewboard-adapter</b> | ✓                                        |                           |                             |

| Services         | TeamForge Application Server (server-01) | EventQ Server (server-02) | Database Server (server-03) |
|------------------|------------------------------------------|---------------------------|-----------------------------|
| subversion       | ✓                                        |                           |                             |
| cvs              | ✓                                        |                           |                             |
| binary           | ✓                                        |                           |                             |
| binary-database  | ✓                                        |                           |                             |
| <b>cliserver</b> | ✓                                        |                           |                             |
| eventq           |                                          | ✓                         |                             |
| mongodb          |                                          | ✓                         |                             |
| redis            |                                          | ✓                         |                             |
| rabbitmq         |                                          | ✓                         |                             |
| ctfcore-database |                                          |                           | ✓                           |
| ctfcore-datamart |                                          |                           | ✓                           |

### Do this step by step on TeamForge Application Server (server-01)

1. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.

See [the Red Hat installation guide](#) for help.

2. Check your basic networking setup. See [Set up Networking](#) for more information.
3. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge.

If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is "permissive", continue with the next step. If not, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

4. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.
- b) Install the repository package.


```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

#### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.
- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**

 **Note:** In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDROM
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

5. Install the TeamForge application packages.

```
yum install teamforge
```


6. Rename the sample site configuration file from the installation package.

```
cd /opt/collabnet/teamforge/etc/
cp site-options-oracle.conf site-options.conf
```

7. Set up your site's master configuration file.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```


#### host: SERVICES Token

 **Note:** cliserver and reviewboard-adapter are the new services added in TeamForge 17.4.

- ```
server-01:SERVICES=ctfcore mail etl search subversion cvs codesearch
cliserver gerrit
```

```
gerrit-database binary binary-database reviewboard reviewboard-database
reviewboard-adapter
```

- `server-02:SERVICES=eventq mongodb redis rabbitmq`
- `server-03:SERVICES=ctfcore-database ctfcore-datamart`

 **Note:** You may remove the identifiers of components you do not want. For example, remove `binary` and `binary-database` if you are not planning to install binary repository managers such as Nexus. See TeamForge services for more information.

### host: PUBLIC\_FQDN Token

```
server-01:PUBLIC_FQDN=my.app.domain.com
```

Save the `site-options.conf` file.

For further customization of your site configuration:

### SSL Tokens

SSL is enabled by default and a self-signed certificate is auto-generated. Use the following tokens to adjust this behavior.

```
SSL_CERT_FILE=
```

```
SSL_KEY_FILE=
```

```
SSL_CHAIN_FILE=
```

- To generate the SSL certificates, see [Generate SSL certificates](#) on page 718.
- Have the custom SSL certificate and private key for custom SSL certificate in place and provide their absolute paths in these tokens. `SSL_CHAIN_FILE` (intermediate certificate) is optional.
- All SSL certificates including self-signed certificates are added automatically.

### Password Tokens

- TeamForge 7.1 and later support automatic password creation. See [AUTO\\_DATA](#) on page 736 for more information.
- If the token `REQUIRE_PASSWORD_SECURITY` is enabled, then set a value for the token, `PASSWORD_CONTROL_EFFECTIVE_DATE`.



**Warning:** The Password Control Kit (PCK) disables, deletes or expires user accounts that don't meet the password security requirements starting from the date set for the `PASSWORD_CONTROL_EFFECTIVE_DATE` token. If a date is not set, the PCK disables, deletes or expires user accounts immediately. See [PASSWORD\\_CONTROL\\_EFFECTIVE\\_DATE](#) for more information.

- If the token `REQUIRE_RANDOM_ADMIN_PASSWORD` is already set to true, then set the token `ADMIN_EMAIL` with a valid email address.

```
ADMIN_EMAIL=root@{__APPLICATION_HOST__}
```

- If you have LDAP set up for external authentication, you must set the `"REQUIRE_USER_PASSWORD_CHANGE"` site options token to false.

### PostgreSQL Tokens and Settings

- Enable the `POSTGRES_INTERFACE` token in the `site-options.conf` file if the value of `"ifconfig -a"` is not `eth0/enp0*`.



- Make sure the PostgreSQL tokens in the `site-options.conf` file are set as recommended in the following topic: [What are the right PostgreSQL settings for my site?](#) on page 943
8. Download the corresponding version of Oracle client from <http://www.oracle.com/technetwork/database/features/instant-client/index-097480.html> and run the following command:

```
yum localinstall <path to oracle client rpm>
```

9. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



**Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

**Do this step by step on Oracle Database Server (server-03)**

10. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.

See [the Red Hat installation guide](#) for help.

11. Check your basic networking setup. See [Set up Networking](#) for more information.

12. Install Oracle 12c.



**Note:** Make sure your database uses UTF8 or AL32UTF8 encoding. This is needed to support users in Asian languages. See [this Oracle knowledge base article](#).

13. Log on to the Oracle Database Server as a system administrator with 'SYSDBA' privilege and run the following query.

```
alter system set parallel_threads_per_cpu=4;
```

**Do this step on TeamForge Application Server (server-01)**

14. Copy the Oracle Datamart setup script from `/opt/collabnet/teamforge/runtime/scripts/` to the `/tmp` directory of server-03.

```
scp /opt/collabnet/teamforge/runtime/scripts/datamart-oracle-setup.sh
<username>@<server-03>:/tmp
```

**Do this step on Oracle Database Server (server-03)**

15. Copy the Oracle Datamart setup script.

```
mkdir /u1
cp /tmp/datamart-oracle-setup.sh /u1
```

16. Log in as Oracle user and create the site database user and permissions.

See [Set up an Oracle database](#) on page 691 for help.

17. Create the reporting user and schema.



**Note:** Skip this step if you have already set up the datamart setup in the Oracle database. Your responses to the script's prompts must match the values of the equivalent variables in the `site-options.conf` file on `my.app.server`.

```
cd /u1
sh datamart-oracle-setup.sh
```

**Do this step by step on EventQ server (server-02)**

18. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.

See [the Red Hat installation guide](#) for help.

19. Check your basic networking setup. See [Set up Networking](#) for more information.

20. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge.

If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.

a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

b) If the output of the `getenforce` command is "permissive", continue with the next step. If not, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

21. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.  
b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

c) Refresh your repository cache.

```
yum clean all
```

#### TeamForge installation repository configuration for sites without internet access

a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD.

The DVD contains the necessary software and utilities required for installing TeamForge without internet access. In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/  
mkdir cdrom  
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]  
name=RHEL CDRom
```

```
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

22. Install the EventQ packages.

```
yum install teamforge-eventq
```

23. Copy the `site-options.conf` file from the TeamForge Application Server (server-01) to the EventQ Server's `/opt/collabnet/teamforge/etc/` directory.

24. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



**Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

**Do this step by step on TeamForge Application Server (server-01)**

25. Run the TeamForge post installation script. For more information, see [post-install.py](#).



**Important:** If the token `REQUIRE_USER_PASSWORD_CHANGE` is set to true, login to TeamForge user interface, change the admin password and then run the `post-install.py` script.

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

26. Restart TeamForge.

```
setenforce 1
/opt/collabnet/teamforge/bin/teamforge restart
```

27. Verify TeamForge installation.

- a) Reboot the server and make sure all services come up automatically at startup.
- b) Log on to the TeamForge web application using the default Admin credentials.
  - Username: "admin"
  - Password: "admin"
- c) Create a sample project. See [Create a TeamForge project](#)
- d) Write a welcome message to your site's users. See [Create a site-wide broadcast](#).

**Post Installation Tasks**

- [Install Review Board](#)
- [Run TeamForge in SELinux enabled mode](#)
- [Install Memcached](#) on page 694
- [Install EventQ Adapters](#)

**Install Code Search on a separate RHEL/CentOS 7.3 server**

Follow these instructions to install TeamForge Codesearch on a separate RHEL/CentOS 7.3 server.



**Attention:** Installing or upgrading TeamForge needs root privileges. You must log on as root or use a root shell to install or upgrade TeamForge.

1. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.

See [the Red Hat installation guide](#) for help.

2. Check your basic networking setup. See [Set up Networking](#) for more information.
3. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to /tmp.
- b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```


- c) Refresh your repository cache.

```
yum clean all
```

#### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit RPM package: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**


 **Note:** In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d)  **Note:** If the Red Hat Enterprise Linux/CentOS installation DVD is mounted already, skip the following instructions. If not, mount the DVD.

Mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

#### 4. Install Codesearch packages.

```
yum install teamforge-codesearch
```

#### 5. Copy the `site-options.conf` file from the TeamForge Application Server to the `/opt/collabnet/teamforge/etc/site-options.conf` directory of the Codesearch Server.

#### 6. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



#### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

- **SELinux:** [Run TeamForge in SELinux enabled mode](#).

### Verify your TeamForge 17.4 installation

Congratulations: you have just installed your TeamForge 17.4 site. Now you can apply some finishing touches and make sure everything is running smoothly.

#### 1. Log into your site as the administrator.

The value of the `PUBLIC_FQDN` variable in the `site-options.conf` file is the URL to login to TeamForge application.

#### 2. To install your license key.

See [Supply your TeamForge license key from Teamforge user interface](#) on page 493.

#### 3. To verify that the tracker initial load (triggered through `TrackerInitialJob`) completed successfully, run the following SQL from `[RUNTIME_DIR]/scripts/psql-reporting-wrapper` or by selecting the **Datamart** option on the **System Tools > Ad Hoc Database Query** page in the TeamForge web interface.

```
select * from etl_job where job_name ='tracker_initial_etl';
```

The status column should have a value of 1.

#### 4. To create a sample project.

See [Create a TeamForge project](#).

#### 5. To write a welcome message to your site's users,

see [Create a site-wide broadcast](#).

Now that you have successfully installed your TeamForge site in its basic configuration, you can use the instructions under [Maintain your TeamForge 17.4 site](#) on page 691 to help keep your site going.

### Supply your TeamForge license key from Teamforge user interface

Your license key enables you to use CollabNet TeamForge for the period of your contract.

Your license key will only work for the IP address of the machine that your CollabNet TeamForge is running on, as specified in your order form.



**Tip:** These steps are for installing your license key via the web interface. If you prefer, you can install it as a text file instead. See [Supply your TeamForge license key as a text file](#) on page 494.

#### 1. Locate the confirmation email you received from your CollabNet representative when you purchased your contract.

#### 2. Log into your site as the site administrator.



**Note:** The site administrator is different from the root user on the machine where the site is running.


#### 3. Click **Admin > License Key**.

If you have entered a license before, the IP address and current licensed number of users on your site are listed on the **License Key** page. Verify that the IP address is the same as the one you entered in your order form.

4. Click **Enter License Key**.
5. Copy your new license key from the confirmation email and paste it into the **Enter License Key** field.

A license key string looks like this:

```
ALM=100:SCM=80:DEVOPS=10:supervillaininc:144.16.116.25.:302D02150080D7853DB3E5C6F67EAB
```


 **Tip:** Save this license key in case you need to reinstall CollabNet TeamForge.

6. Click **Save**.
7. Verify that the new value for **Licensed Number of Users** matches the total number of licensed users in your contract.

### Supply your TeamForge license key as a text file

Your license key enables you to use CollabNet TeamForge for the period of your contract.


Your license key will only work for the IP address of the machine that your CollabNet TeamForge is running on.

 **Important:** If you are upgrading from a site with a limited number of users to an enterprise-scale site, you must install your license key before starting CollabNet TeamForge. Otherwise, your site could be rendered inoperable.


1. Locate the confirmation email you received from your CollabNet representative when you purchased your contract.
2. Create a text file and copy-paste your license key from the confirmation email into it.

For example, if your organization has 80 users who will use only the source code management features and 100 users who need the full range of application lifecycle management features, your license key string may look like this:

```
alm=100:SCM=80:supervillaininc:144.16.116.25.:302D02150080D7853DB3E5C6F67EABC65BD3AC17
```


 **Tip:** Save this license key in case you need to reinstall CollabNet TeamForge.

3. Save the text file as `/opt/collabnet/teamforge/etc/sflicense.txt`.

 **Tip:** Save your license key somewhere remote too, in case you need to reinstall CollabNet TeamForge and your `sflicense.txt` file is not accessible.

4. Make the license file usable by the application.

```
chmod 0664 /opt/collabnet/teamforge/etc/sflicense.txt
chown <APP_USER>:<APP_GROUP> /opt/collabnet/teamforge/etc/sflicense.txt
```

 **Note:** Change the values of `<APP_USER>` and `<APP_GROUP>` with the values of `APP_USER` and `APP_GROUP` tokens respectively from the `/opt/collabnet/teamforge/runtime/conf/runtime-options.conf` file.

## Upgrade TeamForge

You can upgrade to TeamForge 17.4 on RHEL/CentOS 6.8 or 7.3, either on the same hardware or on a new hardware. Choose the instructions that fit your case.

### Before you begin:

*[Plan your upgrade](#)*

### Remember:

- TeamForge 17.4 supports both RHEL/CentOS 6.8 and 7.3. See *[TeamForge installation requirements](#)* on page 441
- For the ETL service to run as expected in a distributed TeamForge installation, all servers must have the same time zone.

- While you can run both EventQ and TeamForge on the same server, CollabNet recommends such an approach only for testing purposes. It's always recommended to run EventQ on a separate server for optimal scalability. See [EventQ installation requirements](#).
- Installing or upgrading TeamForge needs root privileges. You must log on as root or use a root shell to install or upgrade TeamForge.
- In a distributed setup, stop TeamForge services on all the servers while upgrading to TeamForge 17.4.

### Upgrade to TeamForge 17.4 on new hardware: RHEL/CentOS 7.3

You can upgrade to TeamForge 17.4 on a new RHEL/CentOS 7.3 server with all services on one server.

#### Before you begin:

[Plan your upgrade](#)

#### Remember:

- TeamForge 17.4 supports both RHEL/CentOS 6.8 and 7.3. See [TeamForge installation requirements](#) on page 441
- For the ETL service to run as expected in a distributed TeamForge installation, all servers must have the same time zone.
- While you can run both EventQ and TeamForge on the same server, CollabNet recommends such an approach only for testing purposes. It's always recommended to run EventQ on a separate server for optimal scalability. See [EventQ installation requirements](#).
- Installing or upgrading TeamForge needs root privileges. You must log on as root or use a root shell to install or upgrade TeamForge.
- In a distributed setup, stop TeamForge services on all the servers while upgrading to TeamForge 17.4.

### Upgrade to TeamForge 17.4 on new hardware

To upgrade to TeamForge 17.4, set up a new hardware, then bring your old site's data and convert it.

#### Before you begin:

[Plan your upgrade](#)

#### Remember:

- TeamForge 17.4 supports both RHEL/CentOS 6.8 and 7.3. See [TeamForge installation requirements](#) on page 441
- For the ETL service to run as expected in a distributed TeamForge installation, all servers must have the same time zone.
- While you can run both EventQ and TeamForge on the same server, CollabNet recommends such an approach only for testing purposes. It's always recommended to run EventQ on a separate server for optimal scalability. See [EventQ installation requirements](#).
- Installing or upgrading TeamForge needs root privileges. You must log on as root or use a root shell to install or upgrade TeamForge.
- In a distributed setup, stop TeamForge services on all the servers while upgrading to TeamForge 17.4.

In this setup, the following services run on the TeamForge Application Server (we call this server-01).

- TeamForge Application Server (ctfcore)
- Database Server (ctfcore-database and ctfcore-datamart)
- Codesearch Server (codesearch)
- Mail Server (mail)
- ETL Server (etl)
- Git Integration Server (gerrit and gerrit-database)
- SCM Integration Server (subversion and cvs)
- Search Server (search).
- TeamForge EventQ Server (eventq, mongodb, redis and rabbitmq)
- TeamForge CLI Server (cliserver)

**Do this on the new TeamForge Application Server (ctfcore) - server-01**

1. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.


See [the Red Hat installation guide](#) for help.

2. Check your basic networking setup. See [Set up Networking](#) for more information.

**Do this on the existing TeamForge Application Server**

3. Back up all your custom event handlers and remove all the event handler JAR files before starting your TeamForge 17.4 upgrade process.

- a) Go to **My Workspace > Admin**.
- b) Click **System Tools** from the **Projects** menu.
- c) Click **Customizations**.
- d) Select the custom event handler and click **Delete**.

 **Important:** Post upgrade, you can add custom event handlers again from the backup while making sure that you don't have SOAP50 (deprecated) library used.

4. Stop TeamForge.

If you are upgrading from TeamForge 16.7 or earlier releases:

```
/etc/init.d/collabnet stop all
```

If you are upgrading from TeamForge 16.10 or later releases:

```
/opt/collabnet/teamforge/bin/teamforge stop
```

5. Stop EventQ.

```
/etc/init.d/eventq stop
```

```
/etc/init.d/collabnet-rabbitmq-server stop
```

```
/etc/init.d/collabnet-mongod stop
```

6. Backup your site data.


- a) If you are upgrading from TeamForge 16.7 or earlier to TeamForge 17.4:

**Back up your site database**

```
mkdir -p /tmp/backup_dir
```

```
cd /var/lib
```

```
tar -zcvf /tmp/backup_dir/pgsql.tgz postgres/9.3
```

 **Tip:** /tmp/backup\_dir is just an example. You can use any directory or partition you prefer to store your backup files.

**Back up the file system data**

Make an archive file with the following data directories:

Directory	Contents
/opt/collabnet/teamforge/var	User-created data, such as artifact attachments
/opt/collabnet/reviewboard	Review Board data
/svnroot	Subversion source code repositories
/sf-svnroot	Subversion repository for branding data
/cvsroot	CVS source code repositories (not present on all sites)
/gitroot	GIT source code repositories



```
cp -Rpfv /svnroot /sf-svnroot /cvsroot /opt/collabnet/teamforge/var /
opt/collabnet/reviewboard /tmp/backup_dir
```

#### If Git integration is enabled, do the following

```
mkdir /tmp/backup_dir/gerrit
cp -Rpfv /gitroot /tmp/backup_dir
cp -Rpfv /opt/collabnet/gerrit/ /tmp/backup_dir/gerrit
```

#### Back up EventQ


```
cp -Rpfv /opt/collabnet/eventq/config/*.yaml /opt/collabnet/mongodb /
opt/collabnet/rabbitmq /tmp/backup_dir
```

#### Compress your backup data

```
cd /tmp
tar czvf 167backup.tgz backup_dir
```

- b) If you are upgrading from TeamForge 16.10 or later to TeamForge 17.4:

```
mkdir -p /tmp/backup_dir
cp -Rpfv /opt/collabnet/teamforge/var /tmp/backup_dir
```

-  **Tip:** /tmp/backup\_dir is just an example. You can use any directory or partition you prefer to store your backup files.

#### If Git integration is enabled

```
mkdir /tmp/backup_dir/gerrit
cp -Rpfv /opt/collabnet/gerrit/ /tmp/backup_dir/gerrit
```

#### Back up EventQ

```
cp -Rpfv /opt/collabnet/eventq/config/*.yaml /opt/collabnet/mongodb /
opt/collabnet/rabbitmq /tmp/backup_dir
```

#### Compress your backup data


```
cd /tmp
tar czvf 171backup.tgz backup_dir
```

- c) Back up your SSH keys, if any.

- d) Back up your SSL certificates and keys, if any.

7. Copy the master configuration file from the this server to the /tmp directory on the new TeamForge server.

```
scp /opt/collabnet/teamforge/etc/site-options.conf username@newbox:/tmp
```

-  **Tip:** scp is just an example. You can choose any file transfer method you prefer.

8. Copy the backed up data to the new TeamForge Application Server.

```
scp /tmp/171backup.tgz username@newbox:/tmp or scp /tmp/167backup.tgz
username@newbox:/tmp
```

#### Do this on the new TeamForge Application Server (server-01)

9. Unpack the file system data.

```
cd /tmp
tar xzvf 171backup.tgz or tar xzvf 167backup.tgz
```

10. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge.

If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the /opt/collabnet/teamforge/bin/teamforge provision command. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is "permissive", continue with the next step. If not, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

11. Upgrade the operating system packages.

```
yum upgrade
```

12. Reboot the server.

```
reboot
```

13. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.  
b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```


- c) Refresh your repository cache.

```
yum clean all
```

#### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit RPM package: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**


 **Note:** In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -Uvh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d)  **Note:** If the Red Hat Enterprise Linux/CentOS installation DVD is mounted already, skip the following instructions. If not, mount the DVD.

Mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/  
mkdir cdrom  
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]  
name=RHEL CDRom
```

```
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

14. Install the TeamForge application packages.

```
yum install teamforge
```

15. Install PostgreSQL 9.3 packages.

```
yum install postgresql93-*
```

16. Reload site data.

- a) If you are upgrading from TeamForge 16.7 or earlier to TeamForge 17.4:

```
cd /opt/collabnet/teamforge/
mkdir var
cd var
tar zvxf /tmp/backup_dir/pgsql.tgz

cp -Rpfv /tmp/backup_dir/svnroot /svnroot
cp -Rpfv /tmp/backup_dir/cvsroot /cvsroot
cp -Rpfv /tmp/backup_dir/sf-svnroot /sf-svnroot
cp -Rpfv /tmp/backup_dir/var /opt/collabnet/teamforge/
```

#### If Git integration is enabled

```
cp -Rpfv /tmp/backup_dir/gitroot /gitroot
cp -Rpfv /tmp/backup_dir/gerrit/gerrit/etc /opt/collabnet/gerrit
cp -Rpf /tmp/backup_dir/gerrit/gerrit/.ssh /opt/collabnet/gerrit
cp -Rpf /tmp/backup_dir/gerrit/gerrit/bin /opt/collabnet/gerrit
cp -Rpf /tmp/backup_dir/gerrit/gerrit/index /opt/collabnet/gerrit
```

#### EventQ

```
yes | cp -Rf /tmp/backup_dir/*.yml /opt/collabnet/eventq/config/
yes | cp -Rf /tmp/backup_dir/mongodb /opt/collabnet
yes | cp -Rf /tmp/backup_dir/rabbitmq /opt/collabnet
```

- b) If you are upgrading from TeamForge 16.10 or later to TeamForge 17.4:

```
cp -Rpfv /tmp/backup_dir/var /opt/collabnet/teamforge/
```

#### If Git integration is enabled

```
cp -Rpfv /tmp/backup_dir/gerrit/gerrit/etc /opt/collabnet/gerrit
cp -Rpf /tmp/backup_dir/gerrit/gerrit/.ssh /opt/collabnet/gerrit
cp -Rpf /tmp/backup_dir/gerrit/gerrit/bin /opt/collabnet/gerrit
cp -Rpf /tmp/backup_dir/gerrit/gerrit/index /opt/collabnet/gerrit
```

#### EventQ

```
yes | cp -Rf /tmp/backup_dir*.yml/ /opt/collabnet/eventq/config/
yes | cp -Rf /tmp/backup_dir/mongodb /opt/collabnet
yes | cp -Rf /tmp/backup_dir/rabbitmq /opt/collabnet
```

17. Copy the site-options.conf file to the TeamForge installer directory.


```
cp /tmp/site-options.conf /opt/collabnet/teamforge/etc/
```

18. Set up your site's master configuration file.

- a) Set up your site options. See [Site options change log](#) on page 730 for a list of site option changes. While upgrading to a latest TeamForge release, make sure that obsolete site option tokens, if any, are removed from the `site-options.conf` file of the TeamForge version you are upgrading to.


```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

- b) Configure the services and domain name tokens.

 **Note:** `cliserver` and `reviewboard-adapter` are the new services added in TeamForge 17.4.


```
server-01:SERVICES = ctfcore ctfcore-database ctfcore-datamart mail etl
search codesearch
subversion cvs eventq redis mongodb rabbitmq cliserver gerrit gerrit-
database binary binary-database
reviewboard reviewboard-database reviewboard-adapter
```

```
server-01:PUBLIC_FQDN = my.app.domain.com
```

 **Note:** You may remove the identifiers of components you do not want. For example, remove `binary` and `binary-database` if you are not planning to install binary repository managers such as Nexus.

- c) Set the `MONGODB_APP_DATABASE_NAME` token with EventQ's database name in the `site-options.conf` file.

```
MONGODB_APP_DATABASE_NAME=orchestrate
```


- d)  **Attention:** *SSL* is enabled by default and a self-signed certificate is auto-generated. Use the following tokens to adjust this behavior. To generate the SSL certificates, see [Generate SSL certificates](#) on page 718.

Have the custom SSL certificate and private key for custom SSL certificate in place and provide their absolute paths in these tokens. `SSL_CHAIN_FILE` (intermediate certificate) is optional.


```
SSL_CERT_FILE=
```

```
SSL_KEY_FILE=
```


```
SSL_CHAIN_FILE=
```

 **Important:** All SSL certificates including self-signed certificates are added automatically.

- e) If the token `REQUIRE_PASSWORD_SECURITY` is enabled, then set a value for the token, `PASSWORD_CONTROL_EFFECTIVE_DATE`.

 **Caution:** The Password Control Kit (PCK) disables, deletes or expires user accounts that don't meet the password security requirements starting from the date set for the `PASSWORD_CONTROL_EFFECTIVE_DATE` token. If a date is not set, the PCK disables, deletes or expires user accounts immediately. See [PASSWORD\\_CONTROL\\_EFFECTIVE\\_DATE](#) for more information.

- f) TeamForge 7.1 and later support automatic password creation. See [AUTO\\_DATA](#) on page 736 for more information.
- g) If you have LDAP set up for external authentication, you must set the "`REQUIRE_USER_PASSWORD_CHANGE`" site options token to false.
- h) Make sure the PostgreSQL tokens in the `site-options.conf` file are set as recommended in the following topic: [What are the right PostgreSQL settings for my site?](#) on page 943
- i) Configure the `JBOSS_JAVA_OPTS` `site-options.conf` token. See [JBOSS\\_JAVA\\_OPTS](#) on page 747.

 **Important:** All JVM parameters but `-Xms1024m` and `-Xmx2048m` have been hard-coded in the TeamForge core application. You cannot manually configure any of the following default JVM parameters in the `site-options.conf` file.

- `-XX:+UseParallelGC`
- `-XX:MaxMetaspaceSize=512m`
- `-XX:ReservedCodeCacheSize=128M`
- `-server`
- `-XX:+HeapDumpOnOutOfMemoryError`
- `-XX:HeapDumpPath=/tmp -verbose:gc`
- `-XX:+PrintCodeCache`
- `-Djsse.enableSNIExtension=false`
- `-Dsun.rmi.dgc.client.gcInterval=600000`
- `-Dsun.rmi.dgc.server.gcInterval=600000`
- `-Djava.security.egd=file:/dev/urandom`
- `-Djava.awt.headless=true`.

When you change the default value of a JVM parameter such as `"-XX:HeapDumpPath"`, the JBoss runtime parameters include both the user defined and default values for the JVM parameter. However, JBoss runs with the default value and ignores any user defined value.

j) Save the `site-options.conf` file.

19. As you are upgrading on new hardware, contact [CollabNet Support](#), generate the license key for the new server (IP address) and use it to replace `/opt/collabnet/teamforge/var/etc/sflicense.txt`

20. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).


TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

21. Run the `/var/lib/pgsql/analyze_new_cluster.sh` script.

```
su - postgres -c "/var/lib/pgsql/analyze_new_cluster.sh"
```

22. Update the file permissions on your site's data.

```
/opt/collabnet/teamforge/runtime/scripts/fix_data_permissions.sh
```

 **Note:** This process can take a long time on sites with a lot of data.

23. If you have CVS integrations, synchronize permissions post upgrade. See, [Synchronize TeamForge source control integrations](#) on page 695.

24. Run the TeamForge post installation script. For more information, see [post-install.py](#).

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

25. Verify TeamForge installation.

a) Log on to the TeamForge web application using the default Admin credentials.

- Username: "admin"
- Password: "admin"

b) If your site has custom branding, verify that your branding changes still work as intended.

See [Customize TeamForge](#) .

c) Let your site's users know they've been upgraded.

See [Create a Side-wide Broadcast](#).

26. Remove the backup files, if any, after the TeamForge site is up and running as expected. Remove the repository and the file system backup from the `/tmp/backup_dir` directory.

- [Upgrade Review Board](#)
- [Run TeamForge in SELinux enabled mode](#)
- [Add EventQ to existing projects](#)
- [Run ETL initial load jobs](#)

### Upgrade to TeamForge 17.4 on same hardware: RHEL/CentOS 6.8/7.3

The following instructions are valid for upgrading to TeamForge 17.4 on the same hardware on both RHEL/CentOS 6.8/7.3 platforms. Specific steps, if applicable only for a particular RHEL/CentOS platform, are called out explicitly.

#### Before you begin:

[Plan your upgrade](#)

#### Remember:

- TeamForge 17.4 supports both RHEL/CentOS 6.8 and 7.3. See [TeamForge installation requirements](#) on page 441
- For the ETL service to run as expected in a distributed TeamForge installation, all servers must have the same time zone.
- While you can run both EventQ and TeamForge on the same server, CollabNet recommends such an approach only for testing purposes. It's always recommended to run EventQ on a separate server for optimal scalability. See [EventQ installation requirements](#).
- Installing or upgrading TeamForge needs root privileges. You must log on as root or use a root shell to install or upgrade TeamForge.
- In a distributed setup, stop TeamForge services on all the servers while upgrading to TeamForge 17.4.

### Upgrade to TeamForge 17.4 - All services on the same RHEL/CentOS 6.8 server

Single server setup. In this procedure, we'll assume that you are upgrading on the same server where your existing TeamForge 17.1 site is running. The following instructions are valid for both RHEL/CentOS 6.8/7.3 platforms. Specific steps, if applicable only for a particular RHEL/CentOS platform, are called out explicitly.

#### Before you begin:

[Plan your upgrade](#)

#### Remember:

- TeamForge 17.4 supports both RHEL/CentOS 6.8 and 7.3. See [TeamForge installation requirements](#) on page 441
- For the ETL service to run as expected in a distributed TeamForge installation, all servers must have the same time zone.
- While you can run both EventQ and TeamForge on the same server, CollabNet recommends such an approach only for testing purposes. It's always recommended to run EventQ on a separate server for optimal scalability. See [EventQ installation requirements](#).
- Installing or upgrading TeamForge needs root privileges. You must log on as root or use a root shell to install or upgrade TeamForge.
- In a distributed setup, stop TeamForge services on all the servers while upgrading to TeamForge 17.4.

In this setup, all the following services run on a single RHEL/CentOS 6.8/7.3 server (we call this server-01).


- TeamForge Application Server (ctfcore)
- Codesearch Server (codesearch)
- Mail Server (mail)
- Database Server (ctfcore-database and ctfcore-datamart)
- ETL Server (etl)
- Git Integration Server (gerrit and gerrit-database)

- SCM Integration Server (subversion and cvs)
- Search Server (search)
- TeamForge EventQ Server (eventq, mongodb, redis and rabbitmq)

### Do the following on the TeamForge Application Server (server-01)

1. Back up all your custom event handlers and remove all the event handler JAR files before starting your TeamForge 17.4 upgrade process.

- a) Go to **My Workspace > Admin**.
- b) Click **System Tools** from the **Projects** menu.
- c) Click **Customizations**.
- d) Select the custom event handler and click **Delete**.

 **Important:** Post upgrade, you can add custom event handlers again from the backup while making sure that you don't have SOAP50 (deprecated) library used.

2. Uninstall hotfixes and add-ons, if any, installed on your site.
3. If you have Review Board installed, uninstall it.

```
cd /opt/collabnet/RBInstaller-17.1.5
python ./install.py -u
```

4. Stop TeamForge.

If you are upgrading from TeamForge 16.7 or earlier releases:

```
/etc/init.d/collabnet stop all
```

If you are upgrading from TeamForge 16.10 or later releases:

```
/opt/collabnet/teamforge/bin/teamforge stop
```

5. Stop EventQ.

```
/etc/init.d/eventq stop
/etc/init.d/collabnet-rabbitmq-server stop
/etc/init.d/collabnet-mongod stop
```

6. No backup is required for same hardware upgrades. However, you can create a backup as a precaution. See [Backup and restore TeamForge and EventQ](#) on page 566 to learn more about backing up TeamForge and EventQ database and file system.

7. Upgrade the operating system packages.

```
yum upgrade
```

8. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge.

If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is "permissive", continue with the next step. If not, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

9. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.
- b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```


- c) Refresh your repository cache.

```
yum clean all
```

### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.

- Red Hat Enterprise Linux/CentOS 6.8 64 bit RPM package: **CTF-Disconnected-media-17.4.824-110.rhel6.x86\_64.rpm**
- Red Hat Enterprise Linux/CentOS 7.3 64 bit RPM package: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**


 **Note:** In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -Uvh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d)  **Note:** If the Red Hat Enterprise Linux/CentOS installation DVD is mounted already, skip the following instructions. If not, mount the DVD.

Mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

## 10. Install the following application packages.

- a) TeamForge: To install the TeamForge application packages run the following command:

```
yum install teamforge CN-eventq
```



**Attention:** TeamForge installer has been optimized quite a bit. It's likely that you might come across a lot of warning messages while upgrading from TeamForge 8.2 (or earlier) to TeamForge 17.4 on the



same hardware (when you run the `yum install teamforge` command). You can safely ignore such warning messages and proceed with the upgrade.

Run the following command to install the Binary application packages.


```
yum install cn-binary
```

## 11. Set up your site's master configuration file.

- a) Set up your site options. See [Site options change log](#) on page 730 for a list of site option changes. While upgrading to a latest TeamForge release, make sure that obsolete site option tokens, if any, are removed from the `site-options.conf` file of the TeamForge version you are upgrading to.


```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

- b) TeamForge 16.7 and later versions support automatic `JAVA_HOME` configuration. Remove the `JAVA_HOME` token, if present, from the `site-options.conf` file.
- c) Configure the services and domain name tokens.

 **Note:** `cliserver` and `reviewboard-adapter` are the new services added in TeamForge 17.4.


```
server-01:SERVICES=ctfcore ctfcore-database ctfcore-datamart mail etl
search codesearch
subversion cvs eventq redis mongodb rabbitmq cliserver gerrit gerrit-
database binary binary-database
reviewboard reviewboard-database reviewboard-adapter
```

```
server-01:PUBLIC_FQDN=my.app.domain.com
```

 **Note:** You may remove the identifiers of components you do not want. For example, remove `binary` and `binary-database` if you are not planning to install binary repository managers such as Nexus.

- d) Set the `MONGODB_APP_DATABASE_NAME` token with EventQ's database name in the `site-options.conf` file.

```
MONGODB_APP_DATABASE_NAME=orchestrate
```


- e)  **Attention:** `SSL` is enabled by default and a self-signed certificate is auto-generated. Use the following tokens to adjust this behavior. To generate the SSL certificates, see [Generate SSL certificates](#) on page 718.

Have the custom SSL certificate and private key for custom SSL certificate in place and provide their absolute paths in these tokens. `SSL_CHAIN_FILE` (intermediate certificate) is optional.


```
SSL_CERT_FILE=
```

```
SSL_KEY_FILE=
```

```
SSL_CHAIN_FILE=
```

 **Important:** All SSL certificates including self-signed certificates are added automatically.

- f) If the token `REQUIRE_PASSWORD_SECURITY` is enabled, then set a value for the token, `PASSWORD_CONTROL_EFFECTIVE_DATE`.

 **Caution:** The Password Control Kit (PCK) disables, deletes or expires user accounts that don't meet the password security requirements starting from the date set for the `PASSWORD_CONTROL_EFFECTIVE_DATE` token. If a date is not set, the PCK disables, deletes or expires user accounts immediately. See [PASSWORD\\_CONTROL\\_EFFECTIVE\\_DATE](#) for more information.

- g) TeamForge 7.1 and later support automatic password creation. See [AUTO\\_DATA](#) on page 736 for more information.
- h) If you have LDAP set up for external authentication, you must set the “[REQUIRE\\_USER\\_PASSWORD\\_CHANGE](#)” site options token to false.
- i) Make sure the PostgreSQL tokens in the `site-options.conf` file are set as recommended in the following topic: [What are the right PostgreSQL settings for my site?](#) on page 943
- j) Configure the `JBOSS_JAVA_OPTS` `site-options.conf` token. See [JBOSS\\_JAVA\\_OPTS](#) on page 747.



**Important:** All JVM parameters but `-Xms1024m` and `-Xmx2048m` have been hard-coded in the TeamForge core application. You cannot manually configure any of the following default JVM parameters in the `site-options.conf` file.

- `-XX:+UseParallelGC`
- `-XX:MaxMetaspaceSize=512m`
- `-XX:ReservedCodeCacheSize=128M`
- `-server`
- `-XX:+HeapDumpOnOutOfMemoryError`
- `-XX:HeapDumpPath=/tmp -verbose:gc`
- `-XX:+PrintCodeCache`
- `-Djsse.enableSNIExtension=false`
- `-Dsun.rmi.dgc.client.gcInterval=600000`
- `-Dsun.rmi.dgc.server.gcInterval=600000`
- `-Djava.security.egd=file:/dev/urandom`
- `-Djava.awt.headless=true`

When you change the default value of a JVM parameter such as “`-XX:HeapDumpPath`”, the JBoss runtime parameters include both the user defined and default values for the JVM parameter. However, JBoss runs with the default value and ignores any user defined value.

- k) Save the `site-options.conf` file.

## 12. Deploy services.

- a) Move all backup files and folders including the error folder to `/tmp`.



**Note:** This step is required to make sure that there are no backup files and folders present in `/opt/collabnet/teamforge/var/james/var/mail/` to avoid any delay during provisioning.

```
cd /opt/collabnet/teamforge/var/james/var/mail/
mv * /tmp
```

- b) Run the TeamForge provision command.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



### Note:

The “`provision`” command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter “Y” or “N” to proceed. For more information, see [TeamForge script](#) on page 785.

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge “`provision`” command fails otherwise.

- 13. Run the `/var/lib/pgsql/analyze_new_cluster.sh` script.

```
su - postgres -c "/var/lib/pgsql/analyze_new_cluster.sh"
```

- 14. If you have CVS integrations, synchronize permissions post upgrade. See, [Synchronize TeamForge source control integrations](#) on page 695.

- 15. Run the TeamForge post installation script. For more information, see [post-install.py](#).

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

**16. Verify TeamForge installation.**

- a) Log on to the TeamForge web application using the default Admin credentials.
  - Username: "admin"
  - Password: "admin"
- b) If your site has custom branding, verify that your branding changes still work as intended.  
See [Customize TeamForge](#) .
- c) Let your site's users know they've been upgraded.  
See [Create a Side-wide Broadcast](#).

**17. Remove the backup files, if any, after the TeamForge site is up and running as expected. Remove the repository and the file system backup from the /tmp/backup\_dir directory.**

**TeamForge Avatar display issue on RHEL/CentOS 6.8**

TeamForge Avatar image is not displayed properly post Review Board installation on RHEL/CentOS 6.8. Run the following commands to work around this issue:

```
yum erase python-imaging
yum install teamforge
service httpd restart
```

- [Upgrade Review Board](#)
- [Run TeamForge in SELinux enabled mode](#)
- [Add EventQ to existing projects](#)
- [Run ETL initial load jobs](#)

**Upgrade TeamForge, EventQ, Database and Datamart on Three Separate Servers**

Three-server setup, same hardware upgrade, with TeamForge, EventQ, Database and Datamart on three separate servers.

**Dos and Don'ts**

Check [this](#) list while installing or upgrading TeamForge.

**Three Server Setup**

Here's how the services are distributed across three separate RHEL/CentOS 6.8/7.3 servers.


Services	TeamForge Application Server (server-01)	EventQ Server (server-02)	Database Server (server-03)
ctfcore	✓		
mail	✓		
search	✓		
codesearch	✓		
etl	✓		
gerrit	✓		
gerrit-database			✓
reviewboard	✓		
reviewboard-database	✓		

Services	TeamForge Application Server (server-01)	EventQ Server (server-02)	Database Server (server-03)
<b>reviewboard-adapter</b>	✓		
subversion	✓		
cvs	✓		
binary	✓		
binary-database			✓
<b>cliserver</b>	✓		
eventq		✓	
mongodb		✓	
redis		✓	
rabbitmq		✓	
ctfcore-database			✓
ctfcore-datamart			✓

No backup is required for some hardware upgrades. However, you can create a backup as a precaution. See [Back up and restore TeamForge and EventQ](#) on page 566 to learn more about backing up TeamForge and EventQ database and file system.

#### Do this step by step on TeamForge Application Server (server-01)

1. Back up all your custom event handlers and remove all the event handler JAR files before starting your TeamForge 17.4 upgrade process.
  - a) Go to **My Workspace > Admin**.
  - b) Click **System Tools** from the **Projects** menu.
  - c) Click **Customizations**.
  - d) Select the custom event handler and click **Delete**.

 **Important:** Post upgrade, you can add custom event handlers again from the backup while making sure that you don't have SOAP50 (deprecated) library used.

2. Uninstall hotfixes and add-ons, if any, installed on your site.
3. If you have Review Board installed, uninstall it.

```
cd /opt/collabnet/RBInstaller-17.1.5
python ./install.py -u
```

4. Stop TeamForge.

If you are upgrading from TeamForge 16.7 or earlier releases:

```
/etc/init.d/collabnet stop all
```

If you are upgrading from TeamForge 16.10 or later releases:

```
/opt/collabnet/teamforge/bin/teamforge stop
```

5. Upgrade the operating system packages.

```
yum upgrade
```

6. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge.

If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is "permissive", continue with the next step. If not, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

7. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.
- b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

#### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 6.8 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel6.x86\_64.rpm**
- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -Uvh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
```

```
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

8. Install the TeamForge application packages.

```
yum install teamforge
```



**Attention:** TeamForge installer has been optimized quite a bit. It's likely that you might come across a lot of warning messages while upgrading from TeamForge 8.2 (or earlier) to TeamForge 17.4 on the same hardware (when you run the `yum install teamforge` command). You can safely ignore such warning messages and proceed with the upgrade.

9. Install the Binary application packages.

```
yum install cn-binary
```

**Do this step by step on the Database Server (server-03)**

10. Upgrade the operating system packages.

```
yum upgrade
```

11. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is either "Disabled" or "Permissive", SELinux is already disabled.

- c) If not disabled, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

12. Configure your TeamForge installation repository.

#### **TeamForge installation repository configuration for sites with internet access**

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.

- b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

#### **TeamForge installation repository configuration for sites without internet access**

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 6.8 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel6.x86\_64.rpm**
- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -Uvh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDROM
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

### 13. Stop TeamForge.

If you are upgrading from TeamForge 16.7 or earlier releases:

```
/etc/init.d/collabnet stop all
```

If you are upgrading from TeamForge 16.10 or later releases:

```
/opt/collabnet/teamforge/bin/teamforge stop
```


### 14. Install the TeamForge database packages.

```
yum install teamforge-database
```


### 15. Set up your site's master configuration file.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

#### a) **host: SERVICES Token**

 **Note:** cliserver and reviewboard-adapter are the new services added in TeamForge 17.4.

- ```
server-01:SERVICES=ctfcore mail etl search subversion cvs codesearch
cliserver gerrit
binary reviewboard reviewboard-database reviewboard-adapter
```
- ```
server-02:SERVICES=eventq mongodb redis rabbitmq
```
- ```
server-03:SERVICES=ctfcore-database ctfcore-datamart gerrit-database
binary-database
```

 **Note:** You may remove the identifiers of components you do not want. For example, remove binary and binary-database if you are not planning to install binary repository managers such as Nexus.

#### b) **host: PUBLIC\_FQDN Token**

```
server-01:PUBLIC_FQDN=my.app.domain.com
```

- c) Set the `MONGODB_APP_DATABASE_NAME` token with EventQ's database name in the `site-options.conf` file.

```
MONGODB_APP_DATABASE_NAME=orchestrate
```

#### d) SSL Tokens

SSL is enabled by default and a self-signed certificate is auto-generated. Use the following tokens to adjust this behavior.

```
SSL_CERT_FILE=
SSL_KEY_FILE=
SSL_CHAIN_FILE=
```

- To generate the SSL certificates, see [Generate SSL certificates](#) on page 718.
- Have the custom SSL certificate and private key for custom SSL certificate in place and provide their absolute paths in these tokens. `SSL_CHAIN_FILE` (intermediate certificate) is optional.
- All SSL certificates including self-signed certificates are added automatically.

#### e) Password Tokens

- TeamForge 7.1 and later support automatic password creation. See [AUTO\\_DATA](#) on page 736 for more information.
- If the token `REQUIRE_PASSWORD_SECURITY` is enabled, then set a value for the token, `PASSWORD_CONTROL_EFFECTIVE_DATE`.



**Warning:** The Password Control Kit (PCK) disables, deletes or expires user accounts that don't meet the password security requirements starting from the date set for the `PASSWORD_CONTROL_EFFECTIVE_DATE` token. If a date is not set, the PCK disables, deletes or expires user accounts immediately. See [PASSWORD\\_CONTROL\\_EFFECTIVE\\_DATE](#) for more information.

- If the token `REQUIRE_RANDOM_ADMIN_PASSWORD` is already set to true, then set the token `ADMIN_EMAIL` with a valid email address.

```
ADMIN_EMAIL=root@{__APPLICATION_HOST__}
```

- If you have LDAP set up for external authentication, you must set the `"REQUIRE_USER_PASSWORD_CHANGE"` `site-options.conf` token to false.

#### f) PostgreSQL Tokens and Settings

- Enable the `POSTGRES_INTERFACE` token in the `site-options.conf` file if the value of "ifconfig - a" is not eth0/enp0\*.
- Make sure the PostgreSQL tokens in the `site-options.conf` file are set as recommended in the following topic: [What are the right PostgreSQL settings for my site?](#) on page 943

#### g) JAVA\_OPTS

Configure the `JBOSS_JAVA_OPTS` `site-options.conf` token. See [JBOSS\\_JAVA\\_OPTS](#) on page 747.



**Note:** All JVM parameters but `-Xms1024m` and `-Xmx2048m` have been hard-coded in the TeamForge core application. You cannot manually configure any of the following default JVM parameters in the `site-options.conf` file.

- `-XX:+UseParallelGC`
- `-XX:MaxMetaspaceSize=512m`
- `-XX:ReservedCodeCacheSize=128M`
- `-server`
- `-XX:+HeapDumpOnOutOfMemoryError`
- `-XX:HeapDumpPath=/tmp -verbose:gc`
- `-XX:+PrintCodeCache`



- -Djsse.enableSNIExtension=false
- -Dsun.rmi.dgc.client.gcInterval=600000
- -Dsun.rmi.dgc.server.gcInterval=600000
- -Djava.security.egd=file:/dev/urandom
- -Djava.awt.headless=true.

When you change the default value of a JVM parameter such as "-XX:HeapDumpPath", the JBoss runtime parameters include both the user defined and default values for the JVM parameter. However, JBoss runs with the default value and ignores any user defined value.

h) Save the `site-options.conf` file.

## 16. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

## 17. Run the `/var/lib/pgsql/analyze_new_cluster.sh` script.

```
su - postgres -c "/var/lib/pgsql/analyze_new_cluster.sh"
```

### Do this step by step on TeamForge Application Server (server-01)

## 18. Copy the `site-options.conf` file from the Database Server to the TeamForge Application Server's `/opt/collabnet/teamforge/etc/` directory.

## 19. Deploy services.

a) Move all backup files and folders including the error folder to `/tmp`.



**Note:** This step is required to make sure that there are no backup files and folders present in `/opt/collabnet/teamforge/var/james/var/mail/` to avoid any delay during provisioning.

```
cd /opt/collabnet/teamforge/var/james/var/mail/
mv * /tmp
```

b) Run the TeamForge provision command.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "Y" or "N" to proceed. For more information, see [TeamForge script](#) on page 785.

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

## 20. Run the `/var/lib/pgsql/analyze_new_cluster.sh` script.

```
su - postgres -c "/var/lib/pgsql/analyze_new_cluster.sh"
```

## 21. Update the file permissions on your site's data.

```
/opt/collabnet/teamforge/runtime/scripts/fix_data_permissions.sh
```



**Note:** This process can take a long time on sites with a lot of data.

## 22. If you have CVS integrations, synchronize permissions post upgrade. See, [Synchronize TeamForge source control integrations](#) on page 695.

## 23. Run the TeamForge post installation script. For more information, see [post-install.py](#).

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

### Do this step by step on EventQ Server (server-02)

#### 24. Stop EventQ.

```
/etc/init.d/eventq stop
/etc/init.d/collabnet-rabbitmq-server stop
/etc/init.d/collabnet-mongod stop
```

#### 25. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is either "Disabled" or "Permissive", SELinux is already disabled.  
c) If not disabled, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

#### 26. Upgrade the operating system packages.

```
yum upgrade
```

#### 27. Configure your TeamForge installation repository.

##### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.  
b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

##### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD.

The DVD contains the necessary software and utilities required for installing TeamForge without internet access. In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

28. Install the EventQ packages.

```
yum install teamforge-eventq CN-eventq
```

29. Copy the `site-options.conf` file from the Database Server (server-03) to the EventQ Server's `/opt/collabnet/teamforge/etc/` directory.

30. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

**Do this step by step on TeamForge Application Server (server-01)**

31. Verify TeamForge installation.

a) Log on to the TeamForge web application using the default Admin credentials.

- Username: "admin"
- Password: "admin"

b) If your site has custom branding, verify that your branding changes still work as intended.

See [Customize TeamForge](#).

c) Let your site's users know they've been upgraded.

See [Create a Side-wide Broadcast](#).

32. Remove the backup files, if any, after the TeamForge site is up and running as expected. Remove the repository and the file system backup from the `/tmp/backup_dir` directory.

33. Fix the TeamForge Avatar display issue on RHEL/CentOS 6.8.

TeamForge Avatar image is not displayed properly post Review Board installation on RHEL/CentOS 6.8. Run the following commands to work around this issue:

```
yum erase python-imaging
yum install teamforge
service httpd restart
```

### Post Upgrade Tasks

- [Upgrade Review Board](#)
- [Run TeamForge in SELinux enabled mode](#)
- [Add EventQ to existing projects](#)
- [Run ETL initial load jobs](#)

## Upgrade TeamForge, EventQ, and Codesearch on Three Separate Servers

Three-server setup, same hardware upgrade, with TeamForge, EventQ, and Codesearch on three separate servers.

### Dos and Don'ts

Check [this](#) list while installing or upgrading TeamForge.

### Three Server Setup

Here's how the services are distributed across three separate RHEL/CentOS 6.8/7.3 servers.


| Services                   | TeamForge Application Server (server-01) | EventQ Server (server-02) | Codesearch Server (server-03) |
|----------------------------|------------------------------------------|---------------------------|-------------------------------|
| ctfcore                    | ✓                                        |                           |                               |
| mail                       | ✓                                        |                           |                               |
| search                     | ✓                                        |                           |                               |
| codesearch                 |                                          |                           | ✓                             |
| etl                        | ✓                                        |                           |                               |
| gerrit                     | ✓                                        |                           |                               |
| gerrit-database            | ✓                                        |                           |                               |
| reviewboard                | ✓                                        |                           |                               |
| reviewboard-database       | ✓                                        |                           |                               |
| <b>reviewboard-adapter</b> | ✓                                        |                           |                               |
| subversion                 | ✓                                        |                           |                               |
| cvs                        | ✓                                        |                           |                               |
| binary                     | ✓                                        |                           |                               |
| binary-database            | ✓                                        |                           |                               |
| <b>cliserver</b>           | ✓                                        |                           |                               |
| eventq                     |                                          | ✓                         |                               |
| mongodb                    |                                          | ✓                         |                               |
| redis                      |                                          | ✓                         |                               |
| rabbitmq                   |                                          | ✓                         |                               |
| ctfcore-database           | ✓                                        |                           |                               |
| ctfcore-datamart           | ✓                                        |                           |                               |

No backup is required for same hardware upgrades. However, you can create a backup as a precaution. See [Back up and restore TeamForge and EventQ](#) on page 566 to learn more about backing up TeamForge and EventQ database and file system.

### Do this step by step on TeamForge Application Server (server-01)

1. Back up all your custom event handlers and remove all the event handler JAR files before starting your TeamForge 17.4 upgrade process.

- a) Go to **My Workspace > Admin**.
- b) Click **System Tools** from the **Projects** menu.
- c) Click **Customizations**.
- d) Select the custom event handler and click **Delete**.

 **Important:** Post upgrade, you can add custom event handlers again from the backup while making sure that you don't have SOAP50 (deprecated) library used.

2. Uninstall hotfixes and add-ons, if any, installed on your site.
3. If you have Review Board installed, uninstall it.

```
cd /opt/collabnet/RBInstaller-17.1.5
python ./install.py -u
```

4. Stop TeamForge.

If you are upgrading from TeamForge 16.7 or earlier releases:

```
/etc/init.d/collabnet stop all
```

If you are upgrading from TeamForge 16.10 or later releases:

```
/opt/collabnet/teamforge/bin/teamforge stop
```

5. Upgrade the operating system packages.

```
yum upgrade
```

6. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge.

If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is "permissive", continue with the next step. If not, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

7. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.
- b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

#### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 6.8 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel6.x86\_64.rpm**

- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
  - In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.
- b) Unpack the disconnected installation package.
- ```
rpm -Uvh <package-name>
```
- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

## 8. Install the TeamForge application packages.

```
yum install teamforge
```



**Attention:** TeamForge installer has been optimized quite a bit. It's likely that you might come across a lot of warning messages while upgrading from TeamForge 8.2 (or earlier) to TeamForge 17.4 on the same hardware (when you run the `yum install teamforge` command). You can safely ignore such warning messages and proceed with the upgrade.

## 9. Install the Binary application packages.

```
yum install cn-binary
```

## 10. Set up your site's master configuration file.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

- a) **host: SERVICES Token**




**Note:** `cliserver` and `reviewboard-adapter` are the new services added in TeamForge 17.4.

- ```
server-01:SERVICES=ctfcore ctfcore-database ctfcore-datamart mail etl
search
subversion cvs cliserver gerrit gerrit-database binary binary-
database reviewboard reviewboard-database
```

```
reviewboard-adapter
```

- `server-02:SERVICES=eventq mongodb redis rabbitmq`
- `server-03:SERVICES=codesearch`

 **Note:** You may remove the identifiers of components you do not want. For example, remove binary and binary-database if you are not planning to install binary repository managers such as Nexus.

#### b) host: PUBLIC\_FQDN Token

```
server-01:PUBLIC_FQDN=my.app.domain.com
```

- c) Set the `MONGODB_APP_DATABASE_NAME` token with EventQ's database name in the `site-options.conf` file.

```
MONGODB_APP_DATABASE_NAME=orchestrate
```

#### d) SSL Tokens

SSL is enabled by default and a self-signed certificate is auto-generated. Use the following tokens to adjust this behavior.

```
SSL_CERT_FILE=
SSL_KEY_FILE=
SSL_CHAIN_FILE=
```

- To generate the SSL certificates, see [Generate SSL certificates](#) on page 718.
- Have the custom SSL certificate and private key for custom SSL certificate in place and provide their absolute paths in these tokens. `SSL_CHAIN_FILE` (intermediate certificate) is optional.
- All SSL certificates including self-signed certificates are added automatically..

#### e) Password Tokens

- TeamForge 7.1 and later support automatic password creation. See [AUTO\\_DATA](#) on page 736 for more information.
- If the token `REQUIRE_PASSWORD_SECURITY` is enabled, then set a value for the token, `PASSWORD_CONTROL_EFFECTIVE_DATE`.



**Warning:** The Password Control Kit (PCK) disables, deletes or expires user accounts that don't meet the password security requirements starting from the date set for the `PASSWORD_CONTROL_EFFECTIVE_DATE` token. If a date is not set, the PCK disables, deletes or expires user accounts immediately. See [PASSWORD\\_CONTROL\\_EFFECTIVE\\_DATE](#) for more information.

- If the token `REQUIRE_RANDOM_ADMIN_PASSWORD` is already set to true, then set the token `ADMIN_EMAIL` with a valid email address.

```
ADMIN_EMAIL=root@{__APPLICATION_HOST__}
```


- If you have LDAP set up for external authentication, you must set the `"REQUIRE_USER_PASSWORD_CHANGE"` `site-options.conf` token to false.

#### f) PostgreSQL Tokens and Settings

- Enable the `POSTGRES_INTERFACE` token in the `site-options.conf` file if the value of "ifconfig - a" is not `eth0/enp0*`.
- Make sure the PostgreSQL tokens in the `site-options.conf` file are set as recommended in the following topic: [What are the right PostgreSQL settings for my site?](#) on page 943

#### g) JAVA\_OPTS

Configure the `JBOSS_JAVA_OPTS` `site-options.conf` token. See [JBOSS\\_JAVA\\_OPTS](#) on page 747.

 **Note:** All JVM parameters but `-Xms1024m` and `-Xmx2048m` have been hard-coded in the TeamForge core application. You cannot manually configure any of the following default JVM parameters in the `site-options.conf` file.


- `-XX:+UseParallelGC`
- `-XX:MaxMetaspaceSize=512m`
- `-XX:ReservedCodeCacheSize=128M`
- `-server`
- `-XX:+HeapDumpOnOutOfMemoryError`
- `-XX:HeapDumpPath=/tmp -verbose:gc`
- `-XX:+PrintCodeCache`
- `-Djsse.enableSNIExtension=false`
- `-Dsun.rmi.dgc.client.gcInterval=600000`
- `-Dsun.rmi.dgc.server.gcInterval=600000`
- `-Djava.security.egd=file:/dev/urandom`
- `-Djava.awt.headless=true`.

When you change the default value of a JVM parameter such as `"-XX:HeapDumpPath"`, the JBoss runtime parameters include both the user defined and default values for the JVM parameter. However, JBoss runs with the default value and ignores any user defined value.

h) Save the `site-options.conf` file.

## 11. Deploy services.

a) Move all backup files and folders including the error folder to `/tmp`.

 **Note:** This step is required to make sure that there are no backup files and folders present in `/opt/collabnet/teamforge/var/james/var/mail/` to avoid any delay during provisioning.

```
cd /opt/collabnet/teamforge/var/james/var/mail/
mv * /tmp
```

b) Run the TeamForge provision command.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "Y" or "N" to proceed. For more information, see [TeamForge script](#) on page 785.


TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

## 12. Run the `/var/lib/pgsql/analyze_new_cluster.sh` script.

```
su - postgres -c "/var/lib/pgsql/analyze_new_cluster.sh"
```

## 13. Update the file permissions on your site's data.

```
/opt/collabnet/teamforge/runtime/scripts/fix_data_permissions.sh
```

 **Note:** This process can take a long time on sites with a lot of data.

## Do this step by step on Codesearch Server (server-03)

### 14. Upgrade the operating system packages.

```
yum upgrade
```

### 15. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.


a) Verify if SELinux is running in enforcing mode.

```
getenforce
```



- b) If the output of the `getenforce` command is either "Disabled" or "Permissive", SELinux is already disabled.
- c) If not disabled, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

16.  **Attention:** Do this if and only if you are upgrading TeamForge on RHEL/CentOS 6.8.

Delete the `python-crypto` package.

```
yum erase python-crypto
```

17. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.
- b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

#### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 6.8 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel6.x86\_64.rpm**
- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -Uvh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
```

```
yum list apr
```

### 18. Stop TeamForge.

If you are upgrading from TeamForge 16.7 or earlier releases:

```
/etc/init.d/collabnet stop all
```

If you are upgrading from TeamForge 16.10 or later releases:

```
/opt/collabnet/teamforge/bin/teamforge stop
```

### 19. Install the Codesearch packages.

```
yum install teamforge-codesearch
```

### 20. Copy the `site-options.conf` file from the TeamForge Application Server(server-01) to the Codesearch Server's `/opt/collabnet/teamforge/etc/` directory.

### 21. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



#### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

### Do this step by step on TeamForge Application Server (server-01)

### 22. If you have CVS integrations, synchronize permissions post upgrade. See, [Synchronize TeamForge source control integrations](#) on page 695.

### 23. Run the TeamForge post installation script. For more information, see [post-install.py](#).

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

### Do this step by step on EventQ Server (server-02)

### 24. Stop EventQ.

```
/etc/init.d/eventq stop
/etc/init.d/collabnet-rabbitmq-server stop
/etc/init.d/collabnet-mongod stop
```

### 25. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.

a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

b) If the output of the `getenforce` command is either "Disabled" or "Permissive", SELinux is already disabled.

c) If not disabled, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

### 26. Copy the `site-options.conf` file from the TeamForge Application Server(server-01) to the EventQ Server's `/opt/collabnet/teamforge/etc/` directory.

### 27. Upgrade the operating system packages.

```
yum upgrade
```

### 28. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.

b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD.

The DVD contains the necessary software and utilities required for installing TeamForge without internet access. In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

## 29. Install the EventQ packages.

```
yum install teamforge-eventq CN-eventq
```

## 30. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

### Do this step by step on TeamForge Application Server (server-01)

#### 31. Verify TeamForge installation.

- a) Log on to the TeamForge web application using the default Admin credentials.
  - Username: "admin"
  - Password: "admin"
- b) If your site has custom branding, verify that your branding changes still work as intended.  
See [Customize TeamForge](#) .
- c) Let your site's users know they've been upgraded.  
See [Create a Side-wide Broadcast](#).

32. Remove the backup files, if any, after the TeamForge site is up and running as expected. Remove the repository and the file system backup from the `/tmp/backup_dir` directory.

#### 33. Fix the TeamForge Avatar display issue on RHEL/CentOS 6.8.

TeamForge Avatar image is not displayed properly post Review Board installation on RHEL/CentOS 6.8. Run the following commands to work around this issue:

```
yum erase python-imaging
yum install teamforge
service httpd restart
```

### Post Upgrade Tasks

- [Upgrade Review Board](#)
- [Run TeamForge in SELinux enabled mode](#)
- [Add EventQ to existing projects](#)
- [Run ETL initial load jobs](#)

### Upgrade TeamForge, EventQ, and Git on Three Separate Servers

Three-server distributed setup, same hardware upgrade with TeamForge, EventQ, and Git services on three separate servers.

#### Dos and Don'ts

Check [this](#) list while installing or upgrading TeamForge.

#### Three Server Setup

Here's how the services are distributed across three separate RHEL/CentOS 6.8/7.3 servers.


| Services        | TeamForge Application Server (server-01) | EventQ Server (server-02) | Git Integration Server (server-03) |
|-----------------|------------------------------------------|---------------------------|------------------------------------|
| ctfcore         | ✓                                        |                           |                                    |
| mail            | ✓                                        |                           |                                    |
| search          | ✓                                        |                           |                                    |
| codesearch      | ✓                                        |                           |                                    |
| etl             | ✓                                        |                           |                                    |
| gerrit          |                                          |                           | ✓                                  |
| gerrit-database |                                          |                           | ✓                                  |

| Services                   | TeamForge Application Server (server-01) | EventQ Server (server-02) | Git Integration Server (server-03) |
|----------------------------|------------------------------------------|---------------------------|------------------------------------|
| reviewboard                | ✓                                        |                           |                                    |
| reviewboard-database       | ✓                                        |                           |                                    |
| <b>reviewboard-adapter</b> | ✓                                        |                           |                                    |
| subversion                 | ✓                                        |                           |                                    |
| cvs                        | ✓                                        |                           |                                    |
| binary                     | ✓                                        |                           |                                    |
| binary-database            | ✓                                        |                           |                                    |
| <b>cliserver</b>           | ✓                                        |                           |                                    |
| eventq                     |                                          | ✓                         |                                    |
| mongodb                    |                                          | ✓                         |                                    |
| redis                      |                                          | ✓                         |                                    |
| rabbitmq                   |                                          | ✓                         |                                    |
| ctfcore-database           | ✓                                        |                           |                                    |
| ctfcore-datamart           | ✓                                        |                           |                                    |

No backup is required for some hardware upgrades. However, you can create a backup as a measure of caution. See [Back up and restore TeamForge and EventQ](#) on page 566 to learn more about backing up TeamForge and EventQ database and file system.

### Do this step by step on TeamForge Application Server (server-01)

1. Back up all your custom event handlers and remove all the event handler JAR files before starting your TeamForge 17.4 upgrade process.
  - a) Go to **My Workspace > Admin**.
  - b) Click **System Tools** from the **Projects** menu.
  - c) Click **Customizations**.
  - d) Select the custom event handler and click **Delete**.

 **Important:** Post upgrade, you can add custom event handlers again from the backup while making sure that you don't have SOAP50 (deprecated) library used.

2. Uninstall hotfixes and add-ons, if any, installed on your site.
3. If you have Review Board installed, uninstall it.

```
cd /opt/collabnet/RBInstaller-17.1.5
python ./install.py -u
```

4. Stop TeamForge.

If you are upgrading from TeamForge 16.7 or earlier releases:

```
/etc/init.d/collabnet stop all
```

If you are upgrading from TeamForge 16.10 or later releases:

```
/opt/collabnet/teamforge/bin/teamforge stop
```

5. Upgrade the operating system packages.

```
yum upgrade
```

6. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge.

If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.

a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

b) If the output of the `getenforce` command is "permissive", continue with the next step. If not, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

7. Configure your TeamForge installation repository.

**TeamForge installation repository configuration for sites with internet access**

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.
- b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

c) Refresh your repository cache.

```
yum clean all
```

**TeamForge installation repository configuration for sites without internet access**

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 6.8 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel6.x86\_64.rpm**
- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -Uvh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

8. Install the TeamForge application packages.

```
yum install teamforge
```



**Attention:** TeamForge installer has been optimized quite a bit. It's likely that you might come across a lot of warning messages while upgrading from TeamForge 8.2 (or earlier) to TeamForge 17.4 on the same hardware (when you run the `yum install teamforge` command). You can safely ignore such warning messages and proceed with the upgrade.

9. Install the Binary application packages.

```
yum install cn-binary
```

10. Set up your site's master configuration file.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

a) **host: SERVICES Token**



**Note:** `cliserver` and `reviewboard-adapter` are the new services added in TeamForge 17.4.

- ```
server-01:SERVICES=ctfcore ctfcore-database ctfcore-datamart mail etl
search
codesearch subversion cvs cliserver binary binary-database
reviewboard reviewboard-database
reviewboard-adapter
```
- ```
server-02:SERVICES=eventq mongodb redis rabbitmq
```
- ```
server-03:SERVICES=gerrit gerrit-database
```



**Note:** You may remove the identifiers of components you do not want. For example, remove `binary` and `binary-database` if you are not planning to install binary repository managers such as Nexus.

b) **host: PUBLIC\_FQDN Token**

```
server-01:PUBLIC_FQDN=my.app.domain.com
```

c) Set the `MONGODB_APP_DATABASE_NAME` token with EventQ's database name in the `site-options.conf` file.

```
MONGODB_APP_DATABASE_NAME=orchestrate
```

d) **SSL Tokens**

SSL is enabled by default and a self-signed certificate is auto-generated. Use the following tokens to adjust this behavior.

```
SSL_CERT_FILE=
SSL_KEY_FILE=
SSL_CHAIN_FILE=
```

- To generate the SSL certificates, see [Generate SSL certificates](#) on page 718.
- Have the custom SSL certificate and private key for custom SSL certificate in place and provide their absolute paths in these tokens. `SSL_CHAIN_FILE` (intermediate certificate) is optional.
- All SSL certificates including self-signed certificates are added automatically.

#### e) Password Tokens

- TeamForge 7.1 and later support automatic password creation. See [AUTO\\_DATA](#) on page 736 for more information.
- If the token `REQUIRE_PASSWORD_SECURITY` is enabled, then set a value for the token, `PASSWORD_CONTROL_EFFECTIVE_DATE`.



**Warning:** The Password Control Kit (PCK) disables, deletes or expires user accounts that don't meet the password security requirements starting from the date set for the `PASSWORD_CONTROL_EFFECTIVE_DATE` token. If a date is not set, the PCK disables, deletes or expires user accounts immediately. See [PASSWORD\\_CONTROL\\_EFFECTIVE\\_DATE](#) for more information.

- If the token `REQUIRE_RANDOM_ADMIN_PASSWORD` is already set to true, then set the token `ADMIN_EMAIL` with a valid email address.

```
ADMIN_EMAIL=root@{__APPLICATION_HOST__}
```

- If you have LDAP set up for external authentication, you must set the `“REQUIRE_USER_PASSWORD_CHANGE”` `site-options.conf` token to false.

#### f) PostgreSQL Tokens and Settings

- Enable the `POSTGRES_INTERFACE` token in the `site-options.conf` file if the value of "ifconfig - a" is not eth0/ens0\*.
- Make sure the PostgreSQL tokens in the `site-options.conf` file are set as recommended in the following topic: [What are the right PostgreSQL settings for my site?](#) on page 943

#### g) JAVA\_OPTS

Configure the `JBOSS_JAVA_OPTS` `site-options.conf` token. See [JBOSS\\_JAVA\\_OPTS](#) on page 747.



**Note:** All JVM parameters but `-Xms1024m` and `-Xmx2048m` have been hard-coded in the TeamForge core application. You cannot manually configure any of the following default JVM parameters in the `site-options.conf` file.

- `-XX:+UseParallelGC`
- `-XX:MaxMetaspaceSize=512m`
- `-XX:ReservedCodeCacheSize=128M`
- `-server`
- `-XX:+HeapDumpOnOutOfMemoryError`
- `-XX:HeapDumpPath=/tmp -verbose:gc`
- `-XX:+PrintCodeCache`
- `-Djsse.enableSNIExtension=false`
- `-Dsun.rmi.dgc.client.gcInterval=600000`
- `-Dsun.rmi.dgc.server.gcInterval=600000`
- `-Djava.security.egd=file:/dev/urandom`
- `-Djava.awt.headless=true`.

When you change the default value of a JVM parameter such as `"-XX:HeapDumpPath"`, the JBoss runtime parameters include both the user defined and default values for the JVM parameter. However, JBoss runs with the default value and ignores any user defined value.

- h) Save the `site-options.conf` file.

### 11. Deploy services.

- a) Move all backup files and folders including the error folder to `/tmp`.



- 👉 **Note:** This step is required to make sure that there are no backup files and folders present in `/opt/collabnet/teamforge/var/james/var/mail/` to avoid any delay during provisioning.

```
cd /opt/collabnet/teamforge/var/james/var/mail/
mv * /tmp
```

- b) Run the TeamForge provision command.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

- 👉 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "Y" or "N" to proceed. For more information, see [TeamForge script](#) on page 785.

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

12. Run the `/var/lib/pgsql/analyze_new_cluster.sh` script.

```
su - postgres -c "/var/lib/pgsql/analyze_new_cluster.sh"
```

13. Update the file permissions on your site's data.

```
/opt/collabnet/teamforge/runtime/scripts/fix_data_permissions.sh
```

- 👉 **Note:** This process can take a long time on sites with a lot of data.

### Do this step by step on Git Integration Server (server-03)

14. Back up the Git file system data.

- a) Make an archive file with the following data directories.

- 👉 **Tip:** `/tmp` is just an example. You can use any directory or partition that you prefer.

Directory	Contents
<code>/gitroot</code>	Git source code repositories

```
cp -Rpfv /gitroot /tmp/gitbackup_dir
cp -Rpfv /opt/collabnet/gerrit/ /tmp/gitbackup_dir/gerrit
```

- b) Back up your SSH keys, if any.

15. Upgrade the operating system packages.

```
yum upgrade
```

16. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is either "Disabled" or "Permissive", SELinux is already disabled.

- c) If not disabled, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

17. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.
- b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.

- Red Hat Enterprise Linux/CentOS 6.8 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel6.x86\_64.rpm**
- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -Uvh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

### 18. Install the Git packages.

```
yum install teamforge-git
```

19. Copy the `site-options.conf` file from the TeamForge Application Server (server-01) to the Git Server's /opt/collabnet/teamforge/etc/ directory.

### 20. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



#### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

21. Run the `/var/lib/pgsql/analyze_new_cluster.sh` script.

```
su - postgres -c "/var/lib/pgsql/analyze_new_cluster.sh"
```

**Do this step by step on TeamForge Application Server (server-01)**

22. If you have CVS integrations, synchronize permissions post upgrade. See, [Synchronize TeamForge source control integrations](#) on page 695.

23. Run the TeamForge post installation script. For more information, see [post-install.py](#).

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

**Do this step by step on EventQ Server (server-02)**

24. Stop EventQ.

```
/etc/init.d/eventq stop
/etc/init.d/collabnet-rabbitmq-server stop
/etc/init.d/collabnet-mongod stop
```

25. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.

a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

b) If the output of the `getenforce` command is either "Disabled" or "Permissive", SELinux is already disabled.

c) If not disabled, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

26. Upgrade the operating system packages.

```
yum upgrade
```

27. Configure your TeamForge installation repository.

**TeamForge installation repository configuration for sites with internet access**

a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.

b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

c) Refresh your repository cache.

```
yum clean all
```

**TeamForge installation repository configuration for sites without internet access**

a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD.

The DVD contains the necessary software and utilities required for installing TeamForge without internet access. In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDROM
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

28. Install the EventQ packages.

```
yum install teamforge-eventq CN-eventq
```

29. Copy the `site-options.conf` file from the TeamForge Application Server (server-01) to the EventQ Server's `/opt/collabnet/teamforge/etc/` directory.

30. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



**Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

**Do this step by step on TeamForge Application Server (server-01)**

31. Verify TeamForge installation.

- a) Log on to the TeamForge web application using the default Admin credentials.
  - Username: "admin"
  - Password: "admin"
- b) If your site has custom branding, verify that your branding changes still work as intended. See [Customize TeamForge](#).
- c) Let your site's users know they've been upgraded. See [Create a Side-wide Broadcast](#).

32. Remove the backup files, if any, after the TeamForge site is up and running as expected. Remove the repository and the file system backup from the `/tmp/backup_dir` directory.

33. Fix the TeamForge Avatar display issue on RHEL/CentOS 6.8.

TeamForge Avatar image is not displayed properly post Review Board installation on RHEL/CentOS 6.8. Run the following commands to work around this issue:

```
yum erase python-imaging
yum install teamforge
service httpd restart
```

## Post Upgrade Tasks

- [Upgrade Review Board](#)
- [Run TeamForge in SELinux enabled mode](#)
- [Add EventQ to existing projects](#)
- [Run ETL initial load jobs](#)

## Upgrade TeamForge, EventQ, and SCM/Git on Three Separate Servers

Three-server distributed setup, same hardware upgrade, with TeamForge, EventQ, and SCM/Git services installed on three separate servers.

### Dos and Don'ts

Check [this](#) list while installing or upgrading TeamForge.

### Three Server Setup

Here's how the services are distributed across three separate RHEL/CentOS 6.8/7.3 servers.


Services	TeamForge Application Server (server-01)	EventQ Server (server-02)	SCM and Git Server (server-03)
ctfcore	✓		
mail	✓		
search	✓		
codesearch	✓		
etl	✓		
gerrit			✓
gerrit-database			✓
reviewboard	✓		
reviewboard-database	✓		
<b>reviewboard-adapter</b>	✓		
subversion			✓
cvs			✓
binary	✓		
binary-database	✓		
<b>cliserver</b>	✓		
eventq		✓	
mongodb		✓	
redis		✓	

Services	TeamForge Application Server (server-01)	EventQ Server (server-02)	SCM and Git Server (server-03)
rabbitmq		✓	
ctfcore-database	✓		
ctfcore-datamart	✓		

No backup is required for same hardware upgrades. However, you can create a backup as a measure of caution. See [Back up and restore TeamForge and EventQ](#) on page 566 to learn more about backing up TeamForge and EventQ database and file system.

### Do this step-by-step on TeamForge Application Server (server-01)

1. Back up all your custom event handlers and remove all the event handler JAR files before starting your TeamForge 17.4 upgrade process.
  - a) Go to **My Workspace > Admin**.
  - b) Click **System Tools** from the **Projects** menu.
  - c) Click **Customizations**.
  - d) Select the custom event handler and click **Delete**.

 **Important:** Post upgrade, you can add custom event handlers again from the backup while making sure that you don't have SOAP50 (deprecated) library used.

2. Uninstall hotfixes and add-ons, if any, installed on your site.
3. If you have Review Board installed, uninstall it.

```
cd /opt/collabnet/RBInstaller-17.1.5
python ./install.py -u
```

4. Stop TeamForge.

If you are upgrading from TeamForge 16.7 or earlier releases:

```
/etc/init.d/collabnet stop all
```

If you are upgrading from TeamForge 16.10 or later releases:

```
/opt/collabnet/teamforge/bin/teamforge stop
```

5. Upgrade the operating system packages.

```
yum upgrade
```

6. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge.

If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is "permissive", continue with the next step. If not, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

7. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.
- b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.

- Red Hat Enterprise Linux/CentOS 6.8 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel6.x86\_64.rpm**
- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -Uvh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

8. Install the TeamForge application packages.

```
yum install teamforge
```



**Attention:** TeamForge installer has been optimized quite a bit. It's likely that you might come across a lot of warning messages while upgrading from TeamForge 8.2 (or earlier) to TeamForge 17.4 on the same hardware (when you run the `yum install teamforge` command). You can safely ignore such warning messages and proceed with the upgrade.


9. Install the Binary application packages.

```
yum install cn-binary
```


10. Set up your site's master configuration file.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

a) **host: SERVICES Token**

 **Note:** cliserver and reviewboard-adapter are the new services added in TeamForge 17.4.

- ```
server-01:SERVICES=ctfcore ctfcore-database ctfcore-datamart mail etl
search
codesearch cliserver binary binary-database reviewboard reviewboard-
database reviewboard-adapter
```
- ```
server-02:SERVICES=eventq mongodb redis rabbitmq
```
- ```
server-03:SERVICES=subversion cvs gerrit gerrit-database
```

 **Note:** You may remove the identifiers of components you do not want. For example, remove binary and binary-database if you are not planning to install binary repository managers such as Nexus.

b) **host: PUBLIC\_FQDN Token**

```
server-01:PUBLIC_FQDN=my.app.domain.com
```

```
server-03:PUBLIC_FQDN=my.scmangit.domain.com
```

- c) Set the `MONGODB_APP_DATABASE_NAME` token with EventQ's database name in the `site-options.conf` file.

```
MONGODB_APP_DATABASE_NAME=orchestrate
```

d) **SSL Tokens**

SSL is enabled by default and a self-signed certificate is auto-generated. Use the following tokens to adjust this behavior.

```
SSL_CERT_FILE=
SSL_KEY_FILE=
SSL_CHAIN_FILE=
```

- To generate the SSL certificates, see [Generate SSL certificates](#) on page 718.
- Have the custom SSL certificate and private key for custom SSL certificate in place and provide their absolute paths in these tokens. `SSL_CHAIN_FILE` (intermediate certificate) is optional.
- All SSL certificates including self-signed certificates are added automatically.

e) **Password Tokens**

- TeamForge 7.1 and later support automatic password creation. See [AUTO\\_DATA](#) on page 736 for more information.
- If the token `REQUIRE_PASSWORD_SECURITY` is enabled, then set a value for the token, `PASSWORD_CONTROL_EFFECTIVE_DATE`.



**Warning:** The Password Control Kit (PCK) disables, deletes or expires user accounts that don't meet the password security requirements starting from the date set for the `PASSWORD_CONTROL_EFFECTIVE_DATE` token. If a date is not set, the PCK disables, deletes or expires user accounts immediately. See [PASSWORD\\_CONTROL\\_EFFECTIVE\\_DATE](#) for more information.

- If the token `REQUIRE_RANDOM_ADMIN_PASSWORD` is already set to true, then set the token `ADMIN_EMAIL` with a valid email address.

```
ADMIN_EMAIL=root@{__APPLICATION_HOST__}
```

- If you have LDAP set up for external authentication, you must set the `“REQUIRE_USER_PASSWORD_CHANGE”` `site-options.conf` token to false.




f) **PostgreSQL Tokens and Settings**

- Enable the `POSTGRES_INTERFACE` token in the `site-options.conf` file if the value of "ifconfig - a" is not `eth0/enp0*`.
- Make sure the PostgreSQL tokens in the `site-options.conf` file are set as recommended in the following topic: *What are the right PostgreSQL settings for my site?* on page 943

g) **JAVA\_OPTS**

Configure the `JBOSS_JAVA_OPTS` `site-options.conf` token. See *JBOSS\_JAVA\_OPTS* on page 747.

 **Note:** All JVM parameters but `-Xms1024m` and `-Xmx2048m` have been hard-coded in the TeamForge core application. You cannot manually configure any of the following default JVM parameters in the `site-options.conf` file.


- `-XX:+UseParallelGC`
- `-XX:MaxMetaspaceSize=512m`
- `-XX:ReservedCodeCacheSize=128M`
- `-server`
- `-XX:+HeapDumpOnOutOfMemoryError`
- `-XX:HeapDumpPath=/tmp -verbose:gc`
- `-XX:+PrintCodeCache`
- `-Djsse.enableSNIExtension=false`
- `-Dsun.rmi.dgc.client.gcInterval=600000`
- `-Dsun.rmi.dgc.server.gcInterval=600000`
- `-Djava.security.egd=file:/dev/urandom`
- `-Djava.awt.headless=true`.

When you change the default value of a JVM parameter such as `"-XX:HeapDumpPath"`, the JBoss runtime parameters include both the user defined and default values for the JVM parameter. However, JBoss runs with the default value and ignores any user defined value.

h) Save the `site-options.conf` file.

**11. Deploy services.**

a) Move all backup files and folders including the error folder to `/tmp`.

 **Note:** This step is required to make sure that there are no backup files and folders present in `/opt/collabnet/teamforge/var/james/var/mail/` to avoid any delay during provisioning.

```
cd /opt/collabnet/teamforge/var/james/var/mail/
mv * /tmp
```

b) Run the TeamForge provision command.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "Y" or "N" to proceed. For more information, see *TeamForge script* on page 785.

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

**12. Run the `/var/lib/pgsql/analyze_new_cluster.sh` script.**

```
su - postgres -c "/var/lib/pgsql/analyze_new_cluster.sh"
```

**Do this step by step on SCM and Git Server (server-03)****13. Stop TeamForge.**

If you are upgrading from TeamForge 16.7 or earlier releases:

```
/etc/init.d/collabnet stop all
```

If you are upgrading from TeamForge 16.10 or later releases:

```
/opt/collabnet/teamforge/bin/teamforge stop
```

**14. Upgrade the operating system packages.**

```
yum upgrade
```

**15. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.**

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is either "Disabled" or "Permissive", SELinux is already disabled.  
c) If not disabled, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

**16. Configure your TeamForge installation repository.**

**TeamForge installation repository configuration for sites with internet access**

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.  
b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

**TeamForge installation repository configuration for sites without internet access**

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 6.8 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel6.x86\_64.rpm**
- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -Uvh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/  
mkdir cdrom  
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
```

```
name=RHEL CDROM
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

17. Install the Source Code application packages.

```
yum install teamforge-scm
```

18. Install the Git application packages.

```
yum install teamforge-git
```

19. Copy the `site-options.conf` file from the TeamForge Application Server (server-01) to the SCM and Git Server's `/opt/collabnet/teamforge/etc/` directory.

20. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



**Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

During teamforge provision, the Register SCM integration process fails on sites that use self-signed certificates. Perform these steps in such cases.

- a) Restart JBoss on the TeamForge Application server.

```
/opt/collabnet/teamforge/bin/teamforge restart -s jboss
```

- b) Reinitialize TeamForge on the SCM and Git server.

```
/opt/collabnet/teamforge/bin/teamforge reinitialize
```

21. Run the `/var/lib/pgsql/analyze_new_cluster.sh` script.

```
su - postgres -c "/var/lib/pgsql/analyze_new_cluster.sh"
```

**Do this step by step on TeamForge Application Server (server-01)**

22. If you have CVS integrations, synchronize permissions post upgrade. See, [Synchronize TeamForge source control integrations](#) on page 695.

23. Run the TeamForge post installation script. For more information, see [post-install.py](#).

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

**Do this step by step on SCM and Git Server (server-03)**

24. Update the file permissions on your site's data.

```
/opt/collabnet/teamforge/runtime/scripts/fix_data_permissions.sh
```



**Note:** This process can take a long time on sites with a lot of data.

**Do this step by step on TeamForge Application Server (server-01)**

25. Update the file permissions on your site's data.

```
/opt/collabnet/teamforge/runtime/scripts/fix_data_permissions.sh
```



**Note:** This process can take a long time on sites with a lot of data.

**Do this step by step on EventQ Server (server-02)****26. Stop EventQ.**

```
/etc/init.d/eventq stop
/etc/init.d/collabnet-rabbitmq-server stop
/etc/init.d/collabnet-mongod stop
```

**27. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.**

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is either "Disabled" or "Permissive", SELinux is already disabled.

- c) If not disabled, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

**28. Upgrade the operating system packages.**

```
yum upgrade
```

**29. Configure your TeamForge installation repository.****TeamForge installation repository configuration for sites with internet access**

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to /tmp.  
b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

**TeamForge installation repository configuration for sites without internet access**

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD.

The DVD contains the necessary software and utilities required for installing TeamForge without internet access. In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

**30.** Install the EventQ packages.

```
yum install teamforge-eventq CN-eventq
```

**31.** Copy the `site-options.conf` file from the TeamForge Application Server (server-01) to the EventQ Server's `/opt/collabnet/teamforge/etc/` directory.

**32.** Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



**Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

**Do this step by step on TeamForge Application Server(server-01)**

**33.** Use the following URL and set the "svn-internal" **SOAP SERVICE PORT** to 8080 and click **Save**.

```
https://<hostname>/sf/sfmain/do/editSystem?systemId=exsy1001
```

Site Administration / Edit System

## Edit Integration

| Type:                        | Subversion                                |
|------------------------------|-------------------------------------------|
| NAME*                        | svn-internal                              |
| DESCRIPTION:                 | Subversion Integration                    |
| SOAP SERVICE HOST:           | localhost                                 |
| SOAP SERVICE PORT:           | 8080                                      |
| USE SECURE CONNECTION (SSL): | <input type="checkbox"/>                  |
| REPOSITORY BASE URL:         | https://localhost/svn/repository-internal |

### 34. Verify TeamForge installation.

- a) Log on to the TeamForge web application using the default Admin credentials.
  - Username: "admin"
  - Password: "admin"
- b) If your site has custom branding, verify that your branding changes still work as intended.  
See [Customize TeamForge](#).
- c) Let your site's users know they've been upgraded.  
See [Create a Side-wide Broadcast](#).

### 35. Fix the TeamForge Avatar display issue on RHEL/CentOS 6.8.

TeamForge Avatar image is not displayed properly post Review Board installation on RHEL/CentOS 6.8. Run the following commands to work around this issue:

```
yum erase python-imaging
yum install teamforge
service httpd restart
```

- [Upgrade Review Board](#)
- [Run TeamForge in SELinux enabled mode](#)
- [Add EventQ to existing projects](#)
- [Run ETL initial load jobs](#)

### Upgrade TeamForge, EventQ, Database, and SCM on Four Separate Servers

Four-server setup, same hardware upgrade, with TeamForge, EventQ, Database (Database and Datamart), and SCM on four separate servers.

## Dos and Don'ts

Check [this](#) list before installing or upgrading TeamForge.

## Four Server Setup

Here's how the services are distributed across four separate RHEL/CentOS 6.8/7.3 servers.


| Services                   | TeamForge Application Server (server-01) | EventQ Server (server-02) | Database Server(server-03) | SCM and Git Server (server-04) |
|----------------------------|------------------------------------------|---------------------------|----------------------------|--------------------------------|
| ctfcore                    | ✓                                        |                           |                            |                                |
| mail                       | ✓                                        |                           |                            |                                |
| search                     | ✓                                        |                           |                            |                                |
| codesearch                 | ✓                                        |                           |                            |                                |
| etl                        | ✓                                        |                           |                            |                                |
| gerrit                     |                                          |                           |                            | ✓                              |
| gerrit-database            |                                          |                           |                            | ✓                              |
| reviewboard                | ✓                                        |                           |                            |                                |
| reviewboard-database       | ✓                                        |                           |                            |                                |
| <b>reviewboard-adapter</b> | ✓                                        |                           |                            |                                |
| subversion                 |                                          |                           |                            | ✓                              |
| cvs                        |                                          |                           |                            | ✓                              |
| binary                     | ✓                                        |                           |                            |                                |
| binary-database            | ✓                                        |                           |                            |                                |
| <b>cliserver</b>           | ✓                                        |                           |                            |                                |
| eventq                     |                                          | ✓                         |                            |                                |
| mongodb                    |                                          | ✓                         |                            |                                |
| redis                      |                                          | ✓                         |                            |                                |
| rabbitmq                   |                                          | ✓                         |                            |                                |
| ctfcore-database           |                                          |                           | ✓                          |                                |
| ctfcore-datamart           |                                          |                           | ✓                          |                                |

No backup is required for same hardware upgrades. However, you can create a backup as a precaution. See [Back up and restore TeamForge and EventQ](#) on page 566 to learn more about backing up TeamForge and EventQ database and file system.

### Do this step by step on TeamForge Application Server (server-01)

1. Back up all your custom event handlers and remove all the event handler JAR files before starting your TeamForge 17.4 upgrade process.

- a) Go to **My Workspace > Admin**.
- b) Click **System Tools** from the **Projects** menu.
- c) Click **Customizations**.
- d) Select the custom event handler and click **Delete**.

 **Important:** Post upgrade, you can add custom event handlers again from the backup while making sure that you don't have SOAP50 (deprecated) library used.

2. Uninstall hotfixes and add-ons, if any, installed on your site.
3. If you have Review Board installed, uninstall it.

```
cd /opt/collabnet/RBInstaller-17.1.5
python ./install.py -u
```

4. Stop TeamForge.

If you are upgrading from TeamForge 16.7 or earlier releases:

```
/etc/init.d/collabnet stop all
```

If you are upgrading from TeamForge 16.10 or later releases:

```
/opt/collabnet/teamforge/bin/teamforge stop
```

5. Upgrade the operating system packages.

```
yum upgrade
```

6. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge.

If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is "permissive", continue with the next step. If not, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

7. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.
- b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

#### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 6.8 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel6.x86\_64.rpm**



- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
  - In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.
- b) Unpack the disconnected installation package.
- ```
rpm -Uvh <package-name>
```
- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

## 8. Install the TeamForge application packages.

```
yum install teamforge
```



**Attention:** TeamForge installer has been optimized quite a bit. It's likely that you might come across a lot of warning messages while upgrading from TeamForge 8.2 (or earlier) to TeamForge 17.4 on the same hardware (when you run the `yum install teamforge` command). You can safely ignore such warning messages and proceed with the upgrade.

## 9. Install the Binary application packages.

```
yum install cn-binary
```

### Do this step by step on Database Server (server-03)

## 10. Upgrade the operating system packages.

```
yum upgrade
```

## 11. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is either "Disabled" or "Permissive", SELinux is already disabled.
- c) If not disabled, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

## 12. Configure your TeamForge installation repository.

### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to /tmp.
- b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.

- Red Hat Enterprise Linux/CentOS 6.8 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel6.x86\_64.rpm**
- Red Hat Enterprise Linux/CentOS 7.3 64: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -Uvh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

## 13. Stop TeamForge.

If you are upgrading from TeamForge 16.7 or earlier releases:

```
/etc/init.d/collabnet stop all
```

If you are upgrading from TeamForge 16.10 or later releases:

```
/opt/collabnet/teamforge/bin/teamforge stop
```


#### 14. Install the TeamForge database packages.

```
yum install teamforge-database
```

#### 15. Set up your site's master configuration file.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

##### a) host: SERVICES Token


 **Note:** cliserver and reviewboard-adapter are the new services added in TeamForge 17.4.

```
server-01:SERVICES=ctfcore mail etl search codesearch cliserver binary
binary-database reviewboard reviewboard-database reviewboard-adapter
```

```
server-02:SERVICES=eventq mongodb redis rabbitmq
```

```
server-03:SERVICES=ctfcore-database ctfcore-datamart
```

```
server-04:SERVICES=subversion cvs gerrit gerrit-database
```

 **Note:** You may remove the identifiers of components you do not want. For example, remove binary and binary-database if you are not planning to install binary repository managers such as Nexus.

##### b) host: PUBLIC\_FQDN Token

```
server-01:PUBLIC_FQDN=my.app.domain.com
```

```
server-04:PUBLIC_FQDN=my.scmangit.domain.com
```

##### c) Set the `MONGODB_APP_DATABASE_NAME` token with EventQ's database name in the `site-options.conf` file.

```
MONGODB_APP_DATABASE_NAME=orchestrate
```

##### d) SSL Tokens

SSL is enabled by default and a self-signed certificate is auto-generated. Use the following tokens to adjust this behavior.

```
SSL_CERT_FILE=
SSL_KEY_FILE=
SSL_CHAIN_FILE=
```

- To generate the SSL certificates, see [Generate SSL certificates](#) on page 718.
- Have the custom SSL certificate and private key for custom SSL certificate in place and provide their absolute paths in these tokens. `SSL_CHAIN_FILE` (intermediate certificate) is optional.
- All SSL certificates including self-signed certificates are added automatically.

##### e) Password Tokens

- TeamForge 7.1 and later support automatic password creation. See [AUTO\\_DATA](#) on page 736 for more information.
- If the token `REQUIRE_PASSWORD_SECURITY` is enabled, then set a value for the token, `PASSWORD_CONTROL_EFFECTIVE_DATE`.



**Warning:** The Password Control Kit (PCK) disables, deletes or expires user accounts that don't meet the password security requirements starting from the date set for the `PASSWORD_CONTROL_EFFECTIVE_DATE` token. If a date is not set, the PCK disables, deletes

or expires user accounts immediately. See [PASSWORD\\_CONTROL\\_EFFECTIVE\\_DATE](#) for more information.

- If the token [REQUIRE\\_RANDOM\\_ADMIN\\_PASSWORD](#) is already set to true, then set the token [ADMIN\\_EMAIL](#) with a valid email address.

```
ADMIN_EMAIL=root@{__APPLICATION_HOST__}
```


- If you have LDAP set up for external authentication, you must set the ["REQUIRE\\_USER\\_PASSWORD\\_CHANGE"](#) `site-options.conf` token to false.

#### f) PostgreSQL Tokens and Settings

- Enable the [POSTGRES\\_INTERFACE](#) token in the `site-options.conf` file if the value of "ifconfig - a" is not eth0/enp0\*.
- Make sure the PostgreSQL tokens in the `site-options.conf` file are set as recommended in the following topic: [What are the right PostgreSQL settings for my site?](#) on page 943

#### g) JAVA\_OPTS

Configure the `JBOSS_JAVA_OPTS` `site-options.conf` token. See [JBOSS\\_JAVA\\_OPTS](#) on page 747.

 **Note:** All JVM parameters but `-Xms1024m` and `-Xmx2048m` have been hard-coded in the TeamForge core application. You cannot manually configure any of the following default JVM parameters in the `site-options.conf` file.

- `-XX:+UseParallelGC`
- `-XX:MaxMetaspaceSize=512m`
- `-XX:ReservedCodeCacheSize=128M`
- `-server`
- `-XX:+HeapDumpOnOutOfMemoryError`
- `-XX:HeapDumpPath=/tmp -verbose:gc`
- `-XX:+PrintCodeCache`
- `-Djsse.enableSNIExtension=false`
- `-Dsun.rmi.dgc.client.gcInterval=600000`
- `-Dsun.rmi.dgc.server.gcInterval=600000`
- `-Djava.security.egd=file:/dev/urandom`
- `-Djava.awt.headless=true`

When you change the default value of a JVM parameter such as `"-XX:HeapDumpPath"`, the JBoss runtime parameters include both the user defined and default values for the JVM parameter. However, JBoss runs with the default value and ignores any user defined value.

h) Save the `site-options.conf` file.

### 16. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

17. Run the `/var/lib/pgsql/analyze_new_cluster.sh` script.


```
su - postgres -c "/var/lib/pgsql/analyze_new_cluster.sh"
```

**Do this step by step on TeamForge Application Server (server-01)**

18. Copy the `site-options.conf` file from the Database Server to the TeamForge Application Server's `/opt/collabnet/teamforge/etc/` directory.

**19. Deploy services.**

- a) Move all backup files and folders including the error folder to /tmp.

 **Note:** This step is required to make sure that there are no backup files and folders present in /opt/collabnet/teamforge/var/james/var/mail/ to avoid any delay during provisioning.

```
cd /opt/collabnet/teamforge/var/james/var/mail/
mv * /tmp
```

- b) Run the TeamForge provision command.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "Y" or "N" to proceed. For more information, see [TeamForge script](#) on page 785.


TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

**20. Run the /var/lib/pgsql/analyze\_new\_cluster.sh script.**

```
su - postgres -c "/var/lib/pgsql/analyze_new_cluster.sh"
```

**21. Update the file permissions on your site's data.**

```
/opt/collabnet/teamforge/runtime/scripts/fix_data_permissions.sh
```

 **Note:** This process can take a long time on sites with a lot of data.

**22. If you have CVS integrations, synchronize permissions post upgrade. See, [Synchronize TeamForge source control integrations](#) on page 695.****23. Run the TeamForge post installation script. For more information, see [post-install.py](#).**

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

**Do this step by step on SCM and Git Integration Server (server-04)****24. Stop TeamForge.**

If you are upgrading from TeamForge 16.7 or earlier releases:

```
/etc/init.d/collabnet stop all
```

If you are upgrading from TeamForge 16.10 or later releases:

```
/opt/collabnet/teamforge/bin/teamforge stop
```

**25. Upgrade the operating system packages.**

```
yum upgrade
```

**26. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.**

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is either "Disabled" or "Permissive", SELinux is already disabled.

- c) If not disabled, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

**27. Configure your TeamForge installation repository.****TeamForge installation repository configuration for sites with internet access**

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to /tmp.  
b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.

- Red Hat Enterprise Linux/CentOS 6.8 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel6.x86\_64.rpm**
- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -Uvh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

## 28. Install the TeamForge SCM and Git packages.

```
yum install teamforge-scm teamforge-git
```

## 29. Copy the master site-options.conf file from the Database Server to the SCM and Git Integration Server's /opt/collabnet/teamforge/etc/ directory.

## 30. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

During teamforge provision, the Register SCM integration process fails on sites that use self-signed certificates. Perform these steps in such cases.

- a) Restart JBoss on the TeamForge Application server.

```
/opt/collabnet/teamforge/bin/teamforge restart -s jboss
```

- b) Reinitialize TeamForge on the SCM server.


```
/opt/collabnet/teamforge/bin/teamforge reinitialize
```

31. Run the `/var/lib/pgsql/analyze_new_cluster.sh` script.

```
su - postgres -c "/var/lib/pgsql/analyze_new_cluster.sh"
```

32. Update the file permissions on your site's data.

```
/opt/collabnet/teamforge/runtime/scripts/fix_data_permissions.sh
```

 **Note:** This process can take a long time on sites with a lot of data.

### Do this step-by-step on EventQ Server (server-02)

33. Stop EventQ.

```
/etc/init.d/eventq stop
/etc/init.d/collabnet-rabbitmq-server stop
/etc/init.d/collabnet-mongod stop
```

34. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is either "Disabled" or "Permissive", SELinux is already disabled.  
c) If not disabled, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

35. Upgrade the operating system packages.

```
yum upgrade
```

36. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.  
b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

#### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD.

The DVD contains the necessary software and utilities required for installing TeamForge without internet access. In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

### 37. Install the EventQ packages.

```
yum install teamforge-eventq CN-eventq
```

### 38. Copy the site-options.conf file from the Database Server (server-03) to the EventQ Server's /opt/collabnet/teamforge/etc/ directory.

### 39. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



#### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

### Do this step by step on TeamForge Application Server (server-01)

### 40. Use the following URL and set the "svn-internal" SOAP SERVICE PORT to 8080 and click Save.

```
https://<hostname>/sf/sfmain/do/editSystem?systemId=exsy1001
```



Admin ▾ Integrations ☰

Site Administration / Edit System

### Edit Integration

Type:	Subversion
NAME*:	svn-internal
DESCRIPTION:	Subversion Integration
SOAP SERVICE HOST:	localhost
SOAP SERVICE PORT:	8080
USE SECURE CONNECTION (SSL):	<input type="checkbox"/>
REPOSITORY BASE URL:	https://localhost/svn/repository-internal

#### 41. Verify TeamForge installation.

- a) Log on to the TeamForge web application using the default Admin credentials.
  - Username: "admin"
  - Password: "admin"
- b) If your site has custom branding, verify that your branding changes still work as intended.  
See [Customize TeamForge](#).
- c) Let your site's users know they've been upgraded.  
See [Create a Side-wide Broadcast](#).

#### Do this step on Database Server (server-03)

42. Remove the backup files, if any, after the TeamForge site is up and running as expected. Remove the repository and the file system backup from the `/tmp/backup_dir` directory.

#### Do this step on TeamForge Application Server (server-01)

43. Fix the TeamForge Avatar display issue on RHEL/CentOS 6.8.

TeamForge Avatar image is not displayed properly post Review Board installation on RHEL/CentOS 6.8. Run the following commands to work around this issue:

```
yum erase python-imaging
yum install teamforge
service httpd restart
```

#### Post Upgrade Tasks

- [Upgrade Review Board](#)
- [Run TeamForge in SELinux enabled mode](#)

- [Add EventQ to existing projects](#)
- [Run ETL initial load jobs](#)

### Upgrade TeamForge, EventQ, and Oracle Database on Three separate servers

Three-server setup, same hardware upgrade with TeamForge, EventQ, and Database and Datamart (Oracle) on three separate servers.

#### Dos and Don'ts

Check [this](#) list before installing or upgrading TeamForge.

#### Three Server Setup

Here's how the services are distributed across three separate RHEL/CentOS 6.8/7.3 servers.

Services	TeamForge Application Server (server-01)	EventQ Server (server-02)	Database Server (server-03)
ctfcore	✓		
mail	✓		
search	✓		
codesearch	✓		
etl	✓		
gerrit	✓		
gerrit-database	✓		
reviewboard	✓		
reviewboard-database	✓		
<b>reviewboard-adapter</b>	✓		
subversion	✓		
cvs	✓		
binary	✓		
binary-database	✓		
<b>cliserver</b>	✓		
eventq		✓	
mongodb		✓	
redis		✓	
rabbitmq		✓	
ctfcore-database			✓

Services	TeamForge Application Server (server-01)	EventQ Server (server-02)	Database Server (server-03)
ctfcore-datamart			✔


### Do this step on Database Server (Oracle) (server-03)

1. Make a dump file of your site database.

To back up the Oracle database, follow the [Oracle backup procedure](#).

### Do this step by step on TeamForge Application Server (server-01)

2. Back up all your custom event handlers and remove all the event handler JAR files before starting your TeamForge 17.4 upgrade process.
  - a) Go to **My Workspace > Admin**.
  - b) Click **System Tools** from the **Projects** menu.
  - c) Click **Customizations**.
  - d) Select the custom event handler and click **Delete**.

 **Important:** Post upgrade, you can add custom event handlers again from the backup while making sure that you don't have SOAP50 (deprecated) library used.

3. Uninstall hotfixes and add-ons, if any, installed on your site.
4. If you have Review Board installed, uninstall it.

```
cd /opt/collabnet/RBInstaller-17.1.5
python ./install.py -u
```

5. Stop TeamForge.


If you are upgrading from TeamForge 16.7 or earlier releases:

```
/etc/init.d/collabnet stop all
```

If you are upgrading from TeamForge 16.10 or later releases:

```
/opt/collabnet/teamforge/bin/teamforge stop
```

6. Back up the file system data.

 **Tip:** /tmp in the following step is just an example. You can use any directory or partition that you prefer.

- a) Make an archive file with the following data directories:

Directory	Contents
/opt/collabnet/teamforge/var	User-created data, such as artifact attachments
/svnroot	Subversion source code repositories
/sf-svnroot	Subversion repository for branding data
/cvsroot	CVS source code repositories (not present on all sites)

```
mkdir -p /tmp/backup_dir
cp -Rpfv /svnroot /sf-svnroot /cvsroot /opt/collabnet/teamforge/var /
tmp/backup_dir
```

- b) If GIT integration is enabled, do the following:

```
mkdir /tmp/backup_dir/gerrit
cp -Rpfv /gitroot /tmp/backup_dir
cp -Rpfv /opt/collabnet/gerrit/ /tmp/backup_dir/gerrit
```

- c) Back up your SSH keys, if any.

d) Back up your SSL certificates and keys, if any.

7. If Git is enabled, create a gerrit database dump file, for PostgreSQL upgrade.

```
su - postgres
/usr/bin/pg_dumpall > /var/lib/pgsql/9.3/backups/
teamforge_gerrit_data_backup.dmp
exit
```

8. Upgrade the operating system packages.

```
yum upgrade
```

9. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge.

If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.

a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

b) If the output of the `getenforce` command is "permissive", continue with the next step. If not, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

10. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.
- b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

c) Refresh your repository cache.

```
yum clean all
```

#### TeamForge installation repository configuration for sites without internet access

a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 6.8 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel6.x86\_64.rpm**
- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

b) Unpack the disconnected installation package.

```
rpm -Uvh <package-name>
```

c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

## 11. Install the TeamForge application packages.

```
yum install teamforge
```



**Attention:** TeamForge installer has been optimized quite a bit. It's likely that you might come across a lot of warning messages while upgrading from TeamForge 8.2 (or earlier) to TeamForge 17.4 on the same hardware (when you run the `yum install teamforge` command). You can safely ignore such warning messages and proceed with the upgrade.

## 12. Install the Binary application packages.

```
yum install cn-binary
```

## 13. Set up your site's master configuration file.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

### a) host: SERVICES Token



**Note:** `cliserver` and `reviewboard-adapter` are the new services added in TeamForge 17.4.

```
server-01:SERVICES=ctfcore mail etl search codesearch subversion cvs
cliserver gerrit
gerrit-database binary binary-database reviewboard reviewboard-database
reviewboard-adapter
```

```
server-02:SERVICES=eventq mongodb redis rabbitmq
```

```
server-03:SERVICES=ctfcore-database ctfcore-datamart
```



**Note:** You may remove the identifiers of components you do not want. For example, remove `binary` and `binary-database` if you are not planning to install binary repository managers such as Nexus.

### b) host: PUBLIC\_FQDN Token

```
server-01:PUBLIC_FQDN=my.app.domain.com
```

- c) Set the `MONGODB_APP_DATABASE_NAME` token with EventQ's database name in the `site-options.conf` file.

```
MONGODB_APP_DATABASE_NAME=orchestrate
```

### d) SSL Tokens

SSL is enabled by default and a self-signed certificate is auto-generated. Use the following tokens to adjust this behavior.

```
SSL_CERT_FILE=
SSL_KEY_FILE=
SSL_CHAIN_FILE=
```

- To generate the SSL certificates, see [Generate SSL certificates](#) on page 718.
- Have the custom SSL certificate and private key for custom SSL certificate in place and provide their absolute paths in these tokens. `SSL_CHAIN_FILE` (intermediate certificate) is optional.
- All SSL certificates including self-signed certificates are added automatically.

#### e) Password Tokens

- TeamForge 7.1 and later support automatic password creation. See [AUTO\\_DATA](#) on page 736 for more information.
- If the token `REQUIRE_PASSWORD_SECURITY` is enabled, then set a value for the token, `PASSWORD_CONTROL_EFFECTIVE_DATE`.



**Warning:** The Password Control Kit (PCK) disables, deletes or expires user accounts that don't meet the password security requirements starting from the date set for the `PASSWORD_CONTROL_EFFECTIVE_DATE` token. If a date is not set, the PCK disables, deletes or expires user accounts immediately. See [PASSWORD\\_CONTROL\\_EFFECTIVE\\_DATE](#) for more information.

- If the token `REQUIRE_RANDOM_ADMIN_PASSWORD` is already set to true, then set the token `ADMIN_EMAIL` with a valid email address.

```
ADMIN_EMAIL=root@{__APPLICATION_HOST__}
```

- If you have LDAP set up for external authentication, you must set the `"REQUIRE_USER_PASSWORD_CHANGE"` `site-options.conf` token to false.

#### f) PostgreSQL Tokens and Settings

- Enable the `POSTGRES_INTERFACE` token in the `site-options.conf` file if the value of "ifconfig -a" is not `eth0/enp0*`.
- Make sure the PostgreSQL tokens in the `site-options.conf` file are set as recommended in the following topic: [What are the right PostgreSQL settings for my site?](#) on page 943

#### g) JAVA\_OPTS

Configure the `JBOSS_JAVA_OPTS` `site-options.conf` token. See [JBOSS\\_JAVA\\_OPTS](#) on page 747.



**Note:** All JVM parameters but `-Xms1024m` and `-Xmx2048m` have been hard-coded in the TeamForge core application. You cannot manually configure any of the following default JVM parameters in the `site-options.conf` file.

- `-XX:+UseParallelGC`
- `-XX:MaxMetaspaceSize=512m`
- `-XX:ReservedCodeCacheSize=128M`
- `-server`
- `-XX:+HeapDumpOnOutOfMemoryError`
- `-XX:HeapDumpPath=/tmp -verbose:gc`
- `-XX:+PrintCodeCache`
- `-Djsse.enableSNIExtension=false`
- `-Dsun.rmi.dgc.client.gcInterval=600000`
- `-Dsun.rmi.dgc.server.gcInterval=600000`
- `-Djava.security.egd=file:/dev/urandom`
- `-Djava.awt.headless=true`.

When you change the default value of a JVM parameter such as "-XX:HeapDumpPath", the JBoss runtime parameters include both the user defined and default values for the JVM parameter. However, JBoss runs with the default value and ignores any user defined value.


h) Save the `site-options.conf` file.

14. Download the corresponding version of the Oracle client from <http://www.oracle.com/technetwork/database/features/instant-client/index-097480.html>

```
yum localinstall <path to oracle client rpm>
```

15. Deploy services.

a) Move all backup files and folders including the error folder to `/tmp`.

 **Note:** This step is required to make sure that there are no backup files and folders present in `/opt/collabnet/teamforge/var/james/var/mail/` to avoid any delay during provisioning.

```
cd /opt/collabnet/teamforge/var/james/var/mail/
mv * /tmp
```

b) Run the TeamForge provision command.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "Y" or "N" to proceed. For more information, see [TeamForge script](#) on page 785.


TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

16. Run the `/var/lib/pgsql/analyze_new_cluster.sh` script.

```
su - postgres -c "/var/lib/pgsql/analyze_new_cluster.sh"
```

17. Update the file permissions on your site's data.

```
/opt/collabnet/teamforge/runtime/scripts/fix_data_permissions.sh
```

 **Note:** This process can take a long time on sites with a lot of data.

18. If you have CVS integrations, synchronize permissions post upgrade. See, [Synchronize TeamForge source control integrations](#) on page 695.

19. Run the TeamForge post installation script. For more information, see [post-install.py](#).

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

20. Restart TeamForge.

```
setenforce 1
/opt/collabnet/teamforge/bin/teamforge restart
```

### Do this step by step on EventQ Server (server-02)

21. Stop EventQ.

```
/etc/init.d/eventq stop
/etc/init.d/collabnet-rabbitmq-server stop
/etc/init.d/collabnet-mongod stop
```

22. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.

a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

b) If the output of the `getenforce` command is either "Disabled" or "Permissive", SELinux is already disabled.

c) If not disabled, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

### 23. Upgrade the operating system packages.

```
yum upgrade
```

### 24. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to /tmp.
- b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

#### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in /tmp.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**
- In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d) If not mounted already, mount the Red Hat Enterprise Linux/CentOS installation DVD.

The DVD contains the necessary software and utilities required for installing TeamForge without internet access. In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDROM
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

- f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

### 25. Install the EventQ packages.

```
yum install teamforge-eventq CN-eventq
```



26. Copy the `site-options.conf` file from the TeamForge Application Server (server-01) to the EventQ Server's `/opt/collabnet/teamforge/etc/` directory.

27. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



**Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

**Do this step by step on TeamForge Application Server (server-01)**

28. Verify TeamForge installation.

- a) Log on to the TeamForge web application using the default Admin credentials.
  - Username: "admin"
  - Password: "admin"
- b) If your site has custom branding, verify that your branding changes still work as intended. See [Customize TeamForge](#).
- c) Let your site's users know they've been upgraded. See [Create a Side-wide Broadcast](#).

29. Fix the TeamForge Avatar display issue on RHEL/CentOS 6.8.

TeamForge Avatar image is not displayed properly post Review Board installation on RHEL/CentOS 6.8. Run the following commands to work around this issue:

```
yum erase python-imaging
yum install teamforge
service httpd restart
```

**Post Upgrade Tasks**

- [Add EventQ to existing projects](#)
- [Run ETL initial load jobs](#)

**Upgrade from TeamForge 16.3 to 17.4 - All services on the same server**

Single server setup. In this procedure, we'll assume that you are upgrading on the same server where your existing TeamForge 16.3 site is running. The following instructions are valid for both RHEL/CentOS 6.8/7.3 platforms. Specific steps, if applicable only for a particular RHEL/CentOS platform, are called out explicitly.

**Before you begin:**

[Plan your upgrade](#)

**Remember:**


- TeamForge 17.4 supports both RHEL/CentOS 6.8 and 7.3. See [TeamForge installation requirements](#) on page 441
- For the ETL service to run as expected in a distributed TeamForge installation, all servers must have the same time zone.
- While you can run both EventQ and TeamForge on the same server, CollabNet recommends such an approach only for testing purposes. It's always recommended to run EventQ on a separate server for optimal scalability. See [EventQ installation requirements](#).
- Installing or upgrading TeamForge needs root privileges. You must log on as root or use a root shell to install or upgrade TeamForge.
- In a distributed setup, stop TeamForge services on all the servers while upgrading to TeamForge 17.4.

In this setup, the following services run on the TeamForge Application Server (we call this server-01).

- TeamForge Application Server (ctfcore)
- Codesearch Server (codesearch)
- Mail Server (mail)
- Database Server (ctfcore-database and ctfcore-datamart)
- ETL Server (etl)
- Git Integration Server (gerrit and gerrit-database)
- SCM Integration Server (subversion and cvs)
- Search Server (search)
- TeamForge EventQ Server (eventq, mongodb, redis and rabbitmq)

### Do the following on the TeamForge Application Server - server-01

1. Back up all your custom event handlers and remove all the event handler JAR files before starting your TeamForge 17.4 upgrade process.
  - a) Go to **My Workspace > Admin**.
  - b) Click **System Tools** from the **Projects** menu.
  - c) Click **Customizations**.
  - d) Select the custom event handler and click **Delete**.

 **Important:** Post upgrade, you can add custom event handlers again from the backup while making sure that you don't have SOAP50 (deprecated) library used.

2. Uninstall hotfixes and add-ons, if any, installed on your site.
3. If you have Review Board installed, uninstall Review Board.

```
cd /opt/collabnet/RBInstaller-1.7.0.0.7
python ./install.py -u
```

4. Stop TeamForge.

If you are upgrading from TeamForge 16.7 or earlier releases:

```
/etc/init.d/collabnet stop all
```

If you are upgrading from TeamForge 16.10 or later releases:


```
/opt/collabnet/teamforge/bin/teamforge stop
```

5. Stop EventQ (formerly known as Orchestrate).

```
/etc/init.d/orchestrate stop
```

6. No backup is required for same hardware upgrades. However, you can create a backup as a precaution. See [Backup up and restore TeamForge and EventQ](#) on page 566 to learn more about backing up TeamForge and EventQ database and file system.

7. Move the CollabNet-Orchestrate.repo file.

 **Important:** This is required only if you are upgrading from TeamForge 16.3 (or earlier) to TeamForge 16.10 (or later).

```
mv /etc/yum.repos.d/CollabNet-Orchestrate.repo /etc/yum.repos.d/CollabNet-Orchestrate.repo_old
```

8. Run the following commands to upgrade Red Hat/CentOS to the latest version.

- a) Remove the neon-devel package.

```
yum erase neon-devel -y
```

- b) Upgrade the operating system packages.

```
yum upgrade
```

9. If the TeamForge server has SELinux enabled, run it in 'permissive' mode temporarily while installing or upgrading TeamForge.


If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is "permissive", continue with the next step. If not, run the following command to bring it to 'permissive' mode.

```
setenforce 0
```

10.  **Attention:** Do this if and only if you are upgrading TeamForge on RHEL/CentOS 6.8.

Delete the `python-crypto` package.

```
yum erase python-crypto
```

11. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.  
b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4-1.noarch.rpm
```


- c) Refresh your repository cache.

```
yum clean all
```

#### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 6.8 64 bit RPM package: **CTF-Disconnected-media-17.4.824-110.rhel6.x86\_64.rpm**
- Red Hat Enterprise Linux/CentOS 7.3 64 bit RPM package: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**


 **Note:** In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.1-1.el7.noarch.rpm**.

- b) Unpack the disconnected installation package.

```
rpm -Uvh <package-name>
```

- c) Unpack the **compat-ctf-dc-media-1.1-1.el7.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.1-1.el7.noarch.rpm
```

- d)  **Note:** If the Red Hat Enterprise Linux/CentOS installation DVD is mounted already, skip the following instructions. If not, mount the DVD.

Mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/  
mkdir cdrom  
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

- e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDRom
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

12. Install the following application packages.

a) To install the TeamForge application packages run the following command:

```
yum install teamforge CN-eventq
```



**Attention:** TeamForge installer has been optimized quite a bit. It's likely that you might come across a lot of warning messages while upgrading from TeamForge 8.2 (or earlier) to TeamForge 17.4 on the same hardware (when you run the `yum install teamforge` command). You can safely ignore such warning messages and proceed with the upgrade.

Run the following command to install the Binary application packages.

```
yum install cn-binary
```

13. Set up your site's master configuration file.

a) Set up your site options. See [Site options change log](#) on page 730 for a list of site option changes. While upgrading to a latest TeamForge release, make sure that obsolete site option tokens, if any, are removed from the `site-options.conf` file of the TeamForge version you are upgrading to.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

b) Configure the services and domain name tokens.



**Note:** `cliserver` and `reviewboard-adapter` are the new services added in TeamForge 17.4.

```
server-01:SERVICES=ctfcore ctfcore-database ctfcore-datamart etl mail
search codesearch
subversion cvs eventq redis mongodb rabbitmq cliserver gerrit gerrit-
database binary binary-database
reviewboard reviewboard-database reviewboard-adapter
```

```
server-01:PUBLIC_FQDN=my.app.domain.com
```

c) Set the `MONGODB_APP_DATABASE_NAME` token with EventQ's database name in the `site-options.conf` file.

```
MONGODB_APP_DATABASE_NAME=orchestrate
```


d) **Attention:** `SSL` is enabled by default and a self-signed certificate is auto-generated. Use the following tokens to adjust this behavior. To generate the SSL certificates, see [Generate SSL certificates](#) on page 718.

Have the custom SSL certificate and private key for custom SSL certificate in place and provide their absolute paths in these tokens. `SSL_CHAIN_FILE` (intermediate certificate) is optional.

```
SSL_CERT_FILE=
```

```
SSL_KEY_FILE=
```

```
SSL_CHAIN_FILE=
```

 **Important:** All SSL certificates including self-signed certificates are added automatically.

- e) If the token `REQUIRE_PASSWORD_SECURITY` is enabled, then set a value for the token, `PASSWORD_CONTROL_EFFECTIVE_DATE`.



**Caution:** The Password Control Kit (PCK) disables, deletes or expires user accounts that don't meet the password security requirements starting from the date set for the `PASSWORD_CONTROL_EFFECTIVE_DATE` token. If a date is not set, the PCK disables, deletes or expires user accounts immediately. See [PASSWORD\\_CONTROL\\_EFFECTIVE\\_DATE](#) for more information.

- f) TeamForge 7.1 and later support automatic password creation. See [AUTO\\_DATA](#) on page 736 for more information.
- g) If you have LDAP set up for external authentication, you must set the `REQUIRE_USER_PASSWORD_CHANGE` site options token to false.
- h) Make sure the PostgreSQL tokens in the `site-options.conf` file are set as recommended in the following topic: [What are the right PostgreSQL settings for my site?](#) on page 943
- i) Configure the `JBOSS_JAVA_OPTS` `site-options.conf` token. See [JBOSS\\_JAVA\\_OPTS](#) on page 747.



**Important:** All JVM parameters but `-Xms1024m` and `-Xmx2048m` have been hard-coded in the TeamForge core application. You cannot manually configure any of the following default JVM parameters in the `site-options.conf` file.

- `-XX:+UseParallelGC`
- `-XX:MaxMetaspaceSize=512m`
- `-XX:ReservedCodeCacheSize=128M`
- `-server`
- `-XX:+HeapDumpOnOutOfMemoryError`
- `-XX:HeapDumpPath=/tmp -verbose:gc`
- `-XX:+PrintCodeCache`
- `-Djsse.enableSNIExtension=false`
- `-Dsun.rmi.dgc.client.gcInterval=600000`
- `-Dsun.rmi.dgc.server.gcInterval=600000`
- `-Djava.security.egd=file:/dev/urandom`
- `-Djava.awt.headless=true`

When you change the default value of a JVM parameter such as `"-XX:HeapDumpPath"`, the JBoss runtime parameters include both the user defined and default values for the JVM parameter. However, JBoss runs with the default value and ignores any user defined value.

- j) Save the `site-options.conf` file.

#### 14. Deploy services.

- a) Move all backup files and folders including the error folder to `/tmp`.




**Note:** This step is required to make sure that there are no backup files and folders present in `/opt/collabnet/teamforge/var/james/var/mail/` to avoid any delay during provisioning.

```
cd /opt/collabnet/teamforge/var/james/var/mail/
```

```
mv * /tmp
```

- b) Run the TeamForge provision command.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "Y" or "N" to proceed. For more information, see [TeamForge script](#) on page 785.

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

15. Run the `/var/lib/pgsql/analyze_new_cluster.sh` script.

```
su - postgres -c "/var/lib/pgsql/analyze_new_cluster.sh"
```

16. Run the following script to upgrade the Subversion working copies.

```
/opt/collabnet/teamforge/runtime/scripts/svn-upgrade-working-copies.sh
```

17. If you have CVS integrations, synchronize permissions post upgrade. See, [Synchronize TeamForge source control integrations](#) on page 695.

18. Run the TeamForge post installation script. For more information, see [post-install.py](#).

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

19. Verify TeamForge installation.

- a) Log on to the TeamForge web application using the default Admin credentials.
- Username: "admin"
  - Password: "admin"
- b) If your site has custom branding, verify that your branding changes still work as intended. See [Customize TeamForge](#).
- c) Let your site's users know they've been upgraded. See [Create a Side-wide Broadcast](#).

20. Remove the backup files, if any, after the TeamForge site is up and running as expected.

#### **TeamForge Avatar display issue on RHEL/CentOS 6.8**

TeamForge Avatar image is not displayed properly post Review Board installation on RHEL/CentOS 6.8. Run the following commands to work around this issue:

```
yum erase python-imaging
yum install teamforge
service httpd restart
```

- [Upgrade Review Board](#)
- [Run TeamForge in SELinux enabled mode](#)
- [Add EventQ to existing projects](#)
- [Run ETL initial load jobs](#)

## **Back up and restore TeamForge and EventQ**

Here's what it takes to back up all TeamForge and EventQ data comprehensively so that it can be restored in the event of unrecoverable failures.

### **Back up and restore CollabNet TeamForge data**

CollabNet TeamForge stores data in the database and on the file system. Back up all data comprehensively so that it can be restored in the event of unrecoverable failures.

- 👉 **Note:** The items listed in this section address only the data that is either created by or a part of CollabNet TeamForge. Data that is not specific to TeamForge, such as operating system-based content, configuration files, and other third-party applications, will also require a backup and restoration routine to ensure that the entire server can be restored in the event of a catastrophic failure. Contact your application or operating system vendor for specific guidance on backup strategies for their products.

### Back up CollabNet TeamForge data

Save a copy of your TeamForge site's data to a location from where you can quickly retrieve it to your TeamForge 17.4 site.

- 👉 **Important:** The following instructions, though applicable in general to any PostgreSQL version, assume that you run a TeamForge version that runs on PostgreSQL 9.3, which you want to back up. Replace 9.3 with the PostgreSQL version you run, if required.

- If you are upgrading by installing TeamForge 17.4 on new hardware, then you'll need the backed-up site data to complete the upgrade.
- If you are upgrading your site on the same hardware, then you won't need to back up but you should create a backup anyway, as a precaution.

#### 1. Stop TeamForge.

If you are upgrading from TeamForge 16.7 or earlier releases:

```
/etc/init.d/collabnet stop all
```

If you are upgrading from TeamForge 16.10 or later releases:

```
/opt/collabnet/teamforge/bin/teamforge stop
```

#### 2. Back up your site data.

- Make a dump file of your site database.

(Do this on your database server in case you have your DB running on a separate server.)

You have to do a PostgreSQL dump if you are upgrading the PostgreSQL application as part of your TeamForge upgrade.

- 👉 **Note:** These commands are for a PostgreSQL database, which is the default. If your site uses an Oracle database, follow the [Oracle backup procedure](#) instead.

```
su - postgres
/usr/bin/pg_dumpall > /opt/collabnet/teamforge/var/pgsql/9.3/backups/
teamforge_data_backup.dmp
exit
mkdir /tmp/backup_dir
cp /opt/collabnet/teamforge/var/pgsql/9.3/backups/
teamforge_data_backup.dmp /tmp/backup_dir/
```

- 👉 **Note:** If your reporting database is running on a separate port, back up your reporting database too:

```
/usr/bin/pg_dumpall -p <reports_database_port> > /var/lib/
pgsql/9.3/backups/teamforge_reporting_data_backup.dmp
cp /opt/collabnet/teamforge/var/pgsql/9.3/backups/
teamforge_reporting_data_backup.dmp /tmp/backup_dir/
```

- Make an archive file with the following data directories:

Directory	Contents
/opt/collabnet/teamforge/var	User-created data, such as artifact attachments
/opt/collabnet/reviewboard	Review Board data
/svnroot	Subversion source code repositories

Directory	Contents
/sf-svnroot	Subversion repository for branding data
/cvsroot	CVS source code repositories (not present on all sites)

```
cp -Rpf /svnroot /sf-svnroot /cvsroot /gitroot /opt/collabnet/
teamforge/var /opt/collabnet/teamforge/var/reviewboard/data /tmp/
backup_dir
cd /tmp
tar czvf backup.tgz backup_dir
```

If Git integration is enabled, do the following:

```
cp -Rpfv /gitroot /tmp/backup_dir
cp -Rpfv /opt/collabnet/gerrit/ /tmp/backup_dir/gerrit
```

Compress your backup data.

```
cd /tmp
tar czvf backup.tgz backup_dir
```


- c) Back up your SSH keys, if any.
- d) Back up your SSL certificates and keys, if any.

### Back up CollabNet TeamForge data using `backup-data.py`

Use the `backup-data.py` utility to compress a copy of your site data to a location where you can quickly retrieve it.

This backup method requires shutting down your site briefly. If you cannot tolerate a shutdown, you might consider another backup/restore method, such as the NetApp Snapshot utility.

1. Stop TeamForge.

 **Note:** If TeamForge is running on multiple servers, stop all the servers.

```
/etc/init.d/collabnet stop
```

2. Run the backup script.

```
cd /opt/collabnet/teamforge/runtime/scripts/
./backup-data.py --destination=<directory name>
```

CollabNet TeamForge creates the directory and stores the following data in it, in compressed format:

- Subversion repositories
- CVS repositories
- The data directory ( /var)
- The CollabNet TeamForge database.

3. Start TeamForge.

```
/etc/init.d/httpd start
/etc/init.d/postgresql-9.2 start
/etc/init.d/collabnet start
```

### Restore backed-up CollabNet TeamForge data

Use the `restore-data.py` utility to bring back site data that has been backed up.

1. Stop TeamForge.

```
/opt/collabnet/teamforge/bin/teamforge stop
```

2. Run the restore script.



```
cd /opt/collabnet/teamforge/runtime/scripts/
./restore-data.py --source=<directory name>
```

Where, <directory-name> is the directory to which you backed up the data with the `backup-data.py` script.


TeamForge unpacks the backed-up Subversion and CVS repositories, the data directory, and the TeamForge database.

### 3. Start TeamForge.

```
/opt/collabnet/teamforge/bin/teamforge start
```

### Back up a PostgreSQL database

To back up a PostgreSQL database, use the `db.py` script or the `pg_dump` command depending on your setup.

 **Note:** You can use `db.py` to back up PostgreSQL data in all setups except the dedicated database server setup where `ctf` (TeamForge database) is not installed. In a dedicated database setup, you can back up the PostgreSQL database safely while it is online by using the native `pg_dump` command.

#### Using `db.py`:

- To back up the entire PostgreSQL database, run this command.

```
./db.py -a dump -f /tmp/backup/
```

- To back up only the `ctf` database, run this command.

```
./db.py -a dump -t ctf -f /tmp/backup/
```

- To back up only the `datamart`, run this command.

```
./db.py -a dump -t reporting -f /tmp/backup/
```


#### Using `pg_dump`:

- In a dedicated database setup, you can run one of these `pg_dump` commands:

In this example, the database is dumped into a GNU tar formatted file or a `.dmp` file.


- `pg_dump -Ft -b -o ctfdb > ctfdb.tar`
- `pg_dump -Fc <dbname> -f ctfdb.dmp`
- `pg_dump ctfdb > /tmp/backup_dir/teamforge_data_backup.dmp`

 **Note:** For this example, the name of the CollabNet TeamForge database is assumed to be `ctfdb`.

 **Tip:** See the PostgreSQL `pg_dump` man page for more information and examples.

### Restore a PostgreSQL database

You can restore a PostgreSQL database with the native `pg_restore` or `psql` commands, or by using the `db.py` script.

 **Note:** You can use the `db.py` script to restore PostgreSQL data in all setups except the dedicated database server setup where the TeamForge database is not installed. In a dedicated database setup, you can restore the PostgreSQL database using the `pg_restore` or `psql` command.

#### Using `db.py`:

- To restore the entire PostgreSQL database, run this command.

```
./db.py -a restore -f /tmp/backup/
```

- To restore only the `ctf` database, run this command.

```
./db.py -a restore -t ctf -f /tmp/backup/
```

- To restore only the `datamart`, run this command.

```
./db.py -a restore -t reporting -f /tmp/backup/
```

#### Using `pg_restore` or `psql`:


- In a dedicated database setup, you can use `pg_restore` or `psql`:


- a) Locate the dump file you created when backing up the PostgreSQL database.
- b) Shut down CollabNet TeamForge.
- c) Create a database and user with the names used for CollabNet TeamForge .

```
createuser -U $CTFUSER createdb -E UNICODE -U $CTFUSER ctfdb
```

- d) Restore the database using either the `pg_restore` or `psql` utility. Run one of these commands:

- `pg_restore -d ctfdb ctfdb.tar`
- `pg_restore -d <dbname> ctfdb.dmp`
- `psql ctfdb < /tmp/backup_dir/teamforge_data_backup.dmp`

 **Note:** This example assumes that the name of the TeamForge database is `ctfdb`.

 **Tip:** It may also be necessary to restore ownership of the restored tables to the `ctfuser` database user. The following will work again (assuming the database is called `ctfdb`):

```
for i in `echo "\d" | psql ctfdb | awk {'print $3'}` do echo
"ALTER TABLE $i OWNER TO $SFUSER;" | psql ctfdb done
```

### Back up CollabNet TeamForge data with NetApp

If your data is stored on a NetApp volume, you can use the NetApp `snap` utility to establish a known good data point.

 **Remember:** If your database is also on the NetApp device, the database files will be part of this snapshot.

1. Shut down the PostgreSQL database.
2. Stop the CollabNet TeamForge application server.


```
sudo /opt/collabnet/teamforge/bin/teamforge stop
/opt/collabnet/teamforge/bin/teamforge stop
```


3. Provide a name to the snapshot.  
The data up to the given time will be snapped and stored.
4. Restart CollabNet TeamForge .

```
/etc/init.d/collabnet start all
```

### Restore CollabNet TeamForge data with NetApp

Use the Network Appliance `snap` utility to "roll back" to a known good data point on a NetApp volume.

 **Note:** You cannot snap and restore a `qtree`. You must snap and restore the entire NetApp volume.

 **Caution:** When you snap and restore to a particular snapshot, all data between the present time and that snapshot is lost.

1. Shut down the PostgreSQL database.
2. Stop the CollabNet TeamForge application server.

```
sudo /opt/collabnet/teamforge/bin/teamforge stop
/opt/collabnet/teamforge/bin/teamforge stop
```

3. Using the NetApp `snap` utility, select the snap to be restored.
4. Restart the site.

```
/etc/init.d/httpd start
/etc/init.d/postgresql-9.6 start
/etc/init.d/collabnet start
```

# EventQ

---

New to EventQ? Read through these topics to know more about EventQ that provides traceability for product life cycle activities such as work items, SCM commits, continuous integration (CI) builds, and code reviews.

## Get started

---

New to EventQ? Read through these topics to get oriented.

## Overview

TeamForge EventQ is a TeamForge capability that provides traceability for product life cycle activities such as work items, SCM commits, continuous integration (CI) builds, and code reviews.

### Problem statement and Solution

Organizations currently tend to have modern mixed-vendor, heterogeneous environments with complex lifecycles including work item, commit, review, build/test, deploy, and other tasks that are monitored or managed by stand-alone tools. The tools may be vendor-supplied or open-source, and may reside on-premises or in private, public, or hybrid clouds. These tools, however, lack the ability to associate with one another and do not lend themselves to end-to-end traceability: organizations cannot easily see the connections between activities derived from disparate lifecycle tools. Yet to achieve traceability, organizations are often forced to use an all-in-one, monolithic solution that excludes popular point tools. TeamForge EventQ offers a traceability solution that preserves the advantages of your best-of-breed tools.

While TeamForge EventQ is extensible, this version provides the following packaged adapters for connecting lifecycle activities:

**Table 2: TeamForge EventQ Packaged Adapters**

Type of service	Products
SCM/Version Control	Git, Subversion
Code Review	ReviewBoard, Gerrit, Atlassian Crucible
Build/Test	Jenkins, JetBrains TeamCity
Issue Tracker	CollabNet TeamForge, Atlassian JIRA

### TeamForge EventQ Vision

In a modern and frequently fractured development environment employing tools from different vendors, TeamForge EventQ aims to provide:

- **Aggregation** — Using packaged adapters and an extensible architecture, TeamForge EventQ obtains and stores metadata about key lifecycle activities like work items, commits, builds/tests, and code reviews across vendors and platforms.
- **Traceability** — TeamForge EventQ draws associations between activities across the entire lifecycle and enables traceability through the various steps or stages of the lifecycle.

### Feature overview

TeamForge EventQ aims to aggregate lifecycle metadata across various tools and establishes networks of associations across those lifecycle activities. The use cases for TeamForge EventQ include visibility into development activities, visibility into associations between lifecycle activities, and requirements traceability for auditing purposes.

**Activity Streams:** Activity streams provide a "project chronology" showing the most recent activities at the top and the oldest at the bottom. See [Activity Stream](#) for more information.

## Extensibility

While TeamForge EventQ ships with some adapters, its extensibility ensures that users can write custom adapters to extend TeamForge EventQ to other work item trackers, version control, CI, and code review systems. New product classes may also be extended using Extensible Data Sources (XDS). See [Extending TeamForge EventQ](#) on page 884 for more information.

## Architecture

The TeamForge EventQ architecture aims for extensibility, fault tolerance, and scalability.


### Overview

End users experience TeamForge EventQ as a seamless component of TeamForge. However, under the covers EventQ is a separate application with a distinct architecture. This article intends to give an overview of that architecture to assist with installation and administration of the underlying services as well as an overview of the message queue (MQ) architecture used to integrate with and collect data from external sources.

### Architectural Components

At the highest level, there are three server side components of TeamForge EventQ:

- EventQ application server: EventQ is a Ruby on Rails application that runs on the Phusion Passenger application server with Nginx as a web server. The application server also hosts a number of daemons that read messages from RabbitMQ.
- MongoDB: EventQ persists end user-facing data to MongoDB, a NoSQL variety document-store database.
- RabbitMQ: EventQ relies on a message queue architecture to integrate with external sources. RabbitMQ is used to collect activity data from sources like work item, SCM, build, and code review systems.



### Hosting Considerations

TeamForge EventQ is designed for host separation, where each component resides on different operation systems. Any component may be run on virtualized hardware. In the absence of separate hardware or virtualization, EventQ components may be collocated. Resource requirements for components are documented in the [Install EventQ](#) section.


### Load balancing and clustering

TeamForge EventQ components may be load balanced and/or clustered. The EventQ application server may be installed on multiple hosts and load balanced (session clustering is not required). MongoDB and RabbitMQ both support clustered modes as well. However, such configurations are not documented or supported and are only offered as supplemental services from CollabNet.


### Integration with external sources

TeamForge EventQ uses a MQ architecture to collect data from external sources like work item, SCM, build/CI, and code review systems. The diagram below demonstrates a clustered network of MQ servers that share data collected

from external sources. When a relevant activity occurs, the external source uses an MQ agent to publish a specially formatted message to the MQ network. MQ servers can be clustered and strategically located to prevent loss of messages during outages.



Once a message is published by an external source, it is consumed by a daemon service running on the EventQ application server. The message is parsed and the data is persisted. MQ adapters are configured with "association keys": unique identifiers that enable EventQ to file incoming data into the appropriate source.



### Extensibility

EventQ ships with a handful of packaged adapters to common work item, SCM, build/CI, and review systems. But what if you are using different tools in your process? You can integrate your tools by building adapters that work by publishing MQ messages that are consumed by EventQ. See the Extensibility documentation for more on writing TeamForge EventQ adapters: [Extending TeamForge EventQ](#) on page 884.

### User authentication and authorization


TeamForge EventQ is authenticated and authorized by TeamForge, using TeamForge users, permissions and roles. To access TeamForge EventQ features, TeamForge users need additional permissions on a project by project basis.

TeamForge EventQ requires a valid TeamForge user session for access. If you try to access TeamForge EventQ when you are not logged in or when your session has expired, you will be directed to the TeamForge login page. After you successfully log in, you will be automatically redirected to the target TeamForge EventQ screen.

## Permissions

TeamForge EventQ permissions are set in the "Project Admin" area on a project by project basis. The following EventQ permissions may be used in creating TeamForge roles:

- EventQ READ — The minimum required permission users need to view EventQ based data such as the Activity Stream and EventQ-based associations.
- EventQ EDIT — Required to modify EventQ data sources via the **Project Admin > Tools** interface. Users with the "EventQ EDIT" permission are implicitly granted the "EventQ READ" permission.
- EventQ CREATE — Required to create new EventQ data sources via the **Project Admin > Tools** interface. Users with the "EventQ CREATE" permission are implicitly granted the "EventQ READ" and "EventQ EDIT" permissions.
- EventQ API — Required for read and write access to EventQ HTTP APIs such as the Source API and the XDS schema API. TeamForge Site Administrators do not explicitly require this permission; however, all other project members, including Project Administrators, require this permission to use the EventQ HTTP API.
- REPORTING API — Required to retrieve data from the EventQ Reporting API.

 **Note:** TeamForge EventQ does not immediately reflect permission changes. Permissions are automatically refreshed daily. To reflect permission changes more quickly, log out of TeamForge, wait 10 minutes, and then log back in again.

## Activity Sources

---

Add "sources" that populate activity streams with relevant "activity" data such as work item updates, commits, CI builds, and code reviews.

**What is a source?** - TeamForge is an open development platform with a strong notion associations. Within TeamForge, any two objects can be associated together, and when multiple associations are chained together, they create "traceability". But what about tools outside of TeamForge? EventQ provides just such a mechanism, allowing you to add relevant project tools to the TeamForge project context. This is accomplished through "sources", which represent light-weight integrations to external products. Once established, "sources" capture activity data from external tools, bringing them into the TeamForge project context and associations fabric.

For work items, sources are issue trackers; for SCM commits, sources are defined as the code repository of interest; for CI builds, sources are "jobs" or "build configurations"; for code reviews, sources are defined as the project or repository of interest.

**How do sources work?** - Whenever an "activity" occurs in a tool that has been setup as a project source, a special "adapter" sends notification of the activity to TeamForge EventQ. For example, if a specific Jenkins job is configured as a source, build information will be sent to TeamForge EventQ when that Jenkins job is executed. The resulting activity is mapped to the TeamForge project in question. A TeamForge project can have multiple sources, all sending their activity data to the TeamForge project.

**Associations** - Other than capturing external activities and displaying them in an activity stream, EventQ draws associations between these activities. In some cases, this is entirely automatic, other cases require user input to help EventQ draw accurate associations. See [Associations](#) on page 581.

**Extensibility** - TeamForge EventQ is extensible; users may author their own adapters for new source types and even register new product domains. If your organization uses tools that do not yet have adapters, you can write your own adapter that sends activity data to TeamForge EventQ. See [Extending TeamForge EventQ](#) on page 884.

## Manage build source

Add build sources to bring automated build job results into the TeamForge project and traceability context.

Build sources bring build/CI data into EventQ for archival of metadata, participation in traceability, and activity reporting.

**Create or find an existing Tool to group your build source under**

1. From the desired project context, navigate to **Project Admin > Tools**.



2. Determine whether a “Tool” exists that represents a logical container for your desired source. For instance, if you are adding a Jenkins source, look for a tool representing the instance of Jenkins in question.
  - a) If a Tool exists already, click the Tool’s title to edit the Tool.
  - b) If a Tool does not exist, click **Add Tool**.

#### Edit an existing build source

3. To edit an existing build source, find the Data Source in the list of sources associated with the tool.  
You can edit the display name or associate a different commit source for this build source. If you do not select a commit source, no associations will be drawn for the build activities coming from this build source. TeamForge displays a warning if a build source is not associated with a commit source. See [Manage commit sources](#).
4. Click **Update** to save your changes.



**Note:** While editing an existing source configuration, you can toggle a source from Active to Inactive, which will stop TeamForge from collecting data from that source; toggle back to Active to resume data collection from the source.

#### Create a new build source

5. If this is a new Tool, first select the Tool Type in the drop-down menu. For instance, if you’re adding a Jenkins Tool instance, select Jenkins.
6. Select the **Include Traceability** check box.  
A blank data source appears. If this is an existing Tool with existing sources, click the **Data Source + (Add Source)** icon to add a new blank source.
  - a) Provide a display name for the build source. This display name will help you differentiate the source in lists throughout the user interfaces. The display name can be up to 100 alphanumeric characters in length.
  - b) Select a commit source.  
Select the commit source from the list of existing commit sources (or create a new one).



#### Important:

- You must select a commit source from the list of existing commit sources to let TeamForge create associations between your build and commit activities.
- If you do not select a commit source, no associations will be drawn for the build activities coming from this build source. TeamForge displays a warning if a build source is not associated with a commit source. See [Manage commit sources](#) for more information about adding a commit source.
- For most source types, TeamForge generates a unique "source association key". The source association key uniquely identifies and helps route data from sources properly in TeamForge. You must supply this string while configuring adapters for most source types. You can copy the key by clicking on the small clipboard icon. No association key is necessary when you configure a TeamForge project code repository, though.

7. Click **Update**.

TeamForge saves the new build source and activates it.

## Manage review sources

Add review sources to bring code reviews into the TeamForge project and traceability context.


Review sources bring code review data into EventQ for archival of meta data, participation in traceability, and activity reporting.

#### Create or find an existing Tool to group your review source under

1. From the desired project context, navigate to **Project Admin > Tools**.
2. Determine whether a “Tool” exists that represents a logical container for your desired source. For instance, if you are adding a Review Board source, look for a tool representing the instance of Review Board in question.
  - a) If a Tool exists already, click the Tool’s title to edit the Tool.
  - b) If a Tool does not exist, click **Add Tool**.

#### Edit an existing review source

3. To edit an existing review source, locate the **Data Source** in question on the **Edit Tool** screen. You can edit the display name or commit association prefix for this source.
4. Click **Update** to save your changes.

 **Note:** While editing an existing source configuration, you can toggle a source from Active to Inactive, which will stop TeamForge from collecting data from that source; toggle back to Active to resume data collection from the source.

#### Create a new review source

5. If this is a new Tool, first select the Tool Type in the drop-down menu. For instance, if you're adding a Review Board Tool instance, select Review Board.
6. Select the **Include Traceability** check box.

A blank data source appears. If this is an existing Tool with existing sources, click the **Data Source + (Add Source)** icon to add a new blank source.

- a) Provide a display name for the review source. This display name will help you differentiate the source in lists throughout the user interfaces. The display name can be up to 100 alphanumeric characters in length.
- b) Enter the Commit Association Prefix.

This prefix is used to associate code commits with code reviews using the commit message. To associate a review with a code commit, you must enter the association prefix and review ID in the commit message in this format:

```
[key:review_ID]
```

Where key is any alphanumeric string. A commit association key may be up to 20 characters, including alphanumeric and underscore ("\_") characters.


7. Click **Update**.

TeamForge saves the new review source, and activates it.

## Manage commit sources

Commit sources enable your external repositories to participate in traceability.

Commit sources bring commit data into TeamForge for archival of metadata, participation in traceability, and activity reporting.

 **Note:** TeamForge Project SCM repositories are added automatically as sources, both on upgrade and when new repositories are defined in the TeamForge project context.

#### Create or find an existing Tool to group your commit source under


1. From the desired project context, navigate to **Project Admin > Tools**.
2. Determine whether a "Tool" exists that represents a logical container for your desired source. For instance, if you are adding a source to represent an external Git repository, look for a tool representing the Git server in question.
  - a) If a Tool exists already, click the Tool's title to edit the Tool.
  - b) If a Tool does not exist, click **Add Tool**.

#### Edit an existing commit source

3. To edit an existing commit source, locate the **Data Source** in question on the **Edit Tool** screen.

You can edit the display name for any commit source. Commit sources can be defined with one of two repository types:

- **Project repositories:** housed within TeamForge projects
- **External repositories:** housed outside of TeamForge projects

 **Important:** Once defined, you cannot switch the repository type. For external repositories, you can edit the repository URL. For project repositories, "Source Code view" permission is required and once selected and saved to a source, the project repository selection may not be altered. To work around this constraint, create a new source with a new repository and deactivate the old source.

4. Click **Update** to save your changes.



**Note:** While editing an existing source configuration, you can toggle a source from Active to Inactive, which will stop TeamForge from collecting data from that source; toggle back to Active to resume data collection from the source.

#### Create a new commit source

5. If this is a new Tool, first select the Tool Type in the drop-down menu. For instance, if you're adding an external Git Tool instance, select External Git.

6. Select the **Include Traceability** check box.

A blank data source appears. If this is an existing Tool with existing sources, click the **Data Source + (Add Source)** icon to add a new blank source.

a) Provide a display name for the commit source. This display name will help you differentiate the source throughout the user interfaces where this source appears. The display name can be up to 100 alphanumeric characters in length.

b) Select the repository type. Your source may be either a TeamForge project repository or an "external" repository.

- **Project repositories:** repositories housed within TeamForge projects
- **External repositories:** repositories housed outside of TeamForge projects



**Note:** Sources for Project repositories are created automatically. You need "Source Code view" permission to see available project repositories.

c) For External repositories, enter your repository URL.

This is likely the URI used to check out code or a WebDAV-enabled URL.

- **For Subversion repositories:** run the `svn info` command inside the working copy and copy/paste the value of the "URL" field.
- **For Git repositories:** run the `git remote show origin` inside the working tree and use the value of the "Fetch URL" field, without the "username@".

For example, if your Subversion repository URL is set to `https://forge.example.com/svn/repos/myproject`, TeamForge EventQ collects messages from the "myproject" repository and all repositories under "myproject" such as `https://forge.example.com/svn/repos/myproject/branches/mybranch`.

7. Click **Update**.

TeamForge saves the new commit source, and activates it.



**Note:** TeamForge EventQ displays all the defined sources for a step in the order in which they were defined.

## Manage work item sources

Add work item sources to tickets, issues, defects and other work items to bring work items into the TeamForge project and traceability context.

Work item sources bring tracker data into TeamForge for archival of metadata, participation in traceability, and activity reporting.



**Note:** TeamForge trackers are added automatically as sources, both on upgrade and when new trackers are defined in the TeamForge project context.

#### Create or find an existing Tool to group your work item source under

1. From the desired project context, navigate to **Project Admin > Tools**.

2. Determine whether a "Tool" exists that represents a logical container for your desired source. If you are adding a JIRA<sup>®</sup> source, look for a tool representing the instance of JIRA in question.

- a) If a Tool exists already, click the Tool's title to edit the Tool.
- b) If a Tool does not exist, click **Add Tool**.

#### Edit an existing work item source

3. To edit an existing work item source, locate the **Data Source** in question on the **Edit Tool** screen.  
You can edit the display name, while all other properties are system defined.
4. Click **Update** to save your changes.



**Note:** While editing an existing source configuration, you can toggle a source from Active to Inactive, which will stop TeamForge from collecting data from that source; toggle back to Active to resume data collection from the source.

#### Create a new work item source

5. If this is a new Tool, first select the Tool Type in the drop-down menu. For instance, if you're adding an external JIRA Tool instance, select the appropriate Tool Type.
6. Select the **Include Traceability** check box.  
A blank data source appears. If this is an existing Tool with existing sources, click the **Data Source + (Add Source)** icon to add a new blank source.
  - a) Provide a display name for the work item source. This display name will help you differentiate the source throughout the user interfaces where this source appears. The display name can be up to 100 alphanumeric characters in length.
7. Click **Update**.

TeamForge EventQ saves the new work item source, and activates it.

## Extensible Data Source (XDS) Overview

This section provides a description of Extensible Data Sources (XDS), a means to add arbitrary, user-defined sources.

### XDS Schemas and Sources

EventQ can be configured to accept data from a wide range of sources through the Extensible Data Source (or XDS) feature. In addition to the stock activity APIs defined for commits, builds, reviews, and work items, the XDS feature enables you to create your own activity APIs by defining an "XDS Schema" that describes the expected activity message format for a particular tool or product domain in question. See [Extending TeamForge EventQ](#) on page 884 for more details and links to API documentation.

"XDS Sources" can be added to your project through the "Custom" source type. Once added, XDS activities will be visible in the activity stream. XDS sources can also have associations just like any other source resulting in traceability between XDS and other activities, even between multiple XDS sources.

### Manage XDS sources

Add Extensible Data Sources (XDS) to bring a wide variety of tools into the TeamForge project and traceability context.

XDS sources provide a means to integrate tools that do not fall under the stock activity classes (commit, build, review, and work item). You can add XDS sources to represent a wide range of tools and product domains. See [Extending TeamForge EventQ](#) on page 884 for more on how XDS works.

#### Create or find an existing Tool to group your XDS source under

1. From the desired project context, navigate to **Project Admin > Tools**.
2. Determine whether a "Tool" exists that represents a logical container for your desired source. If you are adding a Chef™ source, look for a tool representing the Chef instance in question.
  - a) If a Tool exists already, click the Tool's title to edit the Tool.
  - b) If a Tool does not exist, click **Add Tool**.

#### Edit an existing XDS source

3. To edit an existing XDS source, locate the **Data Source** in question on the **Edit Tool** screen.  
You can edit the display name, look and feel, and tags. If you save an XDS source without defining the Associated Source, you may edit the XDS source later and add an Associated Source. Once the Associated Source has been defined it may not be altered.
4. Click **Update** to save your changes.



**Note:** While editing an existing source configuration, you can toggle a source from Active to Inactive, which will stop TeamForge from collecting data from that source; toggle back to Active to resume data collection from the source.

### Create a new XDS source

5. If this is a new Tool, first select the Tool Type in the drop-down menu. For instance, if you're adding a Nexus Tool instance, select the matching XDS Schema name from the drop-down. Alternatively, select "Other" if no applicable choice exists yet.

6. Select the **Include Traceability** check box.

A blank data source appears. If this is an existing Tool with existing sources, click the **Data Source + (Add Source)** icon to add a new blank source.

a) Provide a display name for the XDS source that is depicted on all TeamForge EventQ user interfaces. The display name can be up to 100 alphanumeric characters in length.

b) Select an icon and set the icon background color to differentiate your source visually.

c) Select an Associated Source.

- Select a source from the list of existing sources. TeamForge EventQ associates activities from the XDS source to the selected source. For instance, if you want the XDS source in question to associate to builds from job XYZ, select the source corresponding to job XYZ in the **Associated Source** drop-down list.
- If you do not select an associated source, no associations will be drawn to the activities coming from this XDS source.

d) Add tags to the Source as a means to organize and categorize XDS Sources.

Tags are particularly useful as filters for reporting via the Reporting API.



**Caution:** Tags are shared across the entire site. Exercise caution so that no superfluous tags are created which may make reporting more difficult.

7. Click **Update**.

TeamForge EventQ saves the new custom source, and activates it.

## Associations

---

TeamForge EventQ creates associations between activities such as reviews, commits, builds, and work items.

Associations are relationships between activities like work items, commits, builds, and reviews. You can see associations for a particular activity on its TeamForge EventQ detail page or create a traceability graph of associations using the **Trace associations** button from an activity details page.

Most associations in TeamForge EventQ center around commits. That is, builds, work items, and reviews associate directly to commits, while a build-to-review association are implied through a common commit.

Builds are associated to commits by the underlying CI/build system. Typically, CI/build systems obtain source code from a target repository and build based on that specific revision of source code. TeamForge EventQ therefore relies on the CI/build system to report the associated revisions, which then get mapped as associations between commits and builds in TeamForge EventQ.

Reviews associate to commits in one of two ways, depending on the process being employed and the source code review product. Please see [Review detail](#) for a detailed explanation on commit-to-review associations.

When creating code reviews, note that the code review's repository URL must match the one specified while configuring the related SCM source. Otherwise, TeamForge EventQ will not be able to register associations between commits and reviews.

Work items can be directly associated to commits by referencing the work item artifact ID in the commit message. Please see [Work item detail](#) for a detailed explanation on commit-to-work-item associations.

## TeamForge Associations

TeamForge EventQ adds build and review associations to TeamForge objects like commits and artifacts.

TeamForge EventQ draws associations between various lifecycle activities. If TeamForge commits or artifacts are a part of the associations chain, EventQ creates associations in the TeamForge object's "associations" area. This allows TeamForge users to see EventQ derived associations directly inside a TeamForge record.

# Set up Review Board

---

Review Board is a popular code review tool available with TeamForge 17.4 as a fully-integrated add-on.

 **Important:** TeamForge 17.4 supports Review Board 2.5.6.1 on RHEL/CentOS 7.3.


## Install Review Board

---

You must install the Review Board before you can make it available as an integrated application to project managers on your TeamForge site.


 **Important:**

- Install the Review Board application on the TeamForge Application Server or on a separate server of its own. Review Board database can be installed on the TeamForge PostgreSQL Database Server on sites with database running on a separate server.
- To install Review Board successfully, ensure that other repositories such as EPEL (Extra Packages for Enterprise Linux) are disabled apart from the CollabNet and Operating System repositories.
- This procedure is for those who are installing the Review Board for the first time.
- In this scenario, both TeamForge and Review Board use PostgreSQL.
- TeamForge 17.4 supports Review Board 2.5.6.1 on RHEL/CentOS 6.8 and 7.3.

 **Attention:** Installing or upgrading TeamForge needs root privileges. You must log on as root or use a root shell to install or upgrade TeamForge.

**Do this on the main TeamForge Application Server. We'll call this `my.app.box`**

1. Download the Review Board installer zip file from <http://collab.net/downloads/integrations#tab-1> and save it in the `/opt/collabnet` folder

 **Note:** Download the `RBInstaller-17.4.7.zip` file.

2. Unzip the installer zip file.

```
cd /opt/collabnet
unzip RBInstaller-17.4.7.zip
```

3. Make sure that "reviewboard", "reviewboard-database" and "reviewboard-adapter" identifiers have been added to the SERVICES token of the TeamForge Application Server's `site-options.conf` file.

```
server-01:SERVICES = ctfcore ctfcore-database mail search codesearch
cliserver etl
ctfcore-datamart subversion cvs gerrit gerrit-database binary binary-
database
reviewboard reviewboard-database reviewboard-adapter
```

Review Board database can be installed on TeamForge PostgreSQL Database Server on sites with database running on a separate server, in which case you must add the "reviewboard-database" identifier to the SERVICES token of the TeamForge PostgreSQL Database Server's `site-options.conf` file.

4. Modify these values in the `install.conf` file to suit your installation environment.


```
vi /opt/collabnet/RBInstaller-17.4.7/installer/install.conf
```

### Option

**`rb_dir=/opt/collabnet/teamforge/var/reviewboard`**

### Description

The path of the directory where the Review Board files and libraries are installed.

<b>Option</b>	<b>Description</b>
<b>rb_data_dir=/opt/collabnet/teamforge/var/reviewboard/data</b>	The path of the directory where Review Board's database file, review request files and attachments are stored.   <b>Note:</b> The default Review Board data directory has been changed from /opt/collabnet/reviewboard/data to /opt/collabnet/teamforge/var/reviewboard/data in TeamForge 17.4.
<b>domain=&lt;domain name or host name&gt;</b>	The Review Board site information. For example, cu064.cloud.maa.collab.net.
<b>rb_database_type=postgresql</b>	The Review Board database type.
<b>rb_database_host=&lt;reviewboard_db_hostname&gt;</b>	The Review Board database host name.
<b>rb_database_port=&lt;reviewboard_db_port&gt;</b>	The Review Board database port.
<b>ctf_base_url=https://myapp.collab.net</b>	The absolute URL of the TeamForge site that you want to associate to.
<b>ctf_site_var_dir=/opt/collabnet/teamforge/var</b>	The location of the rbctfevents.jar file.

5. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



**Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

6. Run the following grep commands to get the value of HTTPD\_USER, HTTPD\_GROUP and HOME\_DIR\_BASE.

```
grep HTTPD_USER= /opt/collabnet/teamforge/runtime/conf/runtime-
options.conf
grep HTTPD_GROUP= /opt/collabnet/teamforge/runtime/conf/runtime-
options.conf
grep HOME_DIR_BASE= /opt/collabnet/teamforge/runtime/conf/runtime-
options.conf
chown -R <HTTPD_USER>:<HTTPD_GROUP> <HOME_DIR_BASE>/<HTTPD_USER>
```

7. Run the install.py script available in the Review Board installer directory.

```
cd /opt/collabnet/RBInstaller-17.4.7
python ./install.py -i -r --auth-scmuser
```



**Attention:** The Review Board installer pauses abruptly on sites using trusted certificates expecting user intervention. Press Enter to proceed with the installation/upgrade.

8. Set up the initial Review Board data. Run the bootstrap-data.py script available in the Review Board installer directory.

```
cd /opt/collabnet/RBInstaller-17.4.7
python ./bootstrap-data.py
```

9. Run the grep command to get the value of httpd\_user and httpd\_group.

```
grep httpd_user /etc/reviewboard.properties
grep httpd_group /etc/reviewboard.properties
```



```
chown -R <httpd_user>:<httpd_group> /opt/collabnet/teamforge/var/
reviewboard/data
```


#### 10. Restart TeamForge.

```
setenforce 1
/opt/collabnet/teamforge/bin/teamforge restart
```

#### 11. If SCM is installed on a separate box, run the following script to authenticate a scmviewer user against a TeamForge Subversion repository for creating a new review request.

```
python ./svn-auth.py --repo-path=https://<scm_domain>/svn/repos/
<repo_dir_name>
```

You should now have a Review Board instance ready to work with TeamForge. The installer has created two configuration files: `installer/conf/rb-application.xml` and `installer/conf/rb-deploy.xml`. See [Integrate Review Board with your TeamForge site](#) on page 587 for more information.

 **Note:** Run the following scripts available in the Review Board installer directory.

```
cd /opt/collabnet/RBInstaller-17.4.7
```


- To recreate the runtime, run the `python ./install.py -r` command.
- To update the revised password for a scmviewer account, run the `python ./install.py --auth-scmuser` command.
- To change the scmviewer password, see [SCM\\_USER\\_ENCRYPTED\\_PASSWORD](#) on page 770.
- To uninstall Review Board, run the `python ./install.py -u` command.

## Upgrade Review Board

---


Use these instructions to upgrade Review Board to a latest build.

- TeamForge 17.4 supports Review Board 2.5.6.1 on RHEL/CentOS 6.8 and 7.3.
- This procedure is for those who have Review Board already and are upgrading Review Board to a latest build on RHEL/CentOS 6.8 or 7.3.
- You may choose to upgrade Review Board on the same server or on a new server.
- In this scenario, both TeamForge and Review Board use PostgreSQL.

 **Attention:** Installing or upgrading TeamForge needs root privileges. You must log on as root or use a root shell to install or upgrade TeamForge.

#### Do this step by step on the TeamForge Application Server

1. Download the Review Board installer zip file from <http://collab.net/downloads/integrations#tab-1> and save it in the `/opt/collabnet` folder
2. The default Review Board data directory has been changed from `/opt/collabnet/reviewboard/data` to `/opt/collabnet/teamforge/var/reviewboard/data` in TeamForge 17.4. If you are upgrading from TeamForge 17.1 or earlier to TeamForge 17.4 or later, regardless of whether you upgrade Review Board on the same or new hardware, you must back up your Review Board data directory from `/opt/collabnet/reviewboard/data` and restore it to `/opt/collabnet/teamforge/var/reviewboard/data`.

 **Tip:** The Review Board database is backed up already when you have upgraded TeamForge. So, it is not necessary to take a back up of the Review Board database again.

- a) Back up the Review Board data directory.

```
cd /opt/collabnet
tar -zcvf /tmp/reviewboard_data.tgz reviewboard
```

- b) Copy the `/tmp/reviewboard_data.tgz` file to the `/tmp` directory of the new server if you are upgrading Review Board on a new hardware.

3. Make sure that "reviewboard", "reviewboard-database" and "reviewboard-adapter" identifiers have been added to the host:SERVICES token of the TeamForge Application Server's `site-options.conf` file.


```
server-01:SERVICES=ctfcore ctfcore-database mail search codesearch etl
ctfcore-datamart subversion cvs
gerrit gerrit-database binary binary-database reviewboard reviewboard-
database reviewboard-adapter cliserver
```

4. Unzip the `RBInstaller-17.4.7.zip` file.

```
cd /opt/collabnet
unzip RBInstaller-17.4.7.zip
```

5. Modify these values in the `install.conf` file to suit your installation environment.

```
vi /opt/collabnet/RBInstaller-17.4.7/installer/install.conf
```

Option	Description
<b>rb_dir=/opt/collabnet/teamforge/var/reviewboard</b>	The path of the directory where the Review Board files and libraries are installed.
<b>rb_data_dir=/opt/collabnet/teamforge/var/reviewboard/data</b>	The path of the directory where Review Board's database file, review request files and attachments are stored.   <b>Note:</b> The default Review Board data directory has been changed from <code>/opt/collabnet/reviewboard/data</code> to <code>/opt/collabnet/teamforge/var/reviewboard/data</code> in TeamForge 17.4.
<b>domain=&lt;domain name or host name&gt;</b>	The Review Board site information. For example, <code>cu064.cloud.maa.collab.net</code> .
<b>rb_database_type=postgresql</b>	The Review Board database type.
<b>rb_database_host=&lt;reviewboard_db_hostname&gt;</b>	The Review Board database host name.
<b>rb_database_port=&lt;reviewboard_db_port&gt;</b>	The Review Board database port.
<b>ctf_base_url=https://myapp.collab.net</b>	The absolute URL of the TeamForge site that you want to associate to.
<b>ctf_site_var_dir=/opt/collabnet/teamforge/var</b>	The location of the <code>rbctfevents.jar</code> file.

6. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



**Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).


TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.


7. Run the following `grep` commands to get the value of `HTTPD_USER`, `HTTPD_GROUP` and `HOME_DIR_BASE`.

```
grep HTTPD_USER= /opt/collabnet/teamforge/runtime/conf/runtime-
options.conf
grep HTTPD_GROUP= /opt/collabnet/teamforge/runtime/conf/runtime-
options.conf
```

```
grep HOME_DIR_BASE= /opt/collabnet/teamforge/runtime/conf/runtime-
options.conf
sudo chown -R <HTTPD_USER>:<HTTPD_GROUP> <HOME_DIR_BASE>/<HTTPD_USER>
```

#### 8. Restore the Review Board data.

 **Note:** The default Review Board data directory has been changed from `/opt/collabnet/reviewboard/data` to `/opt/collabnet/teamforge/var/reviewboard/data` in TeamForge 17.4. If you are upgrading from TeamForge 17.1 or earlier to TeamForge 17.4 or later, regardless of whether you upgrade Review Board on the same or new hardware, you must back up your Review Board data directory from `/opt/collabnet/reviewboard/data` and restore it to `/opt/collabnet/teamforge/var/reviewboard/data`.

 **Note:** If you are upgrading on a new hardware, ensure that you have already copied the backup of the Review Board data directory to the `/tmp` directory of the new server.

##### a) Restore the Review Board data directory.

```
cd /opt/collabnet/teamforge/var/
tar -zxvf /tmp/reviewboard_data.tgz
```

#### 9. Run the `install.py` script available in the Review Board installer directory.

```
cd /opt/collabnet/RBInstaller-17.4.7
python ./install.py -i -r --auth-scmuser
```

#### 10. Run the `migrate.py` script available in the Review Board installer directory.

```
cd /opt/collabnet/RBInstaller-17.4.7
python ./migrate.py
```

#### 11. Restart TeamForge.

```
setenforce 1
/opt/collabnet/teamforge/bin/teamforge restart
```

#### 12. If SCM is installed on a separate box, run the following script to authenticate a scmviewer user against a TeamForge Subversion repository for creating a new review request.

```
python ./svn-auth.py --repo-path=https://<scm_domain>/svn/repos/
<repo_dir_name>
```

#### 13. Select **My Workspace > Admin** and select **Projects > System Tools > Customizations**.


#### 14. Select `rbctfevents.jar` and click **Enable**.

## Integrate Review Board with your TeamForge site

---

When the Review Board application has been installed on your site, you can make it available for projects on your TeamForge site.

Review Board must be installed and configured before you can integrate it with your TeamForge site. See [Install Review Board](#) on page 583.

 **Important:** You must not create, edit or delete new user accounts while installing the Review Board application and integrating it with TeamForge.

When you have installed Review Board, projects on your site can add Review Board to their set of collaboration tools.


1. Log into TeamForge as an administrator.
2. Go to **My Workspace > Admin**.
3. Click **INTEGRATED APPS** from the **Projects** menu.
4. Click **Create**.

5. Use the **Browse** window to select the application and deployment configuration files that enable the Review Board application to work as a part of TeamForge:

- `rb-application.xml` (Application configuration file): Contains the text strings for the Review Board user interface. This file can be found in `/opt/collabnet/RBInstaller-17.4.7/installer/conf/`.
- `rb-deploy.xml` (Deployment configuration file): Contains the data that Review Board needs to interact with TeamForge. This file can be found in `/opt/collabnet/RBInstaller-17.4.7/installer/conf/`.


Click **Next**.

6. On the **Preview** screen, review the parameters you set in the configuration files.

 **Note:** You may have to revise one or more values to ensure they are valid.

7. Click **Save**.

The Review Board application is now available for all projects on your site. Project administrators can refer to the topic, [Add Review Board to a TeamForge project](#) on page 588, for instructions on adding Review Board to their own project toolbars.

 **Note:** You may need to adjust your site's look and feel to support your integrated application. See [the site admin help](#) for details.

## Add Review Board to a TeamForge project

---

When the site administrator has made the Review Board application available, project administrators can add it as one of their project tools.


When you add Review Board to your project, it works just like the other TeamForge tools, authorization, authentication, go-urls, association, linkification and source code management support.

For this example, we'll call your project "testproject" and we'll assume you have Project Admin rights in that project.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Tools** to see the list of integrated applications available in the site.
3. Click **Add Tool**.
4. Select **Review Board** from the list of tools displayed.


Set the values that make sense for your TeamForge installation, then click **Save**.


Option	Description
<b>Prefix</b>	Specify a unique alphanumeric string that will identify this tool throughout the site. For example, suppose you set your prefix to <code>ZZ</code> . Entering <code>ZZ_1</code> in the <b>Jump to ID</b> redirect you to review request id 1 of the "testproject" project. Another may also add Review Board as a tool, with a prefix of <code>YY</code> . Jumping to <code>YY_1</code> will redirect to review request id 1 of that project.

 **Note:** The tool prefix cannot be changed after you have set it.

### Exclude Repositories

Use a comma to separate the directory names of the repositories that you want to exclude from syncing with Review Board. To include all repositories, enter "None".

 **Important:** You must sync repositories with Review Board (select **Synchronize**

Option	Description
<b>Synchronize Repositories</b>	<p data-bbox="941 189 1347 294"><b>Repositories</b> check box) to have the repositories listed in this <b>Exclude Repositories</b> text box excluded.</p> <p data-bbox="860 304 1445 378">Synchronize TeamForge Subversion repositories with Review Board.</p> <p data-bbox="860 388 1461 711"> <b>Important:</b> Select this check box every time you want to sync repositories to Review Board. When you add Review Board to your project with this check box selected, the SVN repositories are synced once and the check box is cleared. For any new repositories you create, you must edit the Review Board (<b>Project Admin &gt; Project Toolbar &gt; Edit Integrated Application</b>) and sync new repositories with Review Board by selecting this check box.</p>

If a **Review Board** button appears along with the prefix as a tooltip when you mouse over the button, your job is done.

## TeamForge - Review Board integration: FAQs

---

Frequently asked questions about TeamForge - Review Board integration.

### TeamForge - Review Board integration: Install FAQs

Questions about TeamForge - Review Board integration.

**What are the software requirements for installing Review Board as an integrated application in TeamForge 17.4?**

Review Board can run on:

- RHEL/CentOS 7.3

In addition, Review Board needs PostgreSQL 9.6.1. See [TeamForge installation requirements](#) on page 441 topic for more information.

**Can I install TeamForge and Review Board on separate servers?**

No. Review Board must be installed on the same server where TeamForge runs.

### TeamForge - Review Board integration: General usage FAQ

General usage questions about the TeamForge - Review Board integration.

**Which version of Review Board does TeamForge 17.4 support?**

TeamForge 17.4 supports Review Board 2.5.6.1.

**Which repositories does Review Board support?**

Review Board supports only Subversion repository in TeamForge 17.4.

**How do I manage users in Review Board?**

You can manage Review Board users from TeamForge. Whenever you create or edit users in TeamForge, they are synchronized automatically in Review Board.

**Can I specify 'RB' as a prefix in my project?**

No. You cannot specify 'RB' as a prefix in your project. The prefix for Review Board must be unique for every project.


<b>Is it possible to grant TeamForge specific-permissions as part of the system generated Review Board administrator?</b>	No. It is not possible to grant TeamForge specific-permissions as part of the system generated Review Board administrator (integrated application specific role).
<b>Can I use the Review Board 'Search' feature after integrating Review Board with TeamForge?</b>	No. TeamForge does not support the 'Search' feature of Review Board.
<b>What are the additional features available in Review Board after you integrate it with TeamForge?</b>	Review Board uses some of the TeamForge features like object IDs, links, GO URLs, and SVN integration and associations. For more information, see <a href="#">How does an integrated application interact with other TeamForge tools?</a> on page 1025
<b>What are the other TeamForge features which Review Board does not support after you integrate Review Board with TeamForge?</b>	Global search, page component, recent history and project template features of TeamForge are not supported in Review Board.
<b>Where can I find the documentation for Review Board?</b>	You can find the documentation for Review Board <a href="#">here</a> .
<b>How do I rectify incorrect pending review count in Review Board ?</b>	You can reset the counters using the following command and restart the Apache server.  <pre>/opt/collabnet/reviewboard/bin/ rb-site manage /opt/collabnet/ reviewboard/ReviewBoard-2.5.6.1 fixreviewcounts</pre>

## How do I control user access for Review Board?

You can control user access for Review Board based on user roles.

The following roles are defined in the XML application configuration file in TeamForge:

<b>Submit / Edit / View</b>	Using this permission, you can submit new review request, modify your existing review requests, and view the existing requests in Review Board.
<b>Change Status as another user</b>	Using this permission, you can change the review request status of another user.
<b>Submit as another user</b>	Using this permission, you can submit a review request as another user through the post-review command line tool.
<b>Edit as another user</b>	Using this permission, you can modify a review request of another user through the user interface or the post-review command line tool.
<b>Admin staff</b>	When you have this permission along with the Submit/Edit/View permission through a site-wide role, you can access Review Board administration user interfaces.

 **Note:** Site administrators have all permissions on all projects on the site, by default. Site administrators grant these permissions to other users based on their roles.

## Set up Git

---

TeamForge 17.4 supports Git integration based on Gerrit 2.13.x.

### TeamForge—Git integration: an overview

---

TeamForge 17.4 supports an integration with the Git distributed version control tool powered by Gerrit.

Although Git is the world's leading distributed version control system, the enterprise has been slow and tentative in its adoption. Concerned with security breaches, compliance violations and lack of governance, many organizations have chosen to take a "wait and see" approach. With TeamForge, Git is ready for the enterprise. TeamForge lets you realize all the benefits of Git while ensuring the security, governance and manageability your business demands. With TeamForge, you can even manage Git and Subversion together, within each individual project.

Gerrit is an open source code review system designed to work with Git. Gerrit supports various access control mechanisms. The TeamForge Git integration uses Gerrit as a vehicle to bring TeamForge project roles and permissions into Git.



#### Highlights

- Management of Central *Git repository permissions through TeamForge* at the branch, repo, project, *project hierarchy*, Git server and *site levels*.
- Full *protection against source code history rewrites* and accidental branch deletions.
- Powerful rules that decide when a commit can be merged into production (*quality gate wizard for Gerrit*).
- Easy *Code Review policies* to manage complex Gerrit access rights.
- Universal code search across repositories and servers.
- Highly optimized for servers under heavy load (500k and above requests per day per server; 24 and above parallel fetch requests per minute).
- Thoroughly tested and certified for security standards such as SSAE 16 and *PCI-DSS*.
- Multiple authentication methods (SSH and http using public keys and TeamForge credentials).
- Automated TeamForge artifact association with Git commits if the artifacts are mentioned in the commit message.
- Embedded source code browsing.
- *Seamless navigation including SSO between TeamForge workspace, Gerrit's review UI, and git repository browser*.

- Standard Gerrit APIs to connect with *Continuous Integration systems such as Jenkins*.
- *Detailed push notification emails*.
- Native integration into popular IDEs such as Eclipse, NetBeans, Visual Studio, IntelliJ and standalone Git clients such as *GitEye*.
- *Git-Pull-Request-based feature branch development workflow*, in addition to the existing individual commit-centric methodology from Gerrit.

## Install Git on a separate server

---

In this setup, we install the Git integration services on a separate server. It is assumed that TeamForge is installed on a separate server.



**Attention:** Installing or upgrading TeamForge needs root privileges. You must log on as root or use a root shell to install or upgrade TeamForge.

In this setup, TeamForge and Git (`gerrit` and `gerrit-database`) run on two separate servers : **server-01** and **server-02** respectively. It is assumed that TeamForge has been installed already on **server-01**.

### Do this on the Git server (server-02)

1. Install Red Hat Enterprise Linux/CentOS 7.3 and log on as root.

The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.

See [the Red Hat installation guide](#) for help.

2. Check your basic networking setup. See [Set up Networking](#) for more information.
3. If the Git server has SELinux enabled, run it in 'Permissive' mode temporarily while installing or upgrading TeamForge.



**Important:** If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "Permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.

- a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is "Permissive", continue with the next step. If not, run the following command to bring it to 'Permissive' mode.

```
setenforce 0
```

See [Set up SELinux](#) on page 716 to have TeamForge run in SELinux enforcing mode after completing the installation or upgrade.

4. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.
- b) Install the repository package.

```
yum install -y /tmp/collabnet-teamforge-repo-17.4.27-1.noarch.rpm
```

- c) Refresh your repository cache.


```
yum clean all
```

#### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.

- Red Hat Enterprise Linux/CentOS 7.3 64 bit RPM package: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**




 **Note:** In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.0-1.el7.centos.noarch.rpm**.

b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

c) Unpack the **compat-ctf-dc-media-1.0-1.el7.centos.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.0-1.el7.centos.noarch.rpm
```

d)  **Note:** If the Red Hat Enterprise Linux/CentOS installation DVD is mounted already, skip the following instructions. If not, mount the DVD.

Mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDROM
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

f) Verify your yum configuration files.

```
yum list httpd
yum list apr
```

5. Install the Git packages.

```
yum install teamforge-git
```

6. Copy the `site-options.conf` file from the TeamForge Application Server (**server-01**) to the `/opt/collabnet/teamforge/etc/` directory of the Git integration server (**server-02**) and configure `site-options.conf` file.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

a) Configure the services and domain name tokens.

```
server-01:SERVICES=ctfcore ctfcore-database etl ctfcore-datamart mail
search codesearch
eventq redis mongodb rabbitmq subversion cvs binary binary-database
reviewboard reviewboard-database
```

```
server-02:SERVICES=gerrit gerrit-database
```

```
server-01:PUBLIC_FQDN=my.app.domain.com
```

```
server-02:PUBLIC_FQDN=my.git.domain.com
```

👉 **Note:** The above example lists all possible services that could run on the TeamForge Application Server. You may have to adjust the list of services for `server-01:SERVICES` token to suit your site's distributed server setup, if required.

b) Save the `site-options.conf` file.

## 7. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

👉 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

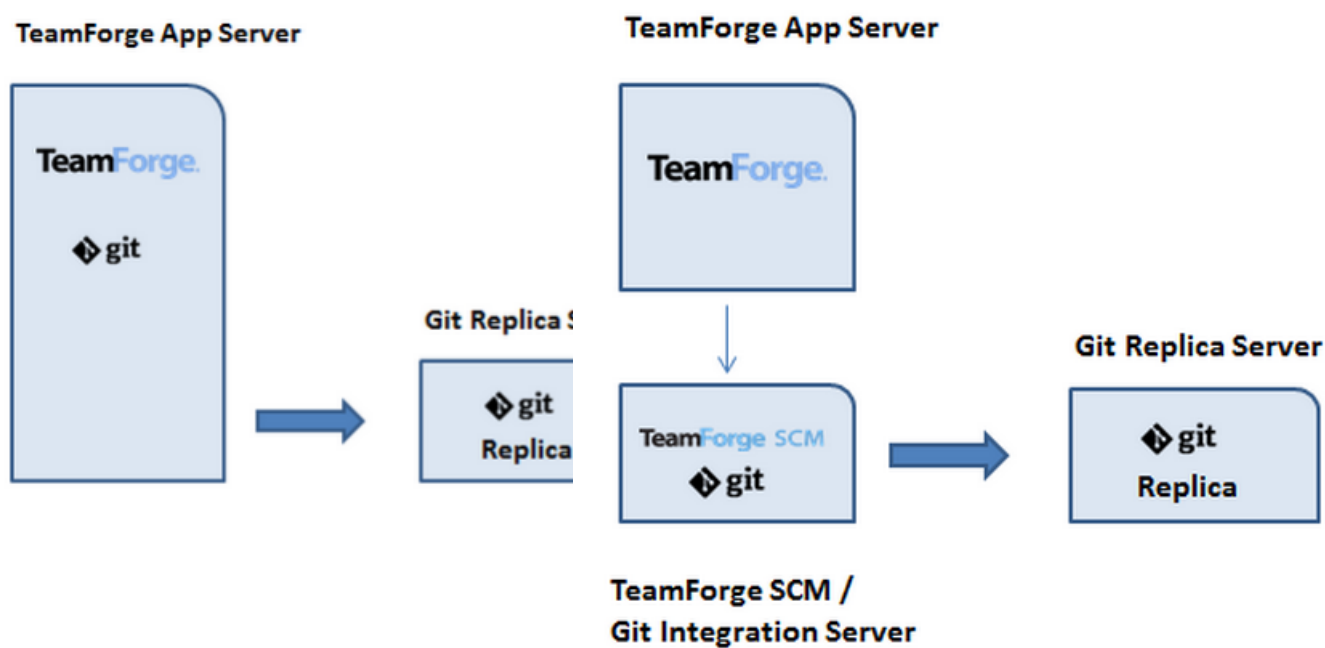
TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

## Set up a Git replica server

On sites distributed across multiple geographic locations, Git replica servers are local and remote mirror servers that can provide up to date copies of the central repositories, which if set up, can address load balancing and fetch performance issues. You can set up one or more Git replica servers (also referred to as slave or mirror servers) with TeamForge 8.1 and later.

### Points to note before you begin

- Every Git replica server has one and only one 'master' Git server.
- It's not possible to set up both Git master and slave on the same server. However, you can have multiple master and slave servers in your TeamForge environment.
- You may have your master Git integration server installed on the TeamForge application server's host or on a separate host dedicated to Git/SCM integration.



- Git replication servers can be set up with TeamForge 8.1 or later only.
- You must upgrade your Git integration server to v8.4.6 or later.
- Have the master Git integration server's `externalSystemId` handy.

- Open the `/opt/collabnet/gerrit/etc/gerrit.config` file on the master Git integration server and note down the `externalSystemId` from the `[teamforge]` section.
- Alternatively, Log on to the TeamForge application server as an Admin user, click **Admin > Integrations > SCM Integrations**, select the master Git integration server, click **Edit** and look for a token such as `exsy####`, for example `exsy1002`, in the browser URL. This is the external system ID of your Git integration server.
- Open the TeamForge application server's `site-options.conf` file and keep the values of the following tokens handy.
  - `SCM_DEFAULT_SHARED_SECRET`
  - Note down the values of the following tokens only if obfuscation is enabled on your TeamForge application server (`OBFUSCATION_ENABLED=true`):

```
OBFUSCATION_ENABLED
OBFUSCATION_KEY
OBFUSCATION_PREFIX
AUTO_DATA
```

1. Install Red Hat Enterprise Linux/CentOS 7.3 and log on to the server as root.



### Important:

- Don't customize your installation. Select only the default packages list.
  - The host must be registered with the Red Hat Network if you are using Red Hat Enterprise Linux.
  - See the [Red Hat Enterprise Linux/CentOS 7 installation guide](#) for help.
2. Check your basic networking setup. See [Set up networking](#) for more information.
  3. If the Git replica server has SELinux enabled, run it in 'Permissive' mode temporarily while installing or upgrading TeamForge. TeamForge create runtime fails otherwise.
    - a) Verify if SELinux is running in enforcing mode.

```
getenforce
```

- b) If the output of the `getenforce` command is "Permissive", continue with the next step.
- c) If not, run the following command to bring it to 'Permissive' mode.

```
setenforce 0
```

4. Upgrade the operating system packages.

```
yum upgrade
```

5. Reboot the server.

```
reboot
```

6. Configure your TeamForge installation repository.

#### TeamForge installation repository configuration for sites with internet access

- a) Contact the [CollabNet Support](#) and download the TeamForge 17.4 installation repository package to `/tmp`.
- b) Install the repository package.


```
yum install -y /tmp/collabnet-teamforge-repo-17.4.27-1.noarch.rpm
```

- c) Refresh your repository cache.

```
yum clean all
```

#### TeamForge installation repository configuration for sites without internet access

- a) Contact the [CollabNet Support](#) to get the auxiliary installer package for TeamForge 17.4 disconnected installation and save it in `/tmp`.
  - Red Hat Enterprise Linux/CentOS 7.3 64 bit RPM package: **CTF-Disconnected-media-17.4.824-110.rhel7.x86\_64.rpm**


 **Note:** In addition to the above CentOS 7.3 64 bit RPM package, you must get the following CentOS 7.3 compatibility RPM, which is required for TeamForge 17.4 disconnected media installation on CentOS 7.3 profile: **compat-ctf-dc-media-1.0-1.el7.centos.noarch.rpm**.

b) Unpack the disconnected installation package.

```
rpm -ivh <package-name>
```

c) Unpack the **compat-ctf-dc-media-1.0-1.el7.centos.noarch.rpm** package if you are installing TeamForge 17.4 on CentOS 7.3.

```
rpm -ivh compat-ctf-dc-media-1.0-1.el7.centos.noarch.rpm
```

d)  **Note:** If the Red Hat Enterprise Linux/CentOS installation DVD is mounted already, skip the following instructions. If not, mount the DVD.

Mount the Red Hat Enterprise Linux/CentOS installation DVD. The DVD contains the necessary software and utilities required for installing TeamForge without internet access.

In the following commands, replace "cdrom" with the identifier for your server's CD/DVD drive, if necessary.

```
cd /media/
mkdir cdrom
mount /dev/cdrom ./cdrom/
```

If there are any spaces in the automount, unmount it first and mount it as a filepath, with no spaces.

e) Create a yum configuration file that points to the Red Hat Enterprise Linux/CentOS installation DVD.

```
vi /etc/yum.repos.d/cdrom.repo
```

Here's a sample yum configuration file.

```
[RHEL-CDROM]
name=RHEL CDROM
baseurl=file:///media/cdrom/Server/
gpgfile=file:///media/cdrom/RPM-GPG-KEY-redhat-release
enabled=1
gpgcheck=0
```

f) Verify your yum configuration files.


```
yum list httpd
yum list apr
```

7. Install the Git packages.

```
yum install teamforge-git
```

8. Set up the `site-options.conf` token for Git.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

 **Note:** In this topic:

- **my.app.domain.com** is an example of the Fully Qualified Domain Name (FQDN) of your TeamForge application server.
- **my.git.domain.com** is an example of the Fully Qualified Domain Name (FQDN) of your Git integration server.
- **my.gitreplica.domain.com** is an example of the Fully Qualified Domain Name (FQDN) of your Git replica server.

a) Modify the HOST tokens.

```
my.gitreplica.domain.com:SERVICES=gerrit gerrit-database
my.app.domain.com:SERVICES=ctfcore ctfcore-database ctfcore-datamart etl
search subversion cvs binary binary-database
```

- b) Turn on the SSL for your site by editing the relevant variables in the `site-options.conf` file. To generate the SSL certificates, see [Generate SSL certificates](#).

```
SSL=on
SSL_CERT_FILE=
SSL_KEY_FILE=
SSL_CHAIN_FILE=
```

 **Note:**

- The `SSL_CHAIN_FILE` is optional.
- If you use certificates that are generated in-house, self-signed, or signed by a non-established Certificate Authority, they must be registered with each client system that will connect to the TeamForge server.
- For the setup discussed in this topic, add the certificate of **my.app.domain.com** to the JVM of **my.git.domain.com** and **my.gitreplica.domain.com**. In addition, add the certificate of **my.gitreplica.domain.com** to the JVM of **my.git.domain.com**. [Click here](#) for more information.

- c) Set the gerrit replication server mode.

```
GERRIT_REPLICATION_MODE=slave
```

- d) Set the external system ID of the master Git integration server. See '[Points to note before you begin](#)'.

```
GERRIT_REPLICATION_MASTER_EXTERNAL_SYSTEM_ID=exsy####
```

- e) Set the obfuscation related tokens as discussed earlier in [Points to note before you begin](#).  
f) Save the `site-options.conf` file.

## 9. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

## 10. Run the TeamForge post installation script. For more information, see [post-install.py](#).

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

Now, the gerrit service is running in replica mode. You can now find the newly created Git replica server listed on TeamForge server by accessing the following url. `http://<TF_HOST>/sf/sfmain/do/listSystems`

Once you have set up one or more Git replica servers, you can replicate repositories. For more information, see [Replicate repositories with Git replica servers](#) on page 597.

## Replicate repositories with Git replica servers

---

Once you have set up one or more Git replica servers, you can replicate repositories.

For more information on how to set up Git replica servers, see [Set up a Git replica server](#) on page 594. This topic discusses how you can start replicating your Git repositories. For more detailed steps like how to create a TeamForge project and repositories, refer to TeamForge documentation.

1. After creating one or more Git replica servers, create a TeamForge project.

My Workspace

MY PROJECTS ALL PROJECTS PROJECT GROUPS PROJECT TEMPLATES

Page 1 of 1 (2 Items)

PROJECT NAME	PARENT PROJECT	PROJECT DESCRIPTION	PROJECT ADMIN	PROJECT MEMBERS	CREATION DATE
Global Development		Project for a global development project		0	06/24/2015

Delete Create Request Membership

## 2. Create a repository.

Global Development / Source Code / Create Repository

### Create Repository

SERVER: Git HQ

DIRECTORY NAME:\* shinyapp

REPOSITORY NAME:\* ShinyApp

DESCRIPTION: Our new flagship product which is developed across the globe and is connected to many build systems (CI).

REPOSITORY CATEGORY:  Default - no review

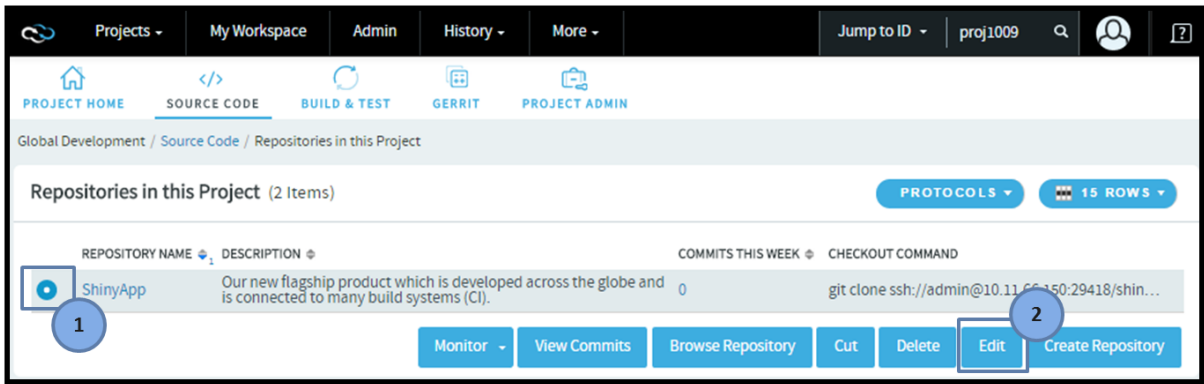
PROTECT HISTORY:

ASSOCIATION:  Required on Commit

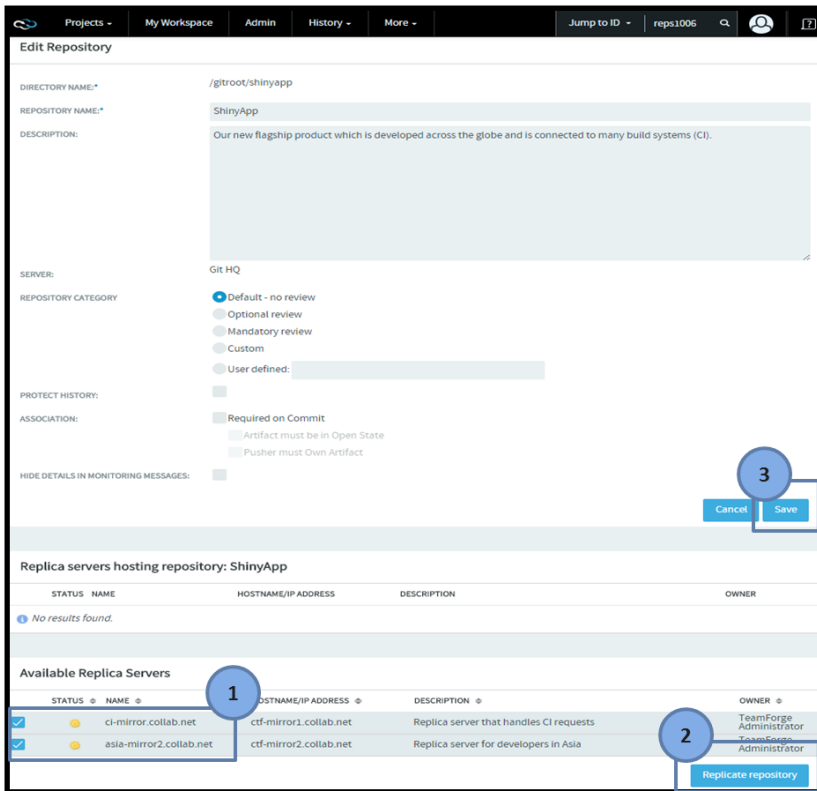
HIDE DETAILS IN MONITORING MESSAGES:

Cancel Save

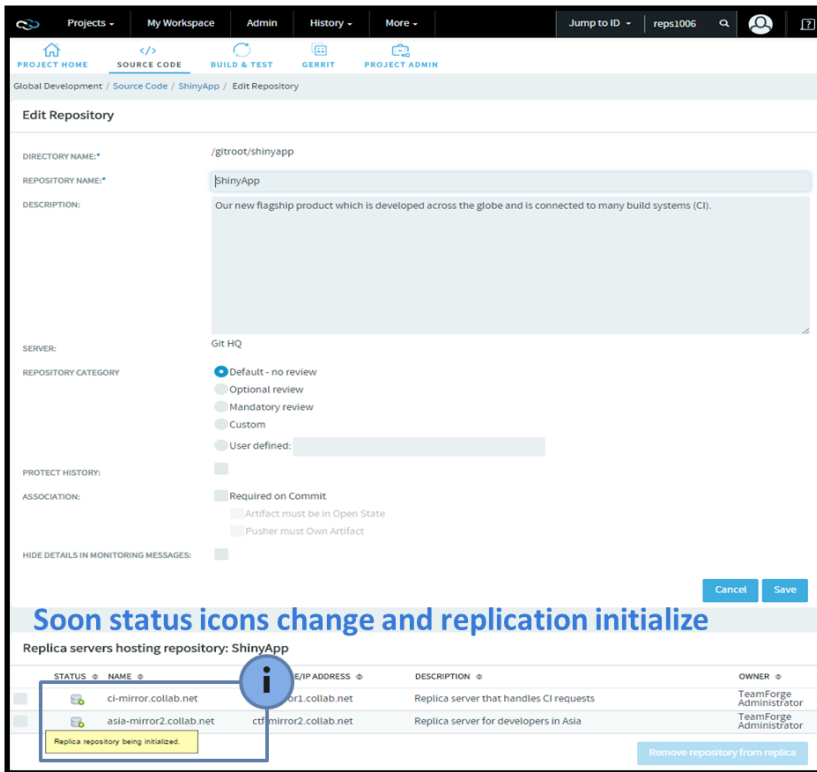
## 3. Edit the newly created repository to set up replication.




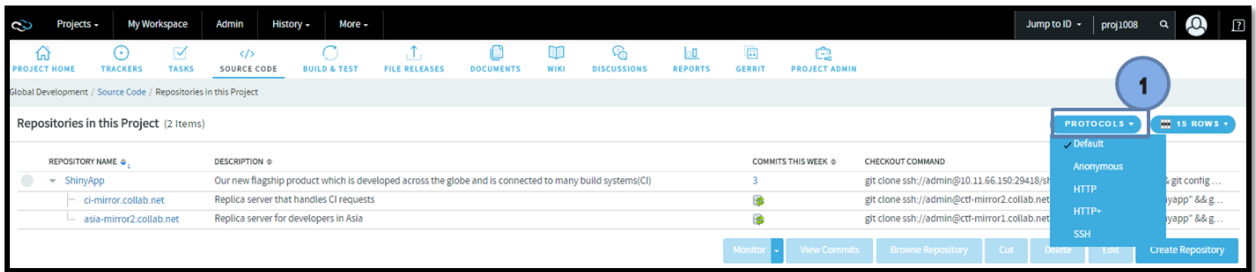
4. Select the target replica server(s).



5. Verify the replication status.

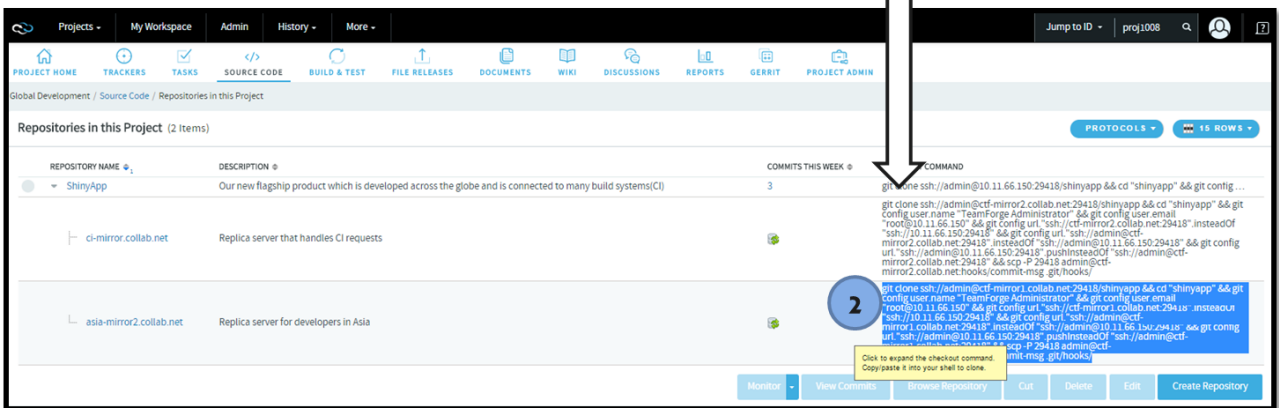
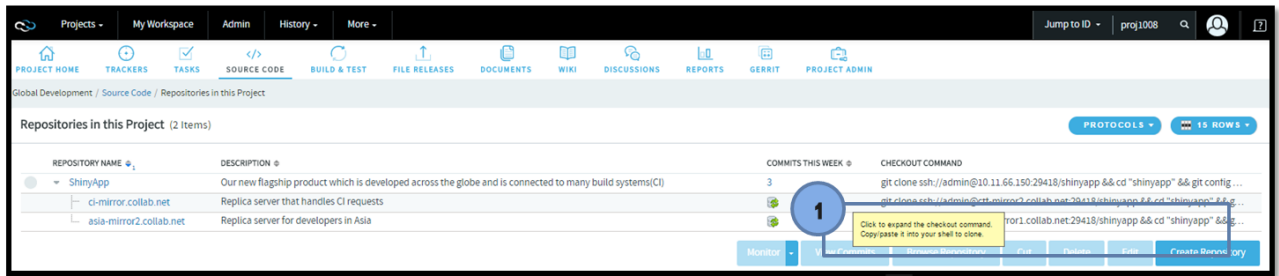


6. Use the clone protocol of your choice.

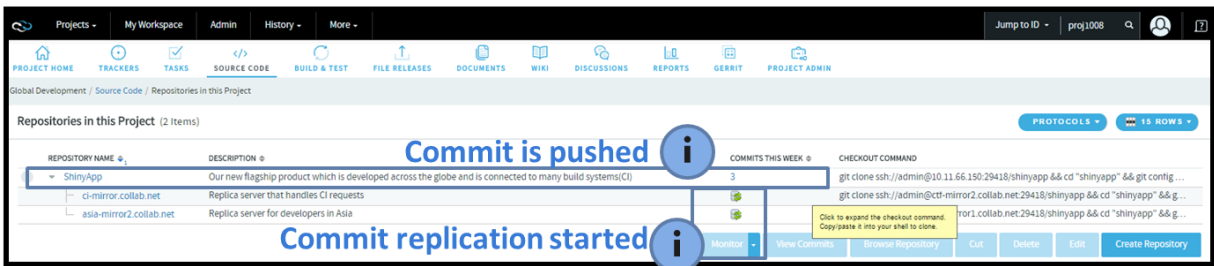


7. Clone the repository from a replica server.

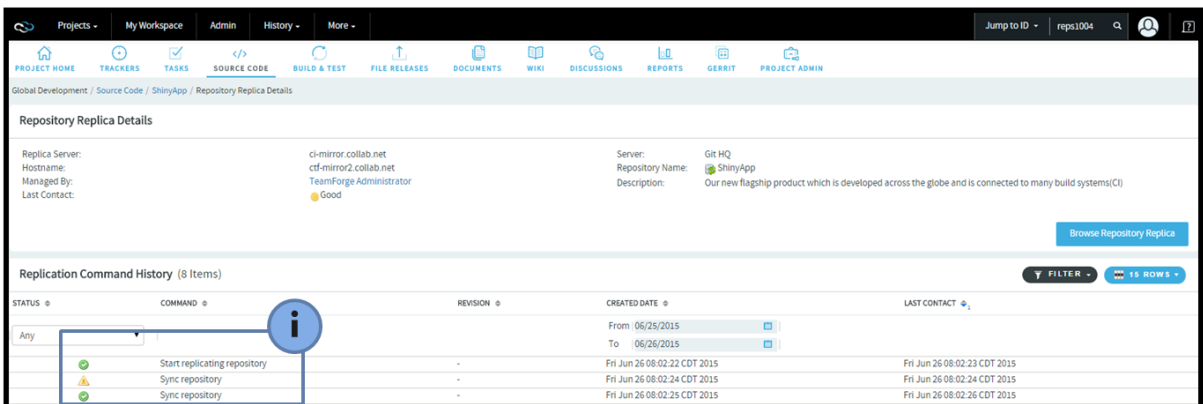




8. Push a commit and verify if it's replicated on the replica servers.



9. Verify the status of synch operations.



## Set up Large File Storage (LFS) for TeamForge repository

LFS is supported by TeamForge-Git/Gerrit integration 16.7.10-2.13.2 and later. LFS is controlled by two levels of configuration in TeamForge. First, integration level LFS configuration that provides default values for a given Gerrit

instance. Second, repository level LFS configuration, which by default derives system level configuration that can be further adjusted.

In practice, it is assumed that the Gerrit integration server is LFS ready by default and one (Project Owner/Site Admin) decides on enabling LFS at the repository level with or without maximum object size limitation. This configuration scenario supports a model where LFS is enabled for specific repositories only while the rest of the system remains unaffected.

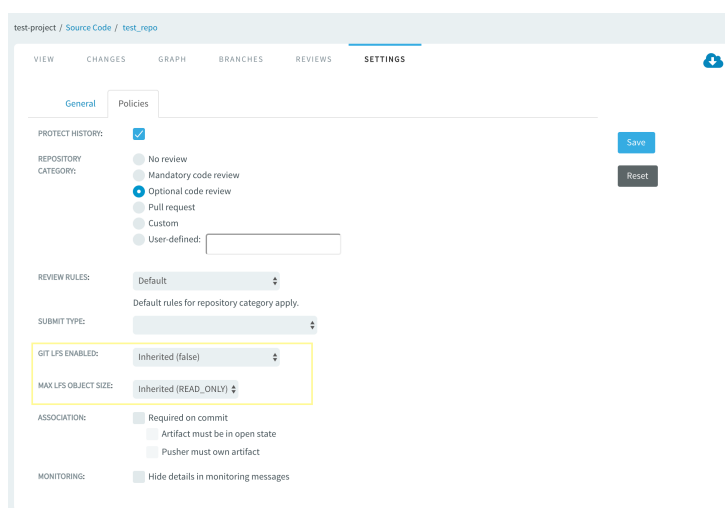
## Enable LFS for a repository

You can enable LFS for both existing and new repositories.

This section provides instructions to set up LFS for both exiting and new repositories.

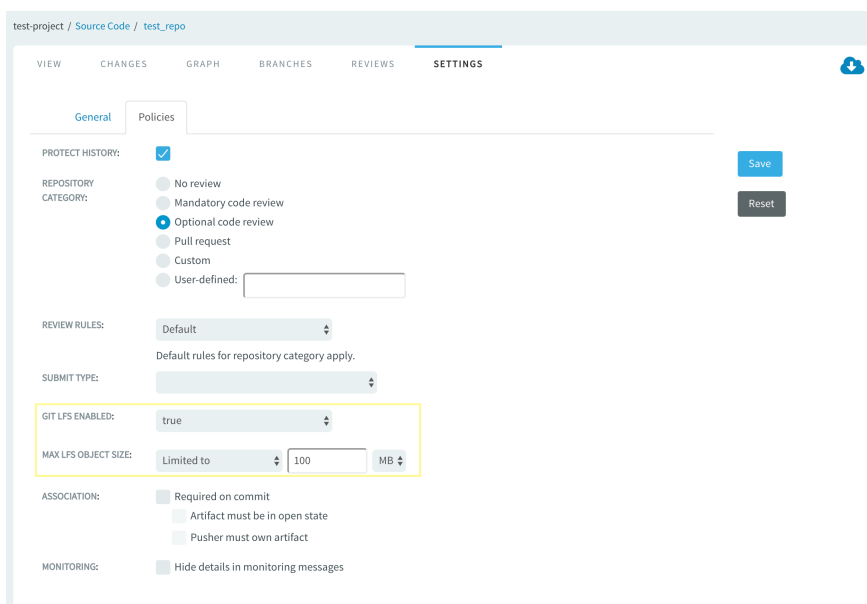
### Setting up LFS for existing repositories

1. Log on to TeamForge, select **Project Home** > **Source Code**, and select (click) a repository.
2. Select **Settings** > **Policies**.
3. To enable LFS, you must select the values for **MAX LFS OBJECT SIZE** and **GIT LFS ENABLED** fields.



- **MAX LFS OBJECT SIZE** (required field): Select one of the values: Inherited, Unlimited, Read-only or Limited to. By default Inherited value is READ\_ONLY. It means that once LFS data is pushed into repository it is always available for fetch/clone operation. Even if you switch to Unlimited, for example, and then decide to go back to READ\_ONLY at a later point in time for a given repository or integration, repository consistency is preserved and data would always be available. This is necessary to prevent situations where crucial binary data is always readable unless you rewrite the repository history to render such binary data unavailable. Select:
  - Inherited that makes this repository inherit the default Git-integration settings. Note that in case of Inherited, current default integration setting is shown for your reference.
  - Unlimited to support unlimited object size (size is not proactively limited by Gerrit but space availability still applies).
  - Limited to to limit maximum object size to a reasonable value, for example, 100MB.
  - Read-only that turns LFS in this repository to read-only mode.
- **GIT LFS ENABLED** (optional field): Select one of the values: Inherited, true or false. The inherited value is false by default. LFS is served from master Gerrit instance over HTTP/HTTPS protocol and therefore you must enable this parameter to extend the checkout URLs with LFS specific part for SSH protocol and replication. This is required as the LFS client, by default, uses the same URL (derives the protocol from it) that is used for fetch/clone/push operations, while pointing to the master Gerrit instance over HTTP/HTTPS to read/write data.

The following illustration shows a typical LFS configuration where the MAX LFS OBJECT SIZE is limited to 100 MB for a repository that's server over SSH:



### Setting up LFS when a new repository is created

4. Log on to TeamForge, select **Project Home > Source Code**, and click **Create Repository**.

5. Select the values for **MAX LFS OBJECT SIZE** and **GIT LFS ENABLED** fields while creating the new repository and click **Save**.

### Set up LFS client and work with large files

Download the Git LFS client for your platform. Instructions in this topic are valid for Git LFS client 1.3 and later. The checkout URLs are automatically extended with LFS part and you do not have to modify the checkout URL manually to have it working out of the box for SSH protocol or replication scenarios.

For downloading Git LFS client, see [Git Large File Storage](#) page.

#### Working with LFS over HTTP/HTTPS (without replication)

1. Use extended checkout URL.

**Example:**

```
git clone -c
```

```
'lfs.url=http://product_developer@main.server.collab.net/gerrit/
test_repo.git/info/lfs'
ssh://product_developer@gerrit.server.collab.net:29418/test_repo &&
cd "test_repo" && git
config user.name "Nancy S." && git config user.email
"nancy@example.com" && git config
url."ssh://gerrit.server.collab.net:29418".insteadOf "ssh://
main.server.collab.net:29418"
&& git config url."ssh://
product_developer@gerrit.server.collab.net:29418".insteadOf
"ssh://product_developer@main.server.collab.net:29418" && git config
url."ssh://
product_developer@main.server.collab.net:29418".pushInsteadOf
"ssh://product_developer@gerrit.server.collab.net:29418" && scp -P
29418
product_developer@gerrit.server.collab.net:hooks/commit-msg .git/
hooks/
```

2. Select the file types you'd like Git LFS to manage. You can configure additional file extensions anytime.

The following command tracks all `.jpg` images in a given working directory.

```
git lfs track *.jpg
```

3. Create a commit by adding the binary file and the technical file (`.gitattributes`) that is modified by Git LFS client.

```
git add IMG_0036.jpg .gitattributes
```

4. Commit and push the file(s) to the remote repository.

```
git commit -sm 'This is LFS test' && git push origin HEAD:master
```

The binary file that is successfully pushed to LFS manifests itself by having a reference file committed to Gerrit and you can check its content by going to TeamForge code browser for the given repository. Here is an example reference file. It contains Git LFS protocol version specification along with Git LFS object SHA and its size.

test-project / Source Code / test\_repo

VIEW CHANGES GRAPH BRANCHES REVIEWS SETTINGS

**a87ef86: This is LFS test**  
 Authored 10/24/2016 4:23 pm UTC (4 minutes ago) by Nancy S  
 P master

This is LFS test

Change-Id: I408b257765f0e8b18b2583e5418f8c35ec220f78  
 Signed-off-by: Nancy S <nancy@example.com>

master

Associations

File	Size	View
.gitattributes	0   1	✕ ↓
IMG_0036.jpg	0   3	✕ ↑


unified  side-by-side

```
1 version https://git-lfs.github.com/spec/v1
2 oid sha256:b5a0872ce5f5c92bc71aef3d482f61aa7ddea1fda0a5f93e3e6defc0b0113237
3 size 106977
```

## Add Git as a linked application

You can set up Git as a linked application to your TeamForge site.

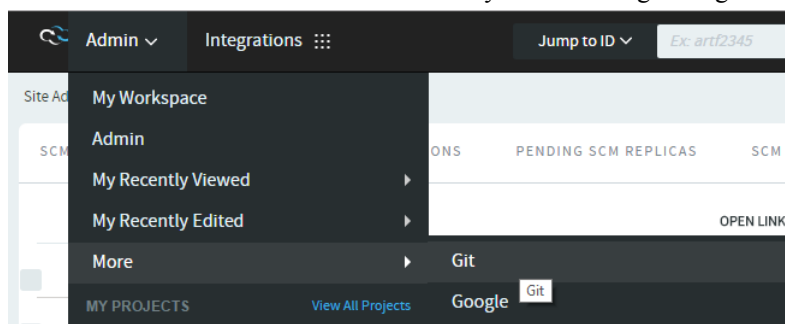
- Set up the URL `http://<TEAMFORGEHOSTNAME>/gerrit/sso/`.

 **Note:** The last "/" matters. Make sure you have it.

- For instructions on setting up a site-wide linked application in TeamForge, see [Create a site-wide linked application](#) on page 361.

Here's an example for Git:

A link for Git is added to the **More** menu in your TeamForge navigation bar.



Clicking **Git** displays the Git console in the main TeamForge window.

- 👉 **Note:** In TeamForge 7.2 and later versions, installing Git for the first time creates a site-wide linked application. From TeamForge 8.0 and later, a project-wide linked application is also created for projects in TeamForge that have at least one CVS repository. However, this behavior can be controlled by the `gerrit.config` property, `teamforge.createTFProjectLinkedApps`.

## Control the code review policy for Git repositories

With versions of the TeamForge–Git integration 6.2 and later, you can control all Gerrit Code Review features directly from TeamForge by specifying a code review policy as part of the TeamForge Git repository's description field.

For more information on Gerrit Code Review, see the [Gerrit documentation](#).

By default, the following code review policy options are provided:

- **Default (no code review):** All Gerrit review features are turned off and read/write access is enforced.
- **Mandatory review:** All code changes must be reviewed and read/write access is enforced.
- **Optional review:** The review feature is turned on but can be bypassed if necessary; read/write access is enforced.
- **Custom:** Access rights must be set manually in the Gerrit web interface; they will not be overridden by TeamForge. This specification is intended for advanced users who are familiar with Gerrit access rights and want to turn off “auto pilot”.

If you have access to the `TeamForgeGerritMappings.xml` file, you can add your own categories. The following animation illustrates the detailed mapping between SCM permissions, repository policies, and Gerrit access rights.

## Mandatory code reviews for Git repositories

When a mandatory review is specified, every change pushed to the repository must pass through a review process before it can get committed (merged) to the repository.


The screenshot shows the 'Create Repository' form with the following details:

- SERVER:** Git (dropdown menu)
- DIRECTORY NAME:** testrepo
- REPOSITORY NAME:** testrepo
- DESCRIPTION:** (empty text area)
- REPOSITORY CATEGORY:**
  - Default - no review
  - Pull request
  - Optional review
  - Mandatory review
  - Custom
  - User defined: (text input field)

Only TeamForge users with the **Source Code Admin** permission can bypass reviews.

Here's a list of permissions and what users with these permissions can do:

- **No access:** Users with no permissions cannot do anything.
- **View only:** Users with read permissions can read branches and push for reviews, and have -1 and +1 for reviews.
- **Commit/View:** Users with commit permissions can do everything read permissions would grant and in addition have -2, +2 for reviews. They can verify and submit permissions but have no right to bypass reviews.
- **Delete/View:** Users with delete permissions can do everything commit permissions would grant.
- **Source Code Admin:** Users with admin permissions can do everything delete permissions would grant and in addition push to and create any branch (bypassing review). They can rewrite history, forge the identity of the Gerrit server, and have the right to push tags, the right to upload merges, and the right to fine tune access rights in Gerrit for the Gerrit project involved.

 **Note:** In TeamForge 6.2, the code review policy for a Git repository is defined in the **Description** field. If you had specified the code review policy in TeamForge 6.2 and have now upgraded TeamForge to version 8.0 or later, you will see the appropriate code review option selected in the TeamForge user interface. The **Description** field will still display the [RepoCategory:<Category\_name>], but you can remove it as it does not have any effect in TeamForge.

## Optional code review for Git repositories


When an optional review is specified, every change submitted to the repository can be pushed for code review or directly pushed to the repository bypassing review. This depends on the TeamForge user having the appropriate

permissions — source code Delete/View or Commit/View permission for the former, or Source Code Admin permission for the latter.

The screenshot shows the 'Create Repository' form. The 'SERVER' dropdown is set to 'Git'. The 'DIRECTORY NAME' and 'REPOSITORY NAME' text boxes both contain 'testrepo'. The 'DESCRIPTION' field is empty. Under 'REPOSITORY CATEGORY', the 'Optional review' radio button is selected. Other options include 'Default - no review', 'Pull request', 'Mandatory review', 'Custom', and 'User defined:'.

Here's a list of permissions and what users with these permissions can do:

- No access: Users with no permissions cannot do anything.
- View only: Users with read permissions can read branches and push for reviews, and have -1 and +1 for reviews.
- Commit/View: Users with commit permissions can do everything read permissions would grant and in addition have -2, +2 for reviews. They can verify and submit permissions, push to/create any branch (bypassing review) and push tags.
- Delete/View: Users with delete permissions can do everything commit permissions would grant and in addition, have the right to rewrite history, upload merges and forge identity.
- Source Code Admin: Users with admin permissions can do everything delete permissions would grant and in addition push to/create any branch (bypassing review). They can rewrite history, forge the identity of the Git server, and have the right to push tags, the right to upload merges, and the right to fine tune access rights in Git for the Git project involved.

 **Note:** In TeamForge 6.2, the code review policy for a Git repository is defined in the **Description** field. If you had specified the code review policy in TeamForge 6.2 and have now upgraded TeamForge to version 8.0 or later, you will see the appropriate code review option selected in the TeamForge user interface. The **Description** field will still display the [RepoCategory:<Category\_name>], but you can remove it as it does not have any effect in TeamForge.

## Default code review for Git repositories

In TeamForge 8.0, the default code review policy — no code review required — is selected unless you choose some other policy.

**Create Repository**

SERVER:

DIRECTORY NAME:

REPOSITORY NAME:


DESCRIPTION:

REPOSITORY CATEGORY

- Default - no review
- Pull request
- Optional review
- Mandatory review
- Custom
- User defined:

Here's a list of permissions and what users with these permissions can do:

- No access: Users with no permissions cannot do anything.
- View only: Users with read permissions can only read branches.
- Commit/View: Users with commit permissions can do everything read permissions would grant and in addition, push to/create any branch and push tags.
- Delete/View: Users with delete permissions can do everything commit permissions would grant and in addition, have the right to rewrite history, upload merges and forge identity.
- Source Code Admin: Users with admin permissions can do everything delete permissions would grant. In addition, they can forge the identity of the Gerrit server, and have the right to fine tune access rights in Git for the Git project involved.

 **Note:** In TeamForge 6.2, the code review policy for a Git repository is defined in the **Description** field. If you had specified the code review policy in TeamForge 6.2 and have now upgraded TeamForge to version 8.0 or later, you will see the appropriate code review option selected in the TeamForge user interface. The **Description** field will still display the [RepoCategory:<Category\_name>], but you can remove it as it does not have any effect in TeamForge.

## Custom code review for Git repositories

When a custom code review is specified, users with the TeamForge Source Code Admin permission can directly fine tune permissions (access rights) in gerrit's web interface. Those changes will not be overridden by TeamForge.



**Create Repository**

SERVER:

DIRECTORY NAME\*:

REPOSITORY NAME\*:

DESCRIPTION:


REPOSITORY CATEGORY

- Default - no review
- Pull request
- Optional review
- Mandatory review
- Custom
- User defined:

For information on manually defining access rights in the Gerrit web interface, see [Update Git repository access permissions in Gerrit](#).

Here's a list of permissions and what users with these permissions can do:

- No access: Users with no permissions cannot do anything.
- View only: Users with read permissions cannot do anything unless added in Gerrit.
- Commit/View: Users with commit permissions cannot do anything unless added in Gerrit.
- Delete/View: Users with delete permissions cannot do anything unless added in Gerrit.
- Source Code Admin: Users with admin permissions have the right to fine tune access rights in Gerrit for the Gerrit project involved.

 **Note:** In TeamForge 6.2, the code review policy for a Git repository is defined in the **Description** field. If you had specified the code review policy in TeamForge 6.2 and have now upgraded TeamForge to version 8.0 or later, you will see the appropriate code review option selected in the TeamForge user interface. The **Description** field will still display the [RepoCategory:<Category\_name>], but you can remove it as it does not have any effect in TeamForge.

## Pull request reviews for Git repositories

Pull requests allow developers to collaborate with each other on a code change before merging it into another branch on a Git repository. Using a pull request, you notify others about a feature or fix change that needs attention.

### Create Repository

---

SERVER: Git ▼


DIRECTORY NAME:\* testrepo

REPOSITORY NAME:\* testrepo

DESCRIPTION:

REPOSITORY CATEGORY

- Default - no review
- Pull request
- Optional review
- Mandatory review
- Custom
- User defined:

 **Note:** In TeamForge 6.2, the code review policy for a Git repository is defined in the **Description** field. If you had specified the code review policy in TeamForge 6.2 and have now upgraded TeamForge to version 8.0 or later, you will see the appropriate code review option selected in the TeamForge user interface. The **Description** field will still display the [RepoCategory:<Category\_name>], but you can remove it as it does not have any effect in TeamForge.

## User defined reviews for Git repositories

Users can define their own code review policy.

**Create Repository**

SERVER:


DIRECTORY NAME:

REPOSITORY NAME:

DESCRIPTION:

REPOSITORY CATEGORY

- Default - no review
- Pull request
- Optional review
- Mandatory review
- Custom
- User defined:

 **Note:** In TeamForge 6.2, the code review policy for a Git repository is defined in the **Description** field. If you had specified the code review policy in TeamForge 6.2 and have now upgraded TeamForge to version 8.0 or later, you will see the appropriate code review option selected in the TeamForge user interface. The **Description** field will still display the [RepoCategory:<Category\_name>], but you can remove it as it does not have any effect in TeamForge.

## TeamForge-Git integration: History protection

History rewrites are non-fast-forward updates of remote `refs` and associated objects. History rewrites happen when a branch in a remote repository gets deleted, previously pushed commits get amended or filtered and forcefully re-pushed, or a remote branch/tag is pointed to an entirely different commit history.

History may get rewritten without leaving any trace of the previous state. Sometimes this behavior may be wanted — for example, in the case of removing code violating intellectual property, removing mistakenly committed large binary files or removing merged feature branches. The TeamForge-Git integration therefore does not disable the history rewrite feature, but instead enables it for SCM Administrators alone. However, since rewriting history might be easily abused and result in accidental data loss, we've introduced the History Protection feature as a safety net and necessity for ensuring proper audit compliance.

History protection archives rewritten changes and keeps backups of deleted branches. If history changes occur, an immutable backup `ref` is created in the remote repository, notification emails are sent to all members of the Gerrit Administrators group, and an event is logged in the audit log. The backed up `ref` can be restored into a new branch with any Git client (without needing physical file access to the Gerrit server). Gerrit site administrators can still decide to remove selected backup `refs` permanently.

## TeamForge-Git integration: Enable history protection

You can turn on history protection for the entire integration server or for an individual Git repository.

- To turn on history protection server-wide (for all the repositories hosted on the Git integration server), do the following:
  - a) Set `GERRIT_FORCE_HISTORY_PROTECTION=true` in `/opt/collabnet/teamforge/runtime/conf/runtime-options.conf`.
  - b) Restart the gerrit service.

```
/etc/init.d/collabnet restart gerrit
```



**Note:** If you enable history protection server-wide, it cannot be overridden (turned off) for an individual Git repository. It will remain in effect for all Git repositories on the integration server.

- To turn on history protection for an individual Git repository in a TeamForge project, select the **Protect History** option while creating the repository. For an existing repository, do the following:
  - a) On the **Source Code** page, select the Git repository and click **Edit**
  - b) Select the **Protect History** check box.
  - c) Click **Save**.

You can turn history protection on or off any time. However, your change will not be reflected in Gerrit immediately. It will be effective after the time that you defined as the regular refresh interval while installing the Git integration.

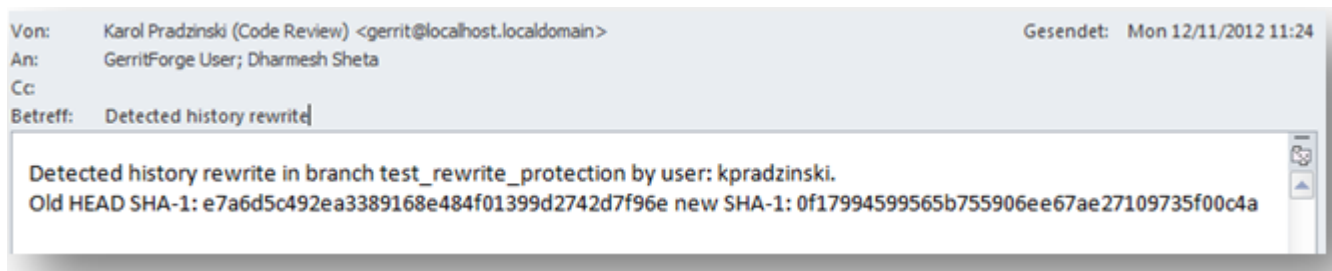
If you want your change to take effect immediately, do this right after you select or clear the **Protect History** check box: as a user with Source Code Admin permission, temporarily remove any user having a project role with any SCM permission, and then add that user back. This will trigger an immediate sync which will enable history protection. After that, the Gerrit Administrator will be able to see History Protection enabled in the Gerrit web interface (by logging in as a Gerrit Administrator and clicking the General link for the project with the name of the Git repository).

## History protection reports

Once history protection is turned on, any non-fast-forward push to a remote repository or deletion of a branch or tag on a remote repository is recorded and reported.

### Email notifications

When history is rewritten, an email is sent to the Administrator group members in Gerrit.



### Gerrit web interface

Every history rewrite event is logged and stored in the Gerrit database and visible in the Gerrit web interface. As Gerrit Administrator, you can —

- See rewritten history from **Project > Rewritten History**.

TeamForge by CollabNet

Projects My Workspace Admin History More Ex: artf2345

Gerrit

All My **Projects** People Plugins

List General Branches Access Dashboards Notification **History Rewrites** Create New Project

Change #, SHA-1, tr.id or owner:email

Project teamforge-auth

Type	Reference Name	User Name	Rewrite Date	Original Commit ID	New Commit ID
rewrite	master	jnicolai	2013 Apr 16 03:46:39	735b638287d7b9eb08f1af99dc5290388dbda858	9b58bf766a119c1674b6282d3e948b372808332a

- Restore history by clicking **Resurrect** and providing a name for the new branch.

TeamForge by CollabNet

Projects My Workspace Admin History More

Gerrit

All My **Projects** People Plugins

List General Branches Access Dashboards Notification **History Rewrites** Create New Project

Change #, SHA-1, tr.id or owner:email

Project teamforge-auth

Type	Reference Name	User Name	Rewrite Date	Original Commit ID	New Commit ID
rewrite	master	jnicolai	2013 Apr 16 03:46:39	735b638287d7b9eb08f1af99dc5290388dbda858	9b58bf766a119c1674b6282d3e948b372808332a

Resurrection

Resurrect refs/rewrite/20130416034639-master-9b58bf766a119c1674b6282d3e948b372808332a-jnicolai as

rescued

OK Cancel

- Permanently remove a branch by clicking **Delete Permanently**.

### Git command line

You can use a standard Git client and run `git fetch && git ls-remote` for information on rewritten and deleted branches.

```
$ git fetch && git ls-remote origin
03412aed0c60a3e4d480ac3d135369431645ab25 HEAD
03412aed0c60a3e4d480ac3d135369431645ab25 refs/heads/master
1814a1b0a4f1351db62a5b5fd74ceff87c3c2076 refs/rewrite/20121102215850-master-03412aed0c60a3e4d480ac3d135369431645ab25
```

You can view entries in `refs/rewrite` (for non-fast-forward pushes) and `refs/delete` using the Git `ls-remote` command only if read access is granted to `refs/*`. Gerrit will prevent any other action such as `delete/force-update` on those special `refs` for all users including administrators.

### Audit log entries

The following events are logged in `/opt/collabnet/gerrit/logs/gerrit.audit.log`:

- Remote branches are deleted.
- History is rewritten (non-fast-forward push).
- Backup branches are resurrected.
- Backup branches are permanently deleted.
- History Protection is turned on or off.

## TeamForge-Git integration: History protection FAQ

Questions on history protection for Git repositories in TeamForge.

**What is history protection?**

History rewrites are non-fast-forward updates of remote `refs` and associated objects. History rewrites happen when a branch in a remote repository gets deleted, previously pushed commits get amended/tree filtered and forcefully re-pushed, or a remote branch/tag is pointed to an entirely different commit history. For more information, see [History Protection](#).

**Is it possible to turn on history protection for all Git repositories hosted on a Git integration server? If yes, how?**

Yes, a user with file system access to the Unix machine where the Git integration is hosted can turn on history protection. First set the property `forceHistoryProtection = true` in `/opt/collabnet/gerrit/etc/gerrit.config` in the `[gerrit]` section. Then restart the `gerrit` service by running `$ service gerrit restart` on a shell.

With TeamForge 7.1 and later, do the following:

- Set `GERRIT_FORCE_HISTORY_PROTECTION=true` in `/opt/collabnet/teamforge/runtime/conf/runtime-options.conf`.
- Restart the `gerrit` service.

```
/etc/init.d/collabnet restart
gerrit
```

The `GERRIT_FORCE_HISTORY_PROTECTION` site-options token is no longer supported in TeamForge 17.4 and later. Remove this token from the `site-options.conf` file while upgrading to TeamForge 17.4 or later. History Protection is enabled by default in TeamForge 17.4 and later. However, Site Administrators can disable History Protection at the site level if need be, after which Project Administrators can choose to have History Protection enabled or disabled for individual repositories.

**I've enabled Protect History for my TeamForge project's Git repository. Will this be effective immediately?**

To enable history protection immediately, a TeamForge user with the Source Code Admin permission must do this right after enabling **Protect History**: temporarily remove any user with a project role with any SCM permission, and then add that user back. This will trigger an immediate sync after which history will be protected for the Git repository. Otherwise, history protection will be enabled after a periodical sync.

**Can I turn off history protection for any particular Git repository when it is enabled server-wide?**

No, when history protection is enabled server-wide for the Git integration server, it cannot be turned off for a particular Git repository.

**Where can I see the backup branches generated by history protection?**

Backup branches are generated based on the type of History Rewrite. For a remote branch that is deleted, this is under `refs/delete`. For a non-fast-forward push, this is under `refs/rewrite` with the branch name containing the timestamp, original branch and the user who rewrote history, for example, `refs/delete/20121112042512-test--david`.

<b>Who can resurrect or permanently delete backup branches?</b>	A user who is a member of the Gerrit Administrator group can resurrect or permanently delete backup branches. By default, the TeamForge site administrator whose credentials are used for running the post-installation script is part of the Gerrit Administrator group.
<b>Who can see backup branches?</b>	By default, a TeamForge user with SCM View (or more) permission can see all backup branches by executing <code>git fetch &amp;&amp; git ls-remote origin</code> . In Gerrit, the user must be part of a group which has at least read access for <code>refs/delete</code> and <code>refs/rewrite</code> for the given Gerrit project (TeamForge Git repository).
<b>Are the backup branches under <code>refs/rewrite</code> and <code>refs/delete</code> protected from Git garbage collection which removes unreferenced objects?</b>	Yes, objects in backup branches under <code>refs/rewrite</code> and <code>refs/deleted</code> are referenced and cannot be cleaned up by Git's garbage collection.
<b>Do backup branches take up a lot of disk space on the Git server?</b>	The backup branches on the Git server are mainly <i>Git objects</i> that are compressed deltas of original file versions. Git regularly compresses these objects to save disk space.
<b>What is the difference between Git <i>reflog</i> and History Protect?</b>	In Git, <code>reflog</code> records all activity on a branch, while History Protect only reports deleted branches/tags and history rewrites (non-fast-forward pushes) For more information, see <i>Git reflog vs TeamForge-Git integration History Protect</i> <a href="#">Git reflog vs Perforce History Protect</a> .

## TeamForge-Git integration: FAQ

---

Use this background information to set up, maintain, support and work with the TeamForge-Git integration.

### TeamForge-Git integration: General usage FAQ

General usage questions about the TeamForge-Git integration.

<b>How can I log into Gerrit?</b>	If your administrator has set up Gerrit as a linked application to Teamforge, you will automatically be logged into Gerrit (SSO) when you click its link. If not, access the URL <code>http(s)://&lt;yourtfinstance&gt;/scm integration server/&gt;gerrit/</code> and provide your TeamForge credentials.
<b>What are the Git protocols that work with the Git repositories managed by TeamForge?</b>	The Git integration currently allows you to access a Git repository using SSH. That said, you must have generated an SSH key pair and uploaded the SSH public key to Teamforge in <b>My Settings &gt; Authorized keys</b> .  Alternatively, you can use <code>http(s)</code> to clone and push to Git repositories. In this case, you can authenticate using your TeamForge user name and password.
<b>How can I use <code>http(s)</code> for accessing Git repositories?</b>	From 8.2.x of the TeamForge-Git integration, there are two possibilities for accessing Git repositories over <code>http(s)</code> .

The clone URL for http(s) access follows this convention:

```
git clone https://
$USERNAME@<yourtfinstance/scm
integration server>/gerrit/p/
<TFreponame>
```

When you run the above command on your Git client, you will be asked to provide your credentials. Use the same credentials you use to log into TeamForge's web interface.

**How do I generate an SSH key pair?**

You can generate an SSH key pair on a Unix machine by running the following shell command:

```
$ ssh-keygen -t rsa
```

(You will be asked to provide the location to store the key pair. The default is the home directory of the logged-in user.)

**After installing a Git client, I am able to clone a Git repository into my local work directory. However, I am not able to "push" anything to the remote repository in spite of having view and commit permissions. What should I do ?**


Right after you clone, but before you commit any changes locally, you will need to configure Git if you haven't already.

```
$ git config --global user.name
"<TeamForge username>"
$ git config --global user.email
"<email used in TeamForge for the
user>"
```

You should now be able to push your changes.

**Is a commit association created in TeamForge after I push my commit to a remote Git repository?**

Yes, when you push a local commit to the remote repository, an association will get created if the commit message contains a reference to a TeamForge item such as a tracker artifact, wiki or document in square brackets, for example [artf1234].

 **Note:** A commit association will not be created if you push your commit to Gerrit's "review branch" (push for review). It will be created once the change is merged into the real branch.

**What happens if the TeamForge site is down or there are some network problems — will the Git integration still work?**

The Git integration will still work, but with the following limitations:

- If the TeamForge site is down, users will not be able to see commit associations created in TeamForge, but still be able to push commits to a Git repository.
- If the Git integration is hosted in LOCAL mode, network-related problems would definitely prevent changes being pushed to a Git repository.
- If the Git integration is hosted in REMOTE mode, the synchronization of roles and permissions will be cached during the period when TeamForge is down; Git will function with the roles and permissions synched already.



**What is a "Jumbo Push"?**

In contrast to Subversion, Git has the concept of local commits that stay in the local environment of a user, and at some point, get pushed to a remote repository all at once. This push checks in changes from all commits into the remote repository. For each of those commits, a commit object appears in the TeamForge (Source Code component). So, one push can have an unlimited number of commits and thus commit objects in TeamForge. You can, however, define the threshold for a single push based on how many commits should generate a commit object. A push' containing commits beyond that threshold is called a "Jumbo Push".

You can configure the Jumbo Push threshold by updating the site option token, [GERRIT\\_GIT\\_PUSH\\_THRESHOLD](#) in the `site-options.conf` file. You have to run the post-installation script after rebuilding the runtime environment. When the Git and Teamforge are hosted on the same server, the runtime involves TeamForge downtime.

**What objects and relationships are mapped between TeamForge and the Git integration?**

See the README (APPENDIX, Relationship and Object mapping section) or [Mappings between TeamForge and Gerrit](#) on page 621.

**When are the objects and relationships synchronized between TeamForge and the Git Integration?**


TeamForge project roles, project role SCM permissions, global groups, SCM repositories, and global group/project role membership are synched in two ways:

- Synchronously: after a regular interval (configurable using the post-installation script)
- Asynchronously: whenever there is a change related to roles or permissions within Teamforge, it triggers the sync between TeamForge and the Git integration.

TeamForge repositories are only synched if there is at least one project role with SCM permissions present in the corresponding TeamForge project.

TeamForge users are provisioned in Gerrit whenever you

- Change their authorized keys in TeamForge
- Log into Gerrit by clicking the linked application link or using TeamForge user name and password
- Access GitWeb (web interface for a Git repository) by clicking a Git repository link in the TeamForge Source Code page

 **Note:** Changes in Gerrit are not synched back to TeamForge.

**Where can I find system logs for the Git integration?**

You can find the logs under `/opt/collabnet/gerrit/logs/`. For more on log files, refer to the README or [Gerrit directory structure, connectivity and more](#) on page 624.

**Can I bypass Gerrit and access a Git repository directly?**

No, Gerrit is used to enforce TeamForge access permissions.

**I deleted a TeamForge Git SCM repository but the corresponding Gerrit project does not get deleted. What's wrong?**

1. Delete the Git repository in TeamForge.
2. Go to the Git repository project in Gerrit.
3. Go to **Projects > General** and delete the Gerrit project.

You can delete the repository even if there are open changes (repository is permanently deleted) with an option to preserve the repository, if required. Select **Preserve Repository** if required.

**How can I import an existing Git repository into Gerrit?**

You can import an existing Git repository into Gerrit as a project admin from a local machine or as a System Admin from the server.

**Option 1:** If you are project admin, create a new repository and configure your account so that it has at least Delete/View SCM permissions for the one in TeamForge. Clone your existing repository and force push its content to the empty TeamForge repository:

```
git clone --mirror [url of repo to
be imported]
cd reponame
git gc
git remote add dest [url to empty
TeamForge Git repository]
git push -f --tags dest refs/heads/
*:refs/heads/*
```

**Option 2:** If you are site admin, create a new repository from the TeamForge UI and perform the following steps from the server as a Gerrit system user to import a repository from the source into TeamForge:

```
# su gerrit
$ cd /tmp
$ git clone --mirror [url of repo to
be imported]
$ cd reponame.git
$ git gc
$ git remote add dest file:///
gitroot/reponame.git/
$ git push -f --tags dest refs/
heads/*:refs/heads/*
```



**Note:** When importing a repository previously hosted on a different Gerrit server, do not push the review branches as Gerrit numerical change numbers are not globally unique and duplication will result in wrong email notifications and problems submitting open reviews.

**In the TeamForge web interface, I see the repository root parameter for Git set to "/tmp". Can I change that?**

For backward compatibility reasons, this parameter has to be set to "/tmp". It does not affect where Gerrit actually stores its Git repositories — this is at /gitroot.

**Do we have default hook scripts available for Git in TeamForge?**

Associating artifacts based on commit messages and blocking commits without a commit message is a core TeamForge mechanism that is supported by Git as well.

**Do we have email alerts for Git in TeamForge? If yes, where do we configure it?**

To add hook scripts, see [Gerrit Code Review - Hooks](#).

Email alerts based on TeamForge commits is a core TeamForge feature, independent of the SCM involved. In addition, Gerrit sends out review emails using the SMTP server specified during installation (it defaults to the TeamForge SMTP server). The mail template is explained in [Gerrit Code Review - Mail Templates](#).

The blog post <http://blogs.collab.net/teamforge/collabnet-gerrit-notifications-for-all-who-miss-the-good-ol-git-push-notifications> explains how to send information on Git pushes to Teamforge forums (which act as mailing lists too).

**Do we have Role Based Access Control and Path Based Permissions for Git in TeamForge?**

We support all SCM permission cluster options for TeamForge project roles, default access permissions, project admin permissions along with the site-wide roles and site admin permissions. However, path-based permissions are not relevant in Git since a Git commit always contains all files. If we did not ship certain files, this would result in a checksum error. Gerrit supports branch-based permissions though.

For more information on branch-based permissions, see [CollabNet's blog post](#).

## Teamforge-Git integration: Technical concepts FAQ

Questions on some of the Teamforge-Git integration's more technical aspects.

**What does the directory structure look like for a Git integration?**

See the README or [Gerrit directory structure, connectivity and more](#) on page 624.

**Where can I find more information about Gerrit?**

For more information on Gerrit, see the [Gerrit Community Documentation](#) page.

**Which ports does the Git Integration use? My organization has a strict firewall policy, and I need to know which ports to make available for the Git integration.**

The Git integration uses 3 ports: 9080,9081, and 29418. See the README or [Gerrit directory structure, connectivity and more](#) on page 624 for more information. For the integration, Git integration uses 3 ports(9080,9081,29418 follow details in README) Only port 29418 should be exposed by the firewall.

**I get a delay whenever I do a Git push, but not when I do a Git fetch. What is wrong?**

When you do a Git push, Gerrit tries to do a reverse lookup of your IP address. If the nameserver configured for your Git integration server cannot do this reverse lookup, it will result in a timeout. You need to configure your nameserver list (`/etc/resolv.conf`) correctly. For further information, see [this core Java bug](#).

**I ran Gerrit manually (without the service script; now my secure config file is gone and Gerrit does not start up. What happened and how can I fix this?**

If Gerrit is run with a different Unix user than `gerrit`, newly created and modified files may not belong to the `gerrit` user any longer. As a consequence, when you try to restart Gerrit using its services script (which switches to the `gerrit` user), Gerrit might not start up due to wrong file permissions. If Gerrit detects that the permissions of its `secure config` file have been tampered with, it even removes this file. You should, therefore, only run Gerrit using the service script provide, and reconfigure it by running the post-

install script again. You can fix incorrect permissions by running the following (as sudo or root):

```
chown -R gerrit.gerrit /opt/
collabnet/gerrit
```

**I deleted the dedicated Teamforge Gerrit user account in TeamForge, and SCM permission synch is no longer working. How can I recover from this situation?**

- The easy, and recommended, approach is to ask CollabNet's Professional Services to undelete the TeamForge user in question.
- Otherwise, you would have to create a new dedicated site admin user in TeamForge, shut down Gerrit, re-run the post-install script and provide the credentials of that user. Then, you would have to start Gerrit again, and log in with as that new user via the web interface. You will see that the user does not have any special admin permissions. If you still have a working user in Gerrit's administrator group, you could add the dedicated Gerrit user to that group using Gerrit's web interface. If not, you would have to manually add the new user to the Gerrit administrator group by shutting down Gerrit, removing all files from its caching directory, inserting the `user id` of the new user into Gerrit's Postgres `reviewdb` DB group/user membership table, and starting Gerrit again. Since this probably requires you to consult CollabNet's Professional Services as well, we strongly recommend the previous option (undeleting the previously removed user).

**Why does the Registered Users group has StreamEvents capability in Git Integration v8.1.x by default?**

- As Gerrit 2.7 **Stream Events** capability is required for user whose account has been used to monitor Gerrit Events on repositories hosted on Gerrit.
- As Jenkins *Gerrit Trigger plug-in* uses such a capability to monitor Gerrit events.
- If you have been using **Jenkins CI with Gerrit Trigger** plug-in to automatically verify code review request and already upgraded to Git Integration v8.1.x, the **Registered Users** group has been given this capability by default. Therefore, you do not need add this capability manually and your Jenkins CI with Gerrit Trigger continues to function as usual.
- In case you're using **Jenkins CI with Gerrit Trigger** plug-in, you may remove the **Stream Events** capability as a user who is part of the Gerrit **Administrators** group.
  - To remove **Stream Events**:
    1. Log on to Gerrit web UI.
    2. Select the **Projects** tab.
    3. Select **All Projects > Access**.
    4. Click **Edit**.
    5. Delete the **Stream Events** drop-down list.
    6. Click **Save Changes**

The screenshot shows the TeamForge interface for configuring project access. The 'Projects' tab is selected, and the 'Access' sub-tab is active. The 'Stream Events' capability is highlighted with a red box, showing it is set to 'ALLOW' for 'Registered Users'. Below this, a 'Commit Message (optional):' text area is visible, and the 'Save Changes' button is also highlighted with a red box.

## TeamForge Git integration reference


Look here for information on the mappings between TeamForge and Gerrit, Gerrit access rights, directory structure, connectivity, logs and configuration properties, and differences compared to vanilla Gerrit.

### Mappings between TeamForge and Gerrit

These tables shows how objects and relationships are mapped between TeamForge and Gerrit.

TeamForge Object	Gerrit Object
TeamForge project	Project
SCM repository in TeamForge project (containing project roles with SCM permissions)	Project
Project Role	Group

TeamForge Object	Gerrit Object
User Group	Group
User	User
Site-wide role (TeamForge 8.0 and later)	Group

 **Note:** In this release, global roles, project groups, project group roles, and project role inheritance are not supported.

TeamForge Relationship	Gerrit Relationship
Git repository is part of a TeamForge project.	Gerrit project corresponding to the Git repository inherits from the Gerrit project corresponding to the TeamForge project (TeamForge-Projects/<TeamForge project id>).
TeamForge project <child> has a parent TeamForge project <parent>.	Gerrit project <child> inherits from the Gerrit project <parent>.
TeamForge project top is a top-level project.	Gerrit project <top> inherits from Gerrit project. TeamForge-Projects which in turn inherits from All-Projects.
User has a TeamForge Project Role.	User is part of the Group which corresponds to the TeamForge Project Role.
User is part of a User Group that is assigned a Project Role.	User is part of a Group (which corresponds to a TeamForge Project Role).
User is part of a User Group.	User is part of a Group (which corresponds to a TeamForge User Group).
Project Role is assigned an SCM permission (such as Admin, Delete and View, View and Commit, View Only, None).	Corresponding group is assigned Gerrit access rights matching the assigned TeamForge SCM permissions. Those access rights are determined by the code review policy of the corresponding TeamForge repository.
Site-wide role is assigned an SCM permission. (TeamForge 8.0 and later only).	Corresponding Gerrit groups are assigned Gerrit access rights matching the assigned TeamForge SCM permissions. Those access rights are determined by the code review policy of the TeamForge repository and hence may vary between repositories.
Guests, All Site Users, All Logged in Users, All Non-Restricted Users or Project Members have SCM permissions associated using TeamForge's Default Access Permissions (TeamForge 8.0 and later only).	Corresponding Gerrit groups are assigned Gerrit access rights matching the assigned TeamForge SCM permissions. Those access rights are determined by the code review policy of the TeamForge repository and hence may vary between repositories.
User is a site admin in TeamForge.	User is part of Gerrit groups. TeamForge: Site Admins. TeamForge: Site-wide Project Admin Access. Private Project - Site-wide Admin Access. Public Project - Site-wide Admin Access. Gated Project - Site-wide Admin Access. Site admins have OWN and READ permissions for all Gerrit projects and the rights granted by the SCM Admin permission

TeamForge Relationship	Gerrit Relationship
	(depends on the code review policy of the Git repository in question).
User is a project admin in TeamForge.	User is part of Gerrit group. TeamForge: Project Admin for <TF project id>. which has OWN and READ permissions for all Git repositories of the corresponding TeamForge project.
User is non restricted in TeamForge (TeamForge 8.0 and later only).	User belongs to Gerrit group. TeamForge: Non-restricted Users.
User is a member of a TeamForge project (TeamForge 8.0 and later only).	User belongs to Gerrit group. TeamForge: Direct Project Member of <TF project id>.
User is member of a user group associated to a TeamForge project role (TeamForge 8.0 and later only).	User belongs to Gerrit group. TeamForge: Project Member of <TF project id>.
User has a site-wide role that has SCM permissions or a site-wide project admin permissions (TeamForge 8.0 and later only).	User is part of Gerrit group. TeamForge : Site-wide Role: <name of TeamForge Site-wide role> and - depending on the prevent inheritance to private projects flag, SCM permissions and project admin permissions - TeamForge : Site-wide Project Admin Access Public Project - Site-wide Admin Access Gated Project - Site-wide Admin Access Private Project - Site-wide Admin Access Public Project - Site-wide Delete Access Gated Project - Site-wide Delete Access Private Project - Site-wide Delete Access Public Project - Site-wide Commit Access Gated Project - Site-wide Commit Access Private Project - Site-wide Commit Access Public Project - Site-wide View Access Gated Project - Site-wide View Access Private Project - Site-wide View Access
User has a TeamForge account.	User belongs to the Gerrit group. Registered Users.
User is not logged into TeamForge yet.	User belongs to the Gerrit group. Anonymous Users.

TeamForge Relationship	Gerrit Relationship
	(as all logged in users do too).

## Gerrit directory structure, connectivity and more

Here's some reference information on Gerrit's scalability and hardware requirements, file/directory structure, database, ports and connectivity.

### Scalability and hardware requirements

- [Click here](#) for estimated requirements and availability assumptions.
- [Click here](#) for hardware considerations.
- [Click here](#) for a detailed description of performance-related settings.
- [Click here](#) for Gerrit performance cheat sheet.

### Gerrit directory structure

Gerrit expects a standardized directory structure under the GERRIT\_SITE directory: `/opt/collabnet/gerrit`, the gerrit user's home directory. The `/etc/default/gerritcodereview` file contains the GERRIT\_SITE variable and points to `/opt/collabnet/gerrit`.

**Sub-directories of GERRIT\_SITE** (`/opt/collabnet/gerrit`):

- GERRIT\_SITE

- `.ssh`: contains the SSH key for the gerrit user; generated during installation and needs to be backed up.

- `bin`: binaries, startup script

- `gerrit.war`: the main Gerrit service

- `gerrit-sync.jar`: the Gerrit-TeamForge synchronization service

- `gerrit.sh`: SYSV-style init script; launches and shuts down; linked to `/etc/init.d/gerrit`

- `cache`: disk cache; does not need to be backed up; can always be re-generated on the fly.

- `etc`: contains all configuration information; needs to be backed up.

- `gerrit.config`: Gerrit's main configuration file

- `secure.config`: contains obfuscated passwords and secrets

- `log4j-ng.properties`: contains logging settings in log4j format

- `TeamForgeGerritMappings.xml`: defines how TeamForge access permissions are mapped to Gerrit access rights

- `GerritSite.css`: .css file for Gerrit Web interface branding

- `GerritSiteFooter.html`: renders the customized footer in the Gerrit Web interface

- `lib`: libraries, potentially customer-specific extensions, treat like the `bin` directory

- `logs`: Gerrit log files; default configuration rotates logs daily, gzips old logs; debug files are rotated after they reach 10 MB; we keep 10 copies. You can make changes in Gerrit's `log4j-ng.properties` file in `/opt/collabnet/gerrit/etc`.

- `audit.log`: audit events

- `system.log`: INFO-level logging

- `sshd.log`: logs connections to Gerrit's SSH port (not system shell SSH)

- `*.gc.log`: information on garbage collector usage

Note: In co-hosted mode, TeamForge log rotation behavior will be used as default.

- `plugins`: The Gerrit plugins directory includes jar files for all CollabNet plugins — TeamForge-Sync, TeamForge-History Protection and so on.

- `static`: The content of this directory is served as `gerrit/static/` to http clients. You can include logos and other custom content you refer to in your branding here. By default, it contains the schema for the TeamForge-to-Gerrit access right mapping file which can be accessed from a browser at `http://<Git Integration server>/gerrit/static/TeamForgeGerritMappings-8.0.0.xsd`



You need to back up this directory if it includes custom content.

- `doc`: contains `README-NG.pdf`, `gerrit.config` and `log4j-ng` properties templates.

- `/gitroot`: default location for Git repositories. The default location can be changed using the setting in `gerrit.config` or by symlinking the directory. You need to back up this directory.

## Database

Gerrit stores runtime information about users, groups, code reviews and commits in a PostgreSQL database called `reviewdb` by default. You need to back up this database. If the Git integration is installed on the TeamForge server, it will use the same PostgreSQL install as TeamForge. If the integration is on a separate server, it will use a local PostgreSQL installation on that server. During installation, a new `gerrit` role and `reviewdb` schema get created. Note that Git does not, at any point, access TeamForge schema within the Postgres database.

To access `reviewdb` from Gerrit, the following line will be added by the installer to `/var/lib/pgsql/9.0/pg_hba.conf`:

```
host    reviewdb    gerrit    samehost    md5
```

In the "free-form" mode, we do not have full control over the Postgres installation and `hb_pga.conf` may contain the following conflicting line:

```
host    all    all    127.0.0.1/32    ident
```

If it exists, you need to comment it out (using `#`) to avoid clashing with this line inserted by installer:

```
host    reviewdb    gerrit    samehost    md5
```

## Ports and connectivity

Gerrit opens three (TCP) ports: 9080, 9081 and 29418. You should expose only port 29418 outside localhost.

### 9080: the Gerrit Web interface

This port will be proxied by Apache `conf` (prefix `/gerrit`) and doesn't need to be externally accessible. Do not expose this port to the outside.

### 9081: Gerrit-sync web service (REST)


In the Local (or co-hosted mode), TeamForge talks to the Gerrit REST API over localhost. Gerrit talks to TeamForge over its default SOAP URL.

The Git integration needs bidirectional connectivity to the TeamForge host. In the Remote (or distributed) mode, TeamForge talks to the Gerrit REST API over an Apache proxy rule (SSL-enabled if Apache is SSL-enabled on the integration server where Gerrit is running). Gerrit talks to TeamForge over its default SOAP URL.

Do not expose this port to the outside.

### 29418: Gerrit SSH access

Developers use the SSH protocol to push and pull source code to and from Gerrit. This port needs to be open to users.

 **Note:** Gerrit ships its own SSH implementation and offers no shell access over this port.

## Access rights in Gerrit 2.13.x

The Git integration maps Gerrit access rights to TeamForge Role Based Access Control (RBAC) permissions.

With the current version of the integration, `gerritforge.mappings` (including custom changes you might have done) is automatically migrated to an XML file located at `/opt/collabnet/gerrit/etc/TeamForgeGerritMappings.xml`. The following table shows how TeamForge RBAC permissions are now mapped to Gerrit access rights by default.

Code review policy	TeamForge permission cluster	Gerrit access right	
Default	SCM None	-	
	SCM View Only	Read	
	SCM Commit/View		Read
			Push
			Create Reference
			Push Annotated Tag (refs/tags/*) Push Signed Tag (refs/tags/*)
	SCM Delete/View		Read
			Push (forcePush)
			Create Reference
			Forge Author Identity
Forge Committer Identity Push Annotated Tag (refs/tags/*) Push Signed Tag (refs/tags/*)			
SCM Admin		Read	
		Push (forcePush)	
		Create Reference	
		Forge Author Identity	
		Forge Committer Identity	
		Forge Server Identity	
		Owner	
		Abandon Push Annotated Tag (refs/tags/*) Push Signed Tag (refs/tags/*)	
Optional Review	SCM None	-	
	SCM View Only		Read
			View Drafts
			Publish Drafts
			Code Review -1,1 Push (refs/for/refs/*) Rebase(refs/for/refs/*)
	SCM Commit/View		Read
			View Drafts
Publish Drafts Code Review -2,2 Verify -1,1 Submit Push Create Reference Rebase (refs/for/refs/*) Push Annotated Tag(refs/tags/*) Push Signed Tag (refs/tags/*)			

Code review policy	TeamForge permission cluster	Gerrit access right
	SCM Delete/View	Read View Drafts Publish Drafts Code Review -2,2 Verify -1,1 Submit Push (forcePush) Create Reference Rebase (refs/for/refs/*) Create References Push Signed Tag (refs/tags/*) Push Annotated Tag (refs/tags/*) Push Merges(refs/for/refs/*) Forge Author Identity Forge Committer Identity
	SCM Admin	Read View Drafts Publish Drafts Delete Drafts Code Review -2,2 Verify -1,1 Submit Push (forcePush) Create Reference Owner Abandon Rebase (refs/for/refs/*) Create References Push Signed Tag (refs/tags/*) Push Annotated Tag (refs/tags/*) Push Merges(refs/for/refs/*) Forge Author Identity Forge Committer Identity Forge Server Identity
Mandatory Review	SCM None	-
	SCM View Only	Read View Drafts Publish Drafts Code Review -2,2 Push (refs/for/refs/*) Rebase (refs/for/refs/*)
	SCM Commit/View	Read View Drafts Publish Drafts Code Review -2,2 Verify -1,1 Submit Push(refs/for/refs/*) Rebase (refs/for/refs/*)

Code review policy	TeamForge permission cluster	Gerrit access right
	SCM Delete/View	Read View Drafts Publish Drafts Code Review -2,2 Verify -1,1 Submit Push(refs/for/refs/*) Rebase (refs/for/refs/*)
	SCM Admin	Read View Drafts Publish Drafts Delete Drafts Code Review -2,2 Verify -1,1 Submit Push (forcePush) Create Reference Owner Abandon Rebase (refs/for/refs/*) Push Annotated Tag(refs/tags/*) Push Signed Tag (refs/tags/*) Create References Push Merges(refs/for/refs/*) Forge Author Identity Forge Committer Identity Forge Server Identity

To make changes to the mappings, modify the `/opt/collabnet/gerrit/etc/TeamForgeGerritMappings.xml` file on the server where your Git integration is hosted. Make sure that the resulting XML structure complies with this schema: <https://forge.collab.net/gerrit/static/TeamForgeGerritMappings-8.0.0.xsd>.

## Gerrit configuration properties

Gerrit provides many configuration options. In addition, CollabNet Gerrit plugins also have configuration options.

For more information on Gerrit's configuration options, see [Gerrit Code Review - Configuration](#). In addition, see [Gerrit performance cheat sheet](#) to know more about tuning Gerrit for optimal performance.

CollabNet Gerrit plugins have these configuration options:

### Section.teamforge

#### **teamforge.cache-path**

Location where Gerrit and CollabNet Gerrit plugin store caches. By default, this is at `/opt/collabnet/gerrit/cache`. We advise that it not be changed.

#### **teamforge.cache-ttl**

Time-to-live for Gerrit caches in seconds. The default value is 300.

#### **teamforge.apiPort**

Port over which TeamForge communicates with the Git integration. The default value is 9081

<b>teamforge.refreshTimeOut</b>	Interval in seconds after which the Git integration synchronizes with TeamForge. The default value is 3600.
<b>teamforge.jumboPushThreshold</b>	The number of commits in one Git push beyond which the Git integration creates only a single commit object in TeamForge. The default value is 30.
<b>teamforge.externalSystemId</b>	ID of the TeamForge external integration system. The value of this property is set by the post-installation script when the Git integration is first installed.
<b>teamforge.url</b>	Host URL of the TeamForge site with which Git is integrated. The value of this property is set by the post-installation script when the Git integration is first installed.
<b>teamforge.allowPushIfTeamForgeConnectionIsDown</b>	TeamForge commit objects are validated prior to creation. When the value of this property is false and connection to TeamForge is down, validation fails. When the value of this property is true, validation and creation of commit objects are postponed until the connection to TeamForge is restored. The default value is false.
<b>teamforge.parallelRemoteCallLimit</b>	TeamForge is able to handle a certain number of parallel connections. This parameter was introduced in order to avoid TeamForge "is out of service" issues. The default value is 9.
<b>teamforge.maxRemoteCallRetry</b>	This parameter was introduced in order to specify the number of retry attempts for calls to TeamForge before connection failure is returned. The default value is 3.
<b>teamforge.credentialsCache</b>	When the value of this property is set to true, users' credentials are cached for the <code>teamforge.credentialsCacheTimeout</code> amount of time and used to authorize actions in case of TeamForge connection outage. The default value is true.
<b>teamforge.credentialsCacheTimeout</b>	Interval (in Seconds) after which the credentials cache expires. The default value is 3600.
<b>teamforge.reconnectInterval</b>	When the "TeamForge connection is down" state is detected, and the number of seconds exceeds the value of this parameter, attempts to restore connection are performed periodically. The default value is 30.
<b>teamforge.repositoryroot</b>	Location where all Git repositories are stored physically. The default value is set to the value of the Gerrit configuration property <code>gerrit.basePath</code> , which is set to <code>/gitroot</code> by default.
<b>teamforge.maxFilesListedInTFCommitObject</b>	Restricts the number of entries in the SCM files list view for a particular TeamForge commit object. This is especially useful for repository initial commit objects as

	they could contain a thousand entries that get processed by TeamForge. The default value is 250.
<b>teamforge.notificationMaxSize</b>	Number of bytes in notification message that will be sent out by git-multimap (part of the <i>notification plugin</i> ). If message is larger than specified limit, it will be truncated. The default value is 25000.
<b>teamforge.notificationMaxPythonExecutors</b>	Number of Python processes used to create <i>git-multimap notification</i> . Each process will create one notification at a time. The default value is 2.
<b>teamforge.syncTeamForgeProjectHierarchy</b>	Turns the <i>Project Hierarchy</i> feature on. New Gerrit installs will have this value set to true, existing ones to false.
<b>teamforge.supportSiteWideRoles</b>	Enables TeamForge site-wide role support. New Gerrit installs will have this value set to true, existing ones to false. This feature requires at least TeamForge 8.0 (will be ignored before).
<b>teamforge.supportDefaultAccessPermissions</b>	Enables TeamForge Default Access Permission support. New Gerrit installs will have this value set to true, existing ones to false. This feature requires at least TeamForge 8.0 (will be ignored before).
<b>teamforge.commitProcessingTimeOut</b>	Maximum time allocated to process each Git commit to create a TeamForge commit object. If processing takes longer, processing of this commit is canceled, no corresponding TeamForge commit object will be created and the next commit will be processed. The default time is 15 min.
<b>teamforge.createTFProjectLinkedApps</b>	If enabled creates Project linked application with target to Gerrit Dashboard for that TeamForge project given project contains at least one Git repository. This feature requires at least TeamForge 8.0 (will be ignored before). The default value is true.
<b>teamforge.teamForgeMenuHeader</b>	Specifies the name of the menu that contains the links back to TeamForge user's Workspace and repositories list for a given TeamForge project. The default value is TeamForge.
<b>teamforge.ensureStreamEventsForRegisteredUsers</b>	If set to true, the <code>RegisteredUsers</code> group will have the <code>StreamEvents</code> global capability assigned during Gerrit startup. The default value is true.
<b>teamforge.ensureAdminRightsForSiteAdmins</b>	If set to true, the <code>TF: Site Admins</code> group will have <code>Administrative Server</code> global capability assigned during Gerrit startup. The default value is true.

## Replication configuration

This feature requires TeamForge 8.1 or later. These properties are ignored if you have TeamForge 8.0 or earlier.

**teamforge.replicationMode**

Sets the server mode (replication master or slave) of the Git integration server. This property is set by the TeamForge installer depending on the value specified in the `site-options.conf` file's `GERRIT_REPLICATION_MODE` token. Therefore, this property should not be edited manually within the `gerrit.config` file. The default value set by the TeamForge installer is master.

**Relocation master configuration****plugin.teamforge-replication.replicationDelay**

The delay (in seconds) between a push to the source repository and the actual replication attempt to the replica server. If further push activities happen between this delay, those will be bundled into the same replication attempt, avoiding bursts of replication attempts in case of repository mass updates. The default value is 15s and should not be set below 3s.

**plugin.teamforge-replication.threads**

The number of threads that are used to push changes for each replica server. The default value is 4.

**plugin.teamforge-replication.replicationRetry**

The maximum wait time before the next replication attempt is performed (upon previous connection failure). It is increased progressively (after each failure per mirror) starting with 1m to the power of 2 and up to the parameter value. For example, if the value is 5m, replication will be reattempted (considering that connection failure still occurs) after 1m then after 2m then after 4m and then after 5m and further attempts will be performed at 5m intervals. The default is 5m.

**plugin.teamforge-replication.sshConnectionTimeout**

The timeout duration for establishing SSH connections during a replication attempt or when an SSH command is performed. This prevents the SSH queue from being blocked while waiting to connect to a mirror that is not responding. The default value is 15s.

**plugin.teamforge-replication.sshCommandTimeout**

The timeout duration for replication SSH command execution (for example, project creation, HEAD change, and so on), after which the command fails. This prevents the SSH queue from being blocked while waiting to connect to a mirror that is not responding. The default is 30s.

**plugin.teamforge-replication.pushTimeout**

The timeout duration for a replication push (push time after SSH connection is established), after which the push fails. This prevents the SSH queue from being blocked while waiting to connect to a mirror that is not responding. The default is 30s.

**Replication mirror configuration****plugin.teamforge-slave.replicaId**

The replica ID of the replication slave created in TeamForge if `GERRIT_REPLICATION_MODE` is set

as slave. This property is set automatically by Gerrit upon start up and hence should not be edited manually.

**plugin.teamforge-slave.allowGroup**

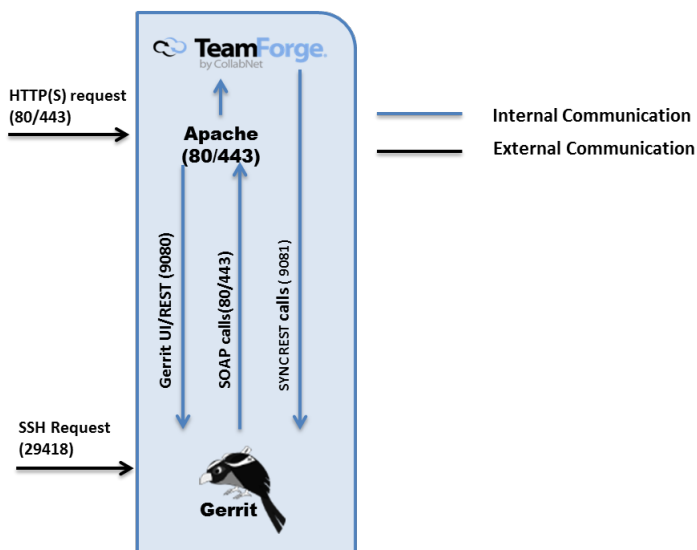
The group or groups that are allowed to push directly to the replication mirror. By default, only Administrator groups can do this.

**Illustrations on TeamForge-Gerrit communication**

The following illustrations help you understand the communication flow between TeamForge and Gerrit in a single host and distributed environments.

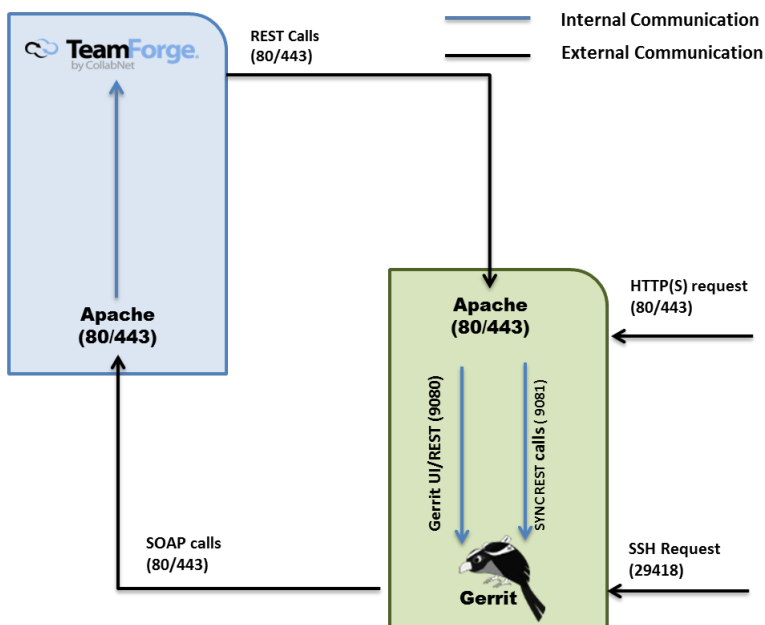
**TeamForge and Git/Gerrit on a single host**

TeamForge App co-hosted with Git Integration



**TeamForge and Git/Gerrit in a distributed two-server setup**


TeamForge App & Git Integration Hosted on Different Hosts





### TeamForge, Git/Gerrit and Replica Server in a three-server distributed setup

TeamForge App & Git Integration Hosted on a Different Host with Replication Server on Yet Another Host



## Integrate TeamForge 17.4 with other tools

---

TeamForge 17.4 supports integrations with third-party tools for versioning, reviewing, searching source code, binary repository management, and test management.

### Adapters: Integration using TeamForge EventQ

---

TeamForge EventQ integrates with different flavors of work item, SCM/commit, build, and review servers using "adapters".

TeamForge EventQ collects metadata about key lifecycle activities from external sources like work item trackers, SCM repositories, CI/build servers, and code review systems. To do so, TeamForge EventQ relies on "adapters" to collect the metadata from sources. TeamForge EventQ packages the following adapters:

Type of service	Product Description
SCM/Version Control	Git, Subversion
Builds/CI server	Jenkins, JetBrains TeamCity
Code Review	ReviewBoard, Atlassian Crucible
Work Item/Issue Tracker	CollabNet TeamForge, Atlassian JIRA


Download adapters from the [TeamForge EventQ Community](#). You can also create adapters as documented in [Extending TeamForge EventQ](#) on page 884.

### Source association keys

TeamForge EventQ uses source association keys to uniquely identify incoming data from activity sources that it monitors.

Adapters supply data to TeamForge EventQ through the message queue (MQ) layer. When a relevant activity occurs in source systems, TeamForge EventQ adapters generate a message with key metadata about the activity and place that message on the designated MQ server. When TeamForge EventQ receives the message, it must associate it to the proper source somehow. To do this, TeamForge EventQ supplies a unique "source association key" for each external source created. Adapters must supply this unique source association key with each message for the data to correlate with the appropriate project source.

Note that commit sources based on TeamForge project repositories do not need source associations keys since they are internal and are identified by other means. All other sources require source association keys.



### Locate a source association key

Source association key for existing activity sources are located on the corresponding source configuration page.

1. From the activity stream, click **Manage Sources**.
2. Select the desired source type toggle.  
TeamForge EventQ displays all existing sources.
3. Click **Edit** on the source whose key you wish to obtain.  
You see the configuration page for the selected source.
4. Locate the source association key for the selected source. On most browsers you can copy the key by clicking the small clipboard icon.

### JIRA adapter overview

The TeamForge EventQ JIRA integration brings associations and traceability to JIRA.

The TeamForge EventQ JIRA adapter enables JIRA as an alternative tracker for use with TeamForge by detecting and storing associations between JIRA issues and TeamForge managed SCM activities, like commits. Traceability is therefore possible between JIRA, TeamForge, and various eventqd third-party tools.

Inside TeamForge, JIRA issues are visible in activity streams and association visualizations. Within JIRA, all the SCM related activities for the particular JIRA issue are tracked and displayed under the **TeamForge** tab on the JIRA issue page.

MLB Defects / MD-1

## Add licensing constraints

Edit
Comment
Assign
More ▾
Start Progress
Done
Admin ▾

the Black Duck Software Knowledgebase. The GNU General Public License (GPL) family of licenses remains the most widely used license group for F/LOSS projects, with over 60% of all projects using one of the GPL licenses. This skewed distribution of license usage has prompted a community call for standardization on a limited set of known and recognized F/LOSS licenses, both to ensure a clear understanding of mutual obligations in case of mixing of code from different projects and to facilitate the process of managing contributions.

---


**Activity**

All
Comments
Work Log
History
Activity
**TeamForge**

**TeamForge Associations** Trace

Date	By	ID	Source	Summary
Jul 28 2015	nicholas_cn	75	WSeries SVN	[Md-1] Jira Task
Jul 17 2015	victor	73	WSeries SVN	[MD-1] Made room for new methods in App.java and added placeholder do...
Jul 17 2015	victor	72	WSeries SVN	[MD-1] Proposing changes to pom.xml. Need to restructure the way our dat...
Jul 16 2015	victor	71	WSeries SVN	[MD-1] adding new tests to handle the scenarios governed by ACs 1-3. Ens...
Jul 16 2015	victor	70	WSeries SVN	[MD-1] Cleaning up code comments.

The **TeamForge** tab provides a summary of associations and also details a full listing of associations. This is a listing of immediate associations, activities with direct relationships to the JIRA issue at hand. To explore the chain of associations, click the **Trace** button. The traceability chain that includes commits, builds, reviews, deploys, and other XDS-based integration is displayed. On the **Trace Associations** page, clicking any node lights up all the association paths to that node. In addition, a small pop-up appears with details about the node in question. You can use **SET AS TARGET** button to expand the associations from the selected point. Another function exists to **SEE** more details about the node, opening a new browser tab.



The TeamForge EventQ JIRA adapter is packaged as a JIRA "add-on". Once installed and configured, the add-on supplies issue tracker "work item" metadata to TeamForge EventQ. The JIRA adapter needs configuration on a per project basis in JIRA, such that JIRA (many) to TeamForge (one) project mappings are established. It is therefore required to configure a TeamForge project URL and a set of TeamForge credentials for each JIRA project. Note that once configured with the aforementioned URL and credentials, the JIRA adapter will create and manage the needed sources and source associations keys. For more about source association keys, see [Source association keys](#) on page 634.

### Install a JIRA adapter

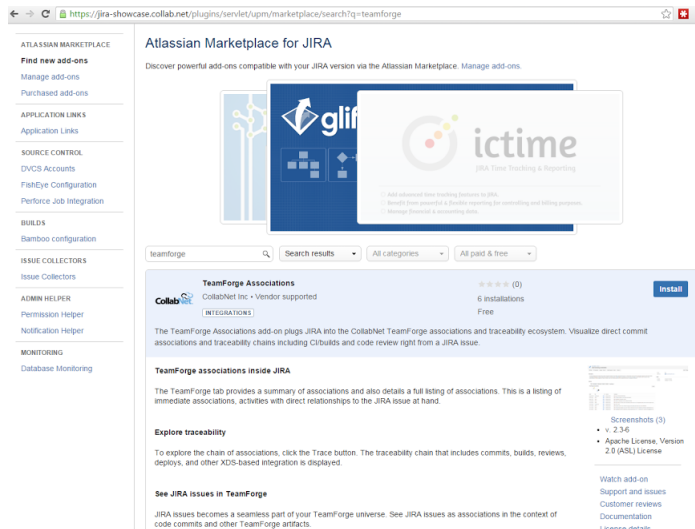
Install the JIRA adapter using the JIRA add-on Manager.

Use the JIRA adapter to notify EventQ of updates to JIRA issues and to visualize the associations between JIRA and other tools. The JIRA adapter must be installed once on each JIRA server you wish to connect to TeamForge EventQ. Install the adapter using the JIRA add-on Manager. Refer to the [Atlassian Marketplace](#) for supported versions.

### Installing the JIRA add-on

- As a privileged JIRA user, navigate to **Administration > Add-ons > Find new add-ons**
- Search the marketplace for the **TeamForge Associations** add-on for JIRA. Type 'TeamForge Associations' and search.

- Click **Install**.



### Upgrading the add-on

- New versions of the TeamForge Associations JIRA add-on will be visible in the **Manage add-ons** section (**Administration > Add-ons > Manage Add-ons**).

### Configure a JIRA adapter

Configure "TeamForge Associations" to notify TeamForge EventQ about work item activities.

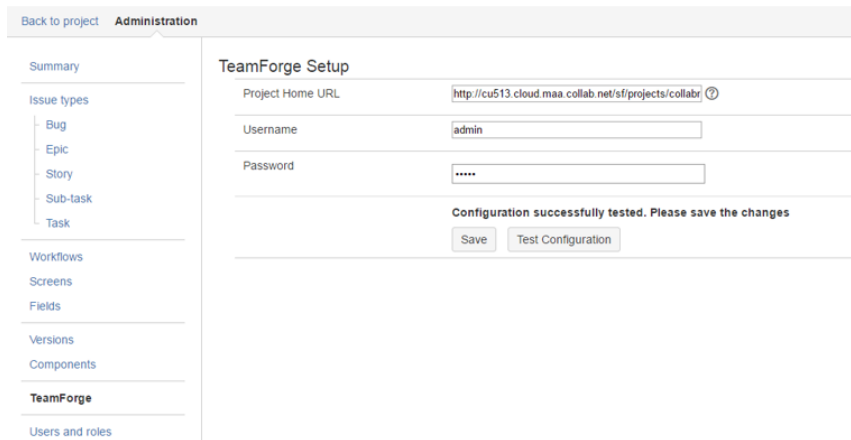
Before you start configuration, make sure that you have installed the TeamForge Associations add-on: verify that **TeamForge Associations** is present in the **User-installed Add-ons** list.

1. Configure the add-on from the **Project Administration** page to notify TeamForge-EventQ about the work item activities.
  - a) Navigate to **Administration > TeamForge**.

The **TeamForge Setup** page appears.

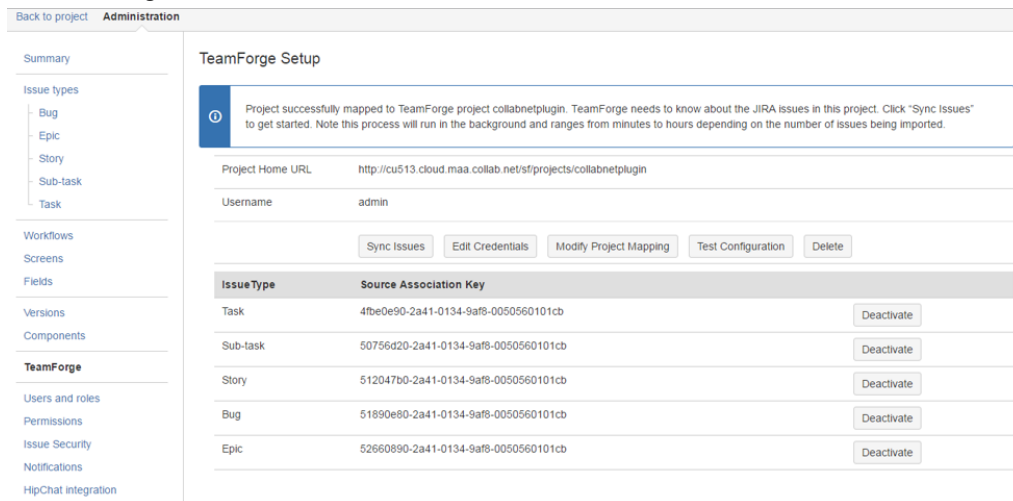
- b) Find the TeamForge project to associate with and enter its url in the **Project Home URL** field.  
Obtain the TeamForge Project Home URL by navigating to the target TeamForge project and clicking **Project Home**. Copy the URL from the browser.
- c) Using your valid TeamForge login credentials, click **Test Configuration**.

The user whose credentials are supplied must have API permissions in both TeamForge and EventQ, or is a Project Admin.



d) Click **Save** to save the configuration.

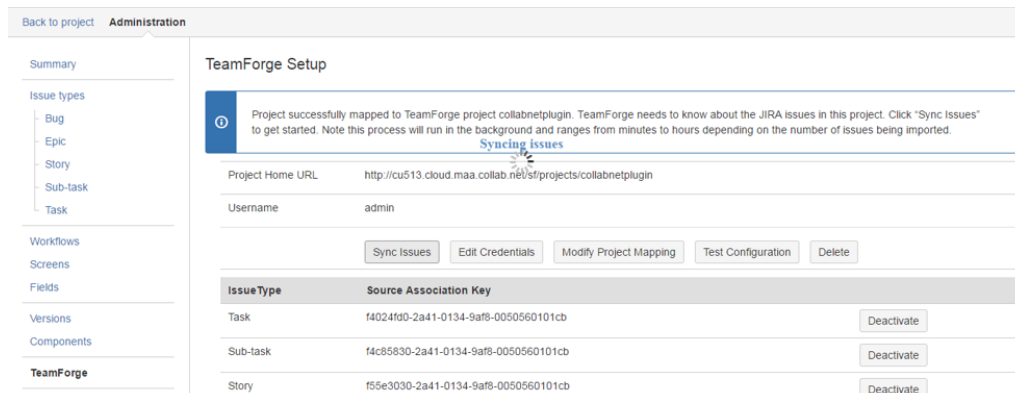
When you save your configuration, for each issue type in the JIRA project a corresponding source is created in the TeamForge EventQ.



### Sync Existing JIRA issues into TeamForge

2. To associate TeamForge objects to existing JIRA issues, click **Sync Issues**.

This step bootstraps JIRA issue data from the current project into the TeamForge EventQ data store. Note that this process may take several minutes (even hours) to complete, depending on the number of issues in the JIRA project.




### Test or modify configuration

3. Test the current configuration by clicking **Test Configuration**.

The system indicates whether the supplied URL and credentials work properly.


4. To modify the supplied credentials, click **Edit Credentials**.
  - a) Modify either the **Username** or **Password** fields.
  - b) Click **Test Configuration** to verify your settings.
  - c) Click **Save** once testing is satisfactory.
5. To modify the TeamForge project mapping, click **Modify Project Mapping**.

- a) Modify the **ProjectHome URL** field to reflect the desired TeamForge project.
- b) If required, update the credentials to match the TeamForge project.
- c) Click **Test Configuration** to verify that the supplied user has adequate permissions in the selected TeamForge project.
- d) Click **Save Mapping** once testing is satisfactory.

 **Note:** Upon successfully re-mapping the JIRA project, Run **Sync Issues** again to populate the TeamForge data store.

### Delete Project Mapping Configuration

6. Click **Delete** if you wish to completely disassociate the JIRA project from the configured TeamForge project.

 **Warning:** Deleting project mapping abandons all existing association data. The JIRA project can be mapped to another project subsequently, but existing associations are lost. Do this if you were testing the integration using a production JIRA project but now wish to remove any test association data.

### Associate JIRA Issues to Version Control Commits

Associations between JIRA® issues and TeamForge-aware version control commits can be created via commit message references. Commits to TeamForge project repositories and external repositories that have been configured as EventQ sources are also supported.

#### Creating associations

1. Using your desired terminal or IDE, instantiate a version control commit to your repository of choice.
2. In the commit message, make reference to the JIRA ID(s) to which you wish to associate the commit surrounded by square brackets.

#### Sample commit messages

```
"[DEMO-123] This commit message will result in an association between
JIRA issue DEMO-123 and this commit."
```

```
"[DEMO-123, DEMO-124] Here I'm associating two JIRA issues with project
identifiers DEMO."
```

#### Viewing Associations inside JIRA

3. Navigate to the desired JIRA issue details page.



4. Click the **TeamForge** tab.
5. A list of immediate associations appears. In other words, these objects are directly associated to the JIRA issue in question.
6. Click **Trace** to view the first three levels of the traceability chain.

**Note:**

- All activities are mapped chronologically.
- Lines indicate direct associations.
- Use the "+" icon to explore further levels of traceability.

**Troubleshooting****My associations do not appear. What should I do?**

Verify the following:


- Only repositories scoped to the configured TeamForge project will result in associations. That is, make sure the version control repository you're using is a part of the TeamForge project that has been mapped to the JIRA project in question.
- When a TeamForge project is first mapped to a JIRA project, the JIRA issues need to be synced into the TeamForge data store. Association will fail if you attempt to create an association between a commit and JIRA issue that has not yet been synced. Please refer to the Sync Issues functionality.
- Make sure the association syntax is correctly using square brackets. Association syntax is case sensitive.

**Remove TeamForge project mappings**

As a JIRA site administrator, you can disable active JIRA-TeamForge project mappings for one or more or all the projects (**Administration > Add-ons > TeamForge Associations**) if you want to prevent a JIRA server that's mirrored into a staging/testing environment from triggering events back to TeamForge (when there are changes to issues in JIRA projects that are mapped to TeamForge) thereby polluting the production TeamForge event data store.

Use this feature to disable all production JIRA-TeamForge mappings in stage environments and then test the TeamForge Associations add-on by creating a new mapping between a staging JIRA server and TeamForge server.

1. Log on to the JIRA server as a site administrator.
2. Select **JIRA Administration > Add-ons**.



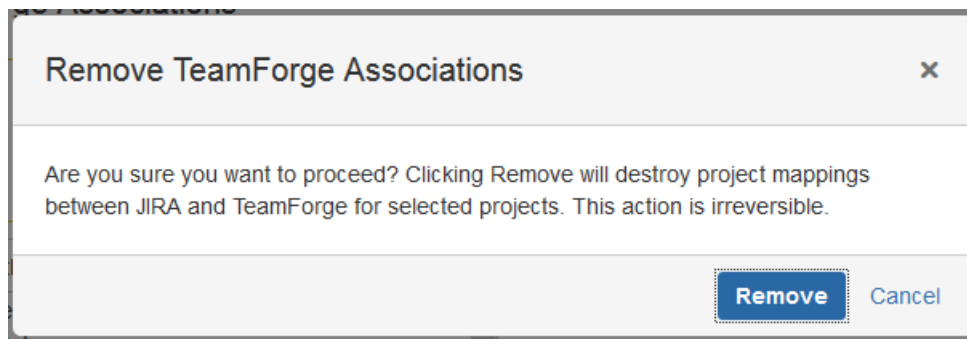
3. Select **TeamForge Associations > Manage**.

4. Select one or more projects from the list to remove the project mapping.

 **Tip:** You can select the **Select All** check box to select all the projects.

5. Click **Remove**.

A confirmation message is displayed.



6. Click **Remove** to delete the mapping.

### Troubleshooting the JIRA adapter

Solutions to common issues with the JIRA adapter:

**I just installed the EventQ JIRA adapter and made a commit against a JIRA ticket but the association doesn't appear.**

EventQ is an event system that associates activities across systems. If the JIRA issue you made a commit against hasn't been modified since installing the EventQ JIRA adapter, EventQ is unaware of the given issue and therefore won't process an association to it. The solution is to make a nominal change to the JIRA issue. This will automatically register the issue with EventQ and associations will start happening after that.

**Associations were working fine but now stopped working or appearing in JIRA.**

First, test the connection from the JIRA project to the TeamForge project. In JIRA: **Project administration > TeamForge > Test Configuration**

If you receive an error, make sure that the credentials used to map the JIRA project to the TeamForge project are still valid. TeamForge credentials can expire, user can be disabled or removed, which would invalidate the integration. The solution is to update the credentials in JIRA and re-test the configuration.

**"Test Configuration" fails with authentication errors** Try the following troubleshooting steps:

**I just re-mapped a JIRA project to a different TeamForge project and made a commit against a JIRA ticket but the association doesn't appear.**

- Check the username and password supplied. Note that this is a TeamForge username and password.
- Check that the Project Home URL is correct and is the intended Project Home URL.
- Check that the user supplied has proper permissions for the TeamForge project in question. The user whose credentials are supplied must have API permissions in both TeamForge and EventQ, or is a project admin.

EventQ is an event system that associates activities across systems. If the JIRA issue you made a commit against hasn't been modified since re-mapping the JIRA project to another TeamForge project, EventQ is unaware of the given issue and therefore won't process an association to it. The solution is to make a nominal change to the JIRA issue. This will automatically register the issue with EventQ and associations will start happening after that.

**I get an error trying to re-map a JIRA project to a different TeamForge project**

Check that the user supplied has proper permissions in both the original and re-mapped TeamForge projects. It is not sufficient for the user to have permission in just one of the two TeamForge projects in question as the re-mapping process touches endpoints in both projects.

## JIRA Cloud integration overview

The TeamForge Integration for JIRA® Webhooks collects issue data from your JIRA Server or JIRA Cloud purposes of associations, traceability, and event reporting.

### Overview

This integration relies on JIRA's webhooks mechanism to send issue data to TeamForge EventQ for purposes of associations, traceability, and event reporting. Being webhooks-based, there is no adapter or software to install; simply configure TeamForge and JIRA webhooks to communicate with each other and start creating associations, traceability, and events reports based on JIRA data within TeamForge. With no adapter to install, TeamForge Webhooks Integration for JIRA is ideal for use with JIRA Cloud.

### Viewing associations and traceability

Because there's no need to install plugins, there is no mechanism to visualize associations between TeamForge objects and JIRA issues from within JIRA. To visualize associations, please use the Source Code tool in TeamForge.

### Prerequisites

This integration is subject to limitations imposed by JIRA Webhooks and JIRA Cloud, see [Atlassian® Documentation](#) for specifics. In particular:

- TeamForge HTTP or HTTPS access must be configured on ports 80 or 443 respectively. JIRA Cloud does not support alternate port configurations for webhooks.
- If using HTTPS, TeamForge must be configured with a **valid SSL certificate signed by a certificate authority** to be an eligible JIRA Cloud webhooks consumer. Self-signed certificates are not supported by JIRA Cloud. This is a limitation *JIRA Cloud* imposes on webhooks consumers.
- TeamForge Integration for JIRA Webhooks is a point-to-point integration, therefore the following ports must be configured for access:


Port	From	To	Description
443/80	JIRA web server	TeamForge web server	JIRA sends webhooks payload to TeamForge

### Configuring the JIRA Cloud Webhooks

The TeamForge Integration for JIRA<sup>®</sup> Webhooks configuration instructions.

This integration documentation walks you through the creation of one or more webhooks triggers in JIRA. Each webhook should be scoped to a specific JIRA project issue type. It is therefore required that a separate webhook be set up for each issue type you wish to collect data from for purposes of associations and event reporting.

1. In TeamForge, navigate to the Project that houses the source code repositories you wish to associate to the JIRA issues in question.
2. Select **Project Home** > **Project Admin** > **Tools** and click **Add Tool**.
3. Select JIRA Cloud from the **Select Tool Type** drop-down list.



4. Type the **Display Name**.

The **Display Name** is used to differentiate your JIRA projects/instances associated to this TeamForge project.

5. Select the **Include Traceability** check box.

6. Type the **Source Name**.

This value differentiates the sources you define for this JIRA Project/instance. It is recommended that you name your new source to match the Issue Type in JIRA you're configuring currently.

7. Click **Generate Webhook**

You will be presented with a unique Webhook URL. This URL should only be used once, as it maps to a specific Project and Issue Type in JIRA.

8. Copy the entire Webhook URL and retain it for future steps.

9. Click **Save** to finish adding your tool (the new JIRA Project/instance).

10. In JIRA, as a privileged user navigate to **System Administration** > **Webhooks**.

11. Click **Create a Webhook** and give it a descriptive name and toggle "Status" to **Enabled**.

12. Paste the Webhook URL into **URL** field.

13. Under the **Events** section, check the following issue events: created, updated, deleted.

**Comment**

- created
- updated
- deleted

**Issue**

- created
- updated
- deleted
- worklog changed

**Worklog**

- created
- updated
- deleted

14. Under the **Events** section, input JQL to narrow the events triggering webhooks payloads to TeamForge as follows:

Supply a Project by entering: **project = [select project name]**

Supply an issue type by entering: **and issuetype = [select issue type]**

- 👉 **Tip:** Notice that a drop-down menu assists with project and issue type selection once the = symbol has been inputted, followed by an empty space.

## Events Issue related events

Events for issues and worklogs. You can specify a JQL query to send only matching issues.

✕ project = "MLB Defects" and issuetype = |

- Bug
- Improvement
- New Feature
- Sub-task
- Task
- Syntax Help

Project related events

A green check mark indicates you've properly configured the webhook.

## Events Issue related events

Events for issues and worklogs. You can specify a JQL query to send only matching issues.

✔ project = "MLB Defects" and issuetype = "New Feature"

### Syntax help

15. Click **Save**. All configuration is now complete for the issue type in question. You should see the newly configured webhook in the listing within the JIRA **System Administration** > **Webhooks** screen.
16. To confirm, create a new issue matching the issue type in question and you should see a new representative activity in the EventQ activity stream.  
Repeat the above steps for each desired issue type. It is recommended that you house all issue types relating to a single JIRA project under the same “Tool” in TeamForge, creating a separate “source” for each issue type.
1. Navigate to the newly created JIRA Cloud Tool in TeamForge: **Project Admin** > **Tools** > [Click the tool created in the steps above]
2. Notice that your existing source is housed here.
3. Click **Data source** + to create a second source to house the second JIRA issue type.
4. Continue with the steps as above.

The screenshot displays the configuration page for JIRA Cloud Webhooks. At the top, the breadcrumb navigation shows 'Project Admin' and a 'Jump to ID' field with the value 'proj1393'. The main configuration area includes the following details:

- DISPLAY NAME:** JIRA Cloud Webhook
- SHOW IN TOOLBAR:**
- INCLUDE TRACEABILITY:**
- DATA SOURCE:**
  - SOURCE NAME:** jira\_wh\_io\_story
  - SOURCE TYPE:** Work Item
  - WEBHOOK URL:** https://showcase.collab.net/orc/api/latest/projects/jchung\_test1/webhook/packets/jira\_cloud/884fa560-d35d-0133-fae8-020026b00001
- SOURCE NAME:** jira\_wh\_io\_task
- SOURCE TYPE:** Work Item
- WEBHOOK URL:** https://showcase.collab.net/orc/api/latest/projects/jchung\_test1/webhook/packets/jira\_cloud/fe9e1920-d35e-0133-fae9-020026b00001
- SOURCE NAME:** jira\_wh\_io\_bug
- SOURCE TYPE:** Work Item
- WEBHOOK URL:** https://showcase.collab.net/orc/api/latest/projects/jchung\_test1/webhook/packets/jira\_cloud/30ea08f0-d363-0133-faea-020026b00001

## Jenkins adapter overview

The TeamForge EventQ Jenkins adapter supplies build data from your Jenkins continuous integration server to TeamForge EventQ for purposes of metadata archival, traceability and event reporting.

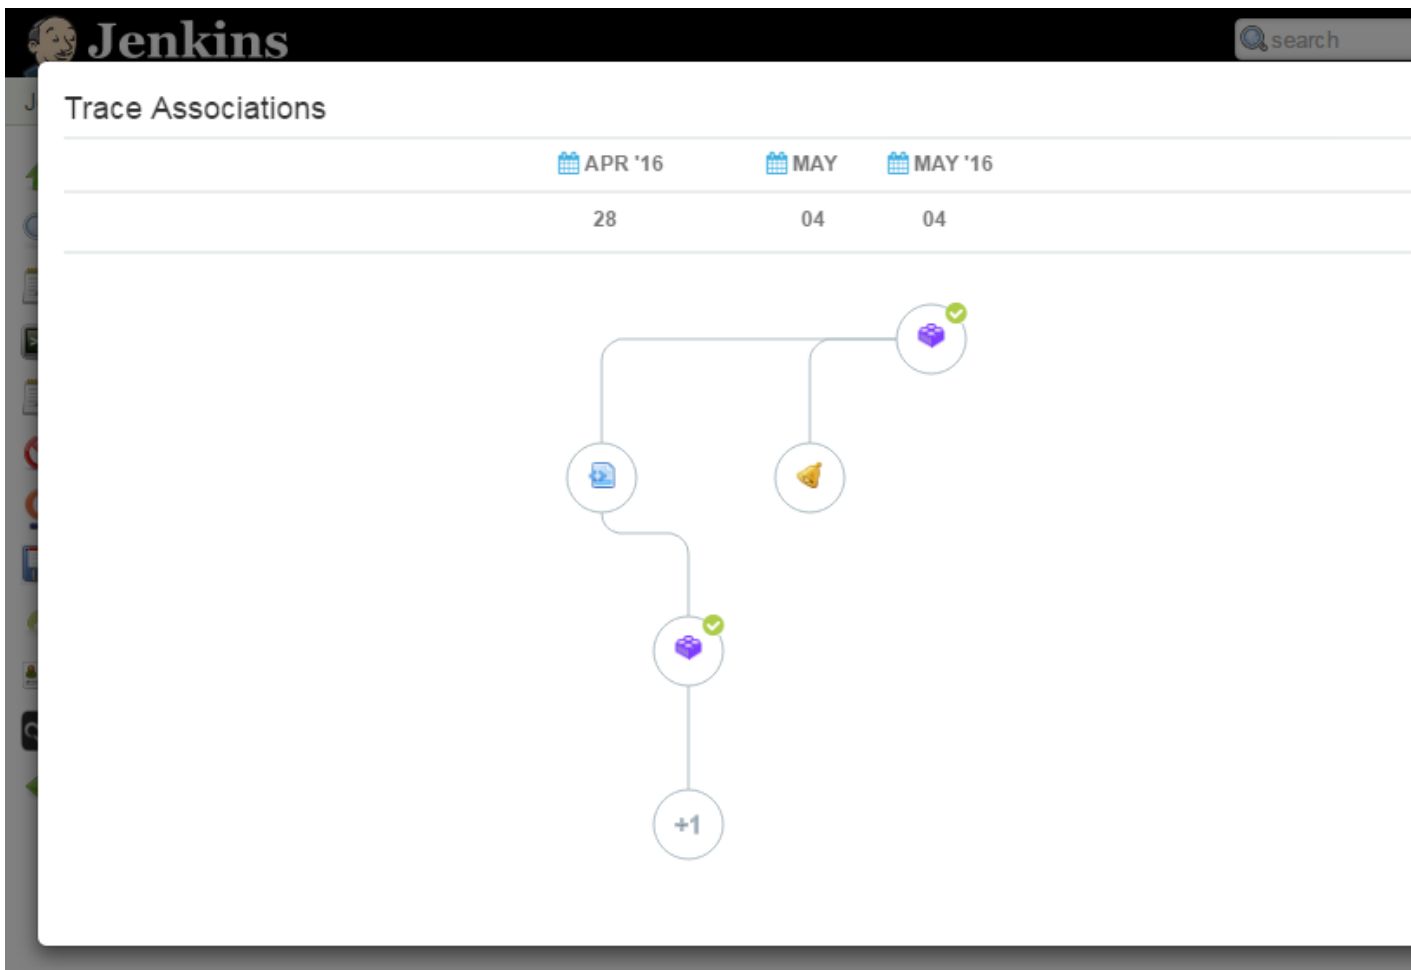
The TeamForge EventQ Jenkins integration brings associations and traceability to Jenkins. The adapter displays associations between Jenkins builds and other tools within the TeamForge landscape like version control commits and binary artifacts. Traceability is therefore possible between Jenkins build, TeamForge, and various orchestrated third-party tools.

Inside TeamForge, Jenkins build details are visible in activity streams and association visualizations. Within Jenkins, associations are visible from the “TeamForge Associations” tab.

The screenshot shows the Jenkins web interface. At the top, there is a search bar and a breadcrumb trail: Jenkins > Jenkins-Artifactory CI > #16 > TeamForge Associations. The left sidebar contains a list of actions: Back to Project, Status, Changes, Console Output, Edit Build Information, Delete Build, Artifactory Build Info, Tag this build, Redeploy Artifacts, See Fingerprints, **TeamForge Associations**, and Previous Build. The main content area is titled 'TeamForge Associations' and features a diagram showing a central node connected to two other nodes, each with a '1' in a black circle. Below the diagram is a table with the following data:

Date	By	ID	Source
04/May/2016	admin	<a href="#">com/example/simple/1.0...</a>	artifactory-
28/Apr/2016	admin	6	jenkins-arti

Clicking “Trace” provides a view of the build in the context of other associated activities. Use cases include root cause analysis and auditable traceability.



The TeamForge EventQ Jenkins adapter is packaged as a Jenkins plug-in. Once installed and configured, the plug-in supplies build results metadata for configured Jenkins jobs to TeamForge EventQ.

### Install a Jenkins adapter

A new CollabNet Plugin for Jenkins (v2.0) is now available that functionally subsumes both CollabNet Plugin v1.2.0 and EventQ Jenkins Adapter v2.0. Install this plugin using the web interface. Existing TeamForge-Jenkins integrations that use either CollabNet Plugin v1.2.0 (or earlier), or EventQ Jenkins Adapter v2.0 (or earlier), or both, must be migrated (using `migrate_jenkins_plugin.sh` script) after installing CollabNet Plugin.

Features of CollabNet Plugin v2.0 include:

- Notify TeamForge EventQ when builds complete. The CollabNet Plugin must be installed once on each Jenkins server you wish to connect to TeamForge EventQ.
- Authenticate users from TeamForge. If setup as the "Build & Test" application, it can even use Single Sign-On.
- Authorization from TeamForge, including the ability to set permissions in Jenkins based on roles in your TeamForge project.
- Upload the build log or workspace artifacts to the TeamForge Documents.
- Upload workspace artifacts to the TeamForge File Release System, as a post-build publishing task or as a build promotion task.
- Open/update/close TeamForge Tracker artifacts based on the Jenkins build status.
- Upload workspace artifacts to the Lab Management Project Build Library. (Requires [CollabNet Lab Management](#)).

[Click here](#) to know more about the requirements for installing CollabNet Plugin.

1. If you are integrating TeamForge and Jenkins for the first time:



- a) Log on to the Jenkins Server as a privileged Jenkins user, navigate to **Manage Jenkins > Manage Plugins > Available**.
  - b) Select the CollabNet Plugin v2.0 and install the plugin.
  - c) Restart your Jenkins server.
2. Existing TeamForge-Jenkins integrations that use CollabNet Plugin v1.2.0 (or earlier):
- a) Log on to the Jenkins Server as a privileged Jenkins user, navigate to **Manage Jenkins > Manage Plugins > Updates**.
  - b) Select the latest CollabNet Plugin and install the plugin.
  - c) *Download* the `migrate_jenkins_plugin.sh` script and save it to `<JENKINS_HOME_DIRECTORY>/jobs/`.
    -  **Tip:** Jenkins default home directory is `/var/lib/jenkins/`.
  - d) Change ownership of the `migrate_jenkins_plugin.sh` file.
 

```
chmod 755 migrate_jenkins_plugin.sh
```
  - e) Run the `migrate_jenkins_plugin.sh` script.
 

```
./migrate_jenkins_plugin.sh
```
3. Existing TeamForge-Jenkins integrations that use either EventQ Jenkins Adapter v2.0 (or earlier) or both CollabNet Plugin v1.2.0 (or earlier) and EventQ Jenkins Adapter v2.0 plugins:
- a) Log on to the Jenkins Server as a privileged Jenkins user, navigate to **Manage Jenkins > Manage Plugins > Installed**.
  - b) Select the EventQ Jenkins Adapter v2.0 and click **Uninstall**.
  - c) If you have Collabnet Plugin v1.2.0 (or earlier) installed already, select the **Updates** tab. Select the **Available** tab otherwise.
  - d) Select the CollabNet Plugin v2.0 and install the plugin.
  - e) *Download* the `migrate_jenkins_plugin.sh` script and save it to `<JENKINS_HOME_DIRECTORY>/jobs/`.
    -  **Tip:** Jenkins default home directory is `/var/lib/jenkins/`.
  - f) Change ownership of the `migrate_jenkins_plugin.sh` file.
 


```
chmod 755 migrate_jenkins_plugin.sh
```
  - g) Run the `migrate_jenkins_plugin.sh` script.
 

```
./migrate_jenkins_plugin.sh
```

### Configure a Jenkins adapter

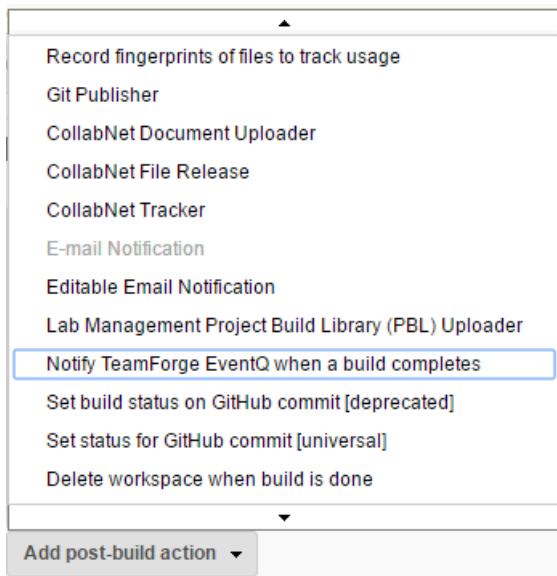
Configure the Jenkins adapter to notify TeamForge EventQ when builds complete.

Before you start the configuration, make sure that you have installed the Jenkins plug-in using the web interface. The TeamForge EventQ Jenkins adapter has a system configuration as well as job level configurations. The system configuration sets defaults to be used by subsequently configured job level configurations.

 **Important:** The source association key should be unique for each job configured.

### Configuring an individual Jenkins job

1. *Create a build source* and have the configuration details handy for the following steps.
2. As a privileged Jenkins user, locate the job you wish to report build data to TeamForge EventQ and navigate to its configuration page.
3. Add a post-build action to **Notify TeamForge EventQ when a build completes**.



- Using the values you obtained in Step 1, copy and paste the Queue Server value into the **Server URL** field.



**Important:** Include the "amqp" protocol and the port information.

- Fill in the **Server Username** and **Server Password** fields. Copy and paste the queue username and password values from the values obtained in Step 1.
- Populate the Build source association key field with the value obtained in Step 1. For more on association keys, see [Source association keys](#).

A screenshot of the 'Post-build Actions' configuration page. The title is 'Notify TeamForge EventQ when a build completes'. There is a red 'X' icon in the top right corner. Below the title, there is a checkbox for 'Optional TeamForge Association view' which is checked. To its right is an unchecked checkbox for 'Override the Global CollabNet Configuration'. Below these are four input fields: 'MQ URI' with the value 'amqp://10.2.17.233:5672', 'Server Username' with the value 'e9aef5c0-c4c7-0134-b993-0050560101a0', 'Server Password' with a masked password '.....', and 'Build source association key' with the value 'e9aef5c0-c4c7-0134-b993-0050560101a0'. Each field has a help icon to its right. At the bottom left, there is a button labeled 'Add post-build action' with a downward arrow. At the bottom center, there are two buttons: 'Save' and 'Apply'.

- By default, the **Optional TeamForge Association View** check box is selected. If required, you can override the global configuration. Select the **Override the Global CollabNet Configuration** check box and type the TeamForge URL and user credentials.

Notify TeamForge EventQ when a build completes

Optional TeamForge Association view

Override the Global CollabNet Configuration

CollabNet TeamForge URL:  ?

For example: https://forge.collab.net

Username:  ?

Password:  ?

MQ URI:  ?

Server Username:  ?

Server Password:  ?

Build source association key:  ?

Add post-build action ▾

8. Save the job configuration.

9. Run a build to test the new configuration and verify configuration. Information and errors will be reported to your Jenkins log and to the build console.

### Configuring global default Jenkins settings


Now that you've set up your first job level configuration, you can save some effort next time around by defining system values for the Server URL, Server Username and Server Password. Note that each job needs a unique Source Association Key, so that will still need to be configured per job.

10. Refer to Step 1 and obtain values for the Server URL, Server Username and Server Password.

11. As a privileged Jenkins user, navigate to **Manage Jenkins > Configure System** screen on the target Jenkins server.

12. On the Jenkins configuration page, find the TeamForge EventQ section.

a) Fill in the **Server URL** field. Copy and paste the Queue Server value into the **Server URL** field in Jenkins.

 **Important:** Include the "ampq" protocol and the port information.

b) Fill in the **Server Username** and **Server Password** fields.

13. Save the configuration.

### Troubleshooting the Jenkins adapter

You may encounter these problems when using the Jenkins adapter:

#### How do I get the queue server hostname and credentials? I need it to set up my adapter.

Adapters use a message queue server to transmit data to TeamForge EventQ. Adapters must therefore be configured with the queue server hostname, username, and password.

When you create a new "Source" in TeamForge EventQ, a unique set of queue server credentials is created for use with your adapters. To create a Source and obtain queue server hostname and credentials, see the [Manage build source](#) documentation.

Note that the same queue server hostname, username and password may be shared by multiple sources; however, the source association must be unique for each source.

#### My build is marked "Unstable"

Builds are marked "Unstable" if the Jenkins system configuration has not been saved and the Jenkins adapter cannot successfully communicate with TeamForge EventQ or the message queue. The console output for the build will include a message like this:

Build information NOT sent: plug-in needs a Jenkins URL

Workaround: As a privileged Jenkins user, navigate to **Manage Jenkins > Configure System** and populate and save the required configuration.

### Automated test results for my Jenkins builds are not showing in TeamForge EventQ

TeamForge EventQ relies on Jenkins to supply automated test results. Jenkins has a post-build action called "Publish JUnit test result report" which must be activated for test results to show up in TeamForge EventQ, regardless of whether JUnit is actually used.

Workaround: To enable this configuration, add the Publish JUnit test result report post-build action to the desired Jenkins job configuration.

### "java.lang.NoClassDefFoundError: javax.mail.internet.MimeMultipart" error is displayed while including 'Collabnet Security Realm' in Jenkins Security Configuration - Why?

If Jenkins runs along with any other applications such as TeamForge and JIRA on the same server, this error occurs. Hence, it is recommended to install Jenkins on a separate server. If you have installed Jenkins on the same server where other applications are installed as well, make sure that you host Jenkins on Tomcat to get rid of this issue. To know how to host Jenkins on Tomcat, check this wiki page: <https://wiki.jenkins-ci.org/display/JENKINS/Tomcat>.

## TeamCity adapter overview

The TeamForge EventQ TeamCity adapter supplies build data from your TeamCity continuous integration server to TeamForge EventQ.

The TeamForge EventQ TeamCity adapter is packaged as a TeamCity plug-in. Once installed and configured, the plug-in supplies build results metadata for selected TeamCity build configurations to TeamForge EventQ. The build results metadata is associated to the proper source using source association keys (see [Source association keys](#) on page 634). Configuration therefore requires that the MQ server URL and credentials be supplied, along with the appropriate source association key.

### Install a TeamCity adapter

Install the TeamCity adapter as a TeamCity plug-in

TeamForge EventQ currently supports TeamCity — 9.0 and 10.0.

The TeamCity adapter must be installed once on each TeamCity server you wish to connect to TeamForge EventQ. The TeamForge EventQ TeamCity adapter installs as a TeamCity plug-in. Installation of the adapter requires command line access to the TeamCity server.

1. Shutdown the TeamCity server. You must shut down the TeamCity server prior to installing the adapter. For more information on how to shut down and restart the TeamCity server, see the JetBrains TeamCity documentation.
2. Download the `EventQ-TeamCity-adapter-[version].zip` adapter file from the [TeamForge EventQ Community](#) web site.
3. Copy the TeamCity adapter zip archive to: `[TeamCity Data Directory]/plugins`

Do not unzip the file. The TeamCity Data Directory location can be set using the `TEAMCITY_DATA_PATH` environment variable. The default data directory is `$HOME/.BuildServer` on Linux and `C:\Users\<user_name>\.BuildServer` on Windows.


See [TeamCity plug-in Information](#) for more information on TeamCity plug-ins.

4. Start the TeamCity server and then configure the adapter.

## Configure a TeamCity adapter

Configure the TeamCity adapter to notify TeamForge EventQ when builds complete.

Before you start configuration, make sure that you have installed the TeamCity adapter. The TeamForge EventQ TeamCity adapter has a system configuration as well as build configuration level settings.


-  **Important:** When a build configuration is copied in TeamCity, the source association key is not carried over. You must add a corresponding build source in TeamForge Application Server, and enter the source association key while configuring the new build configuration.

### Do this on the TeamCity Server

1. Log on to the TeamCity Server.
2. Click **Administration** from the top navigation bar and click **Projects** from the **Project-related Settings** section.
3. If required, click the **Create Project** drop-down list and create a project. You can also choose to add build configuration for an existing project.
4. Select the project for which you wish to add build configuration and click **Create build configuration** drop-down list and create the build configuration.

### Do this on the TeamForge Application Server

5. Log on to the TeamForge Application Server.
6. Create a new build source or edit an existing one and associate the build source with the TeamCity build configuration.
  - a) Select a project and go to **Project Admin > Tools**.
  - b) Click **Add New Tool**.
  - c) Select **TeamCity** from the **Select Tool Type** drop-down list.
  - d) Enter the values for mandatory fields.
  - e) Select the **Include Traceability** check box as it is a mandatory field. A few more fields show up. Enter the build source details such as the source name, source type, and associated commit source and click **Show Configuration** to view the configuration details.

-  **Note:** Copy the configuration details such as the source association key, queue server, queue username, and queue password and keep it handy for configuring **TeamCity** as discussed later in this section.

### Do this on the TeamCity Server

7. Click **Administration** from the top navigation bar.
8. Click **TeamForge EventQ** from the **Server Administration** section.
9. Copy and paste the queue server, queue username and password values which you noted down earlier during EventQ build source configuration. Make sure that you include the "ampq" protocol and the port information as shown in the following illustration.

## DATA SOURCE +

SOURCE NAME: *	source2
SOURCE TYPE	Build/CI
SELECT ASSOCIATED COMMIT SOURCE	svnrepo
SOURCE ASSOCIATION KEY *	a2e50c30-0e0f-0135-f96a-0050560101bf
QUEUE SERVER *	amqp://cu493.cloud.maa.collab.net:5672
QUEUE USERNAME *	a2e50c30-0e0f-0135-f96a-0050560101bf
QUEUE PASSWORD *	44196ba717a2fd301c9df140c37fa68306a24ac6fb592367b078629da08

**TeamForge EventQ**  
Enter an AMQP server


AMQP Server URL:

Username:

Password:

**Save**

10. Click **Save** to save the system configuration.

 **Note:** System configuration steps need to be done only once per TeamCity server.

11. Click **Projects** on the **Administration** page. Select the build configuration for the desired project.

12. Select the **EventQ** tab.

13. Copy and paste the source association key value, which you noted down earlier during EventQ build source configuration. For more information about source association keys, see [Source association keys](#) on page 634.

Overview History Change Log Issue Log Statistics Compatible Agents 1 Pending Changes

**Report builds to TeamForge EventQ**

Association Key:  **Save**

*If no key is provided, then no activity will be reported for this build.*

14. Save the build configuration.

15. Run a build to test the new configuration and verify configuration.

## External Git repository integration overview

With a customized post-receive extension (hook) installed, Git servers external to TeamForge can notify TeamForge EventQ of Git "push" activity.

TeamForge project Git repositories are already wired to supply push activity data to TeamForge EventQ; no special configuration needed. This section instead details integration of Git repositories external to TeamForge.

The TeamForge EventQ Git adapter is packaged as a Git post-receive hook script. Once installed and configured, the hook script supplies push metadata for the selected repository to TeamForge EventQ. The push data is associated to the proper source using source association keys (see [Source association keys](#) on page 634). Configuration therefore

requires that the message queue server URL and credentials be supplied, along with the appropriate source association key.

### Install a Git adapter

Install a Git adapter as a post-receive hook on Mac OS X/Linux.

Minimum requirements:

- Git 1.8.3.1 or later.
- python 2.6. Version 3.0 and later are not compatible.

Each external Git repository you wish to monitor with TeamForge EventQ requires the installation of a customized post-receive hook script as follows.

### Install Git post-receive hooks

1. Contact the [CollabNet Support](#) to download the `eventq-git-adapter-[version].zip` adapter file and unpack it inside your repository's hooks directory. For example, if your repository's hooks directory is located at `/path/to/repo.git/hooks`:

```
cd /path/to/repo.git/hooks
```

Download the zip file to the current location, then:

```
unzip EventQ-git-adapter-1.0.zip
```

This will unpack the following into your hooks directory:

- Files: `LICENSE.txt` `README.txt` `post_receive_orc_publish_amqp.conf.sample` `post_receive_orc_publish_amqp.py`
- Directories: `eventq/` `pika/`

### Configuration

2. In a web browser, log into TeamForge EventQ as a privileged EventQ user and edit (or add) the relevant Commit source you wish to associate with this Git repository.

**Manage Sources > Toggle Commit > Add or Edit a Source > External repository**

3. On the Edit Source screen, locate the section titled "Adapter Configuration Information" and copy all values; these values are used later in the configuration process.
4. Switch back to your command line console and copy `post_receive_orc_publish_amqp.conf.sample` to `post_receive_orc_publish_amqp.conf`.
5. Edit `post_receive_orc_publish_amqp.conf` and modify options for your specific installation, for example:

#### Logging section

- `filename` — Full path to the desired log file, for example, `/tmp/post-commit.log`.
- `level` — The logging level, which is one of these: `debug`, `info`, `warning`, `error`, or `critical`.

#### RabbitMQ (amqp) section

- `host` — The hostname of the `amqp` server. Copy and paste the hostname from the Queue Server value from the Edit Source screen in TeamForge EventQ.
- `port` — The port number of the `amqp` server (default is 5672). Copy and paste the port number the Queue Server value from the Edit Source screen in TeamForge EventQ.

- `username`— Username used to access the `amqp` server. Copy and paste the queue username value from the Edit Source screen in TeamForge EventQ.
- `password` — Password used to access the `amqp` server. Copy and paste the queue password value from the Edit Source screen in TeamForge EventQ.
- `source_association_key` — key provided by TeamForge EventQ. Copy and paste the source association key value from the Edit Source screen in TeamForge EventQ. For more on association keys, see [Source association keys](#) on page 634.

### TeamForge EventQ section

#### Activate post-receive hook

6. Copy or create a `symlink` called "post-receive" linking to `post_receive_orc_publish_amqp.py`.

## External Subversion repository integration overview

With a customized post-commit extension (hook) installed, Subversion servers external to TeamForge can notify TeamForge EventQ of Subversion commit activity.

TeamForge project Subversion repositories are already wired to supply commit activity data to TeamForge EventQ; no special configuration needed. This section instead details integration of Subversion repositories external to TeamForge.

The TeamForge EventQ Subversion adapter is packaged as a Subversion post-commit hook script. Once installed and configured, the hook script supplies commit metadata for the selected repository to TeamForge EventQ. The commit data is associated to the proper source using source association keys (see [Source association keys](#) on page 634). Configuration therefore requires that the message queue server URL and credentials be supplied, along with the appropriate source association key.

### Install a Subversion adapter on Mac OS X/Linux

Install a Subversion adapter as a post-receive hook on Mac OS X/Linux.

Minimum requirements:

- Subversion 1.7 or later.
- python 2.6. Version 3.0 and later are not compatible.

Each external Subversion repository you wish to monitor with TeamForge EventQs requires the installation of a customized post-commit hook script as follows. CollabNet Subversion Edge 4.0 or later is also supported.

### Install Subversion post-commit hooks

1. Contact the [CollabNet Support](#) to download the `eventq-subversion-adapter-[version].zip` adapter file and unpack it inside your repository's hooks directory. For example, if your repository's hooks directory is located at `/path/to/svn/repos/my_repo/hooks`:

```
cd /path/to/svn/repos/my_repo/hooks
```

Download the zip file to the current location

```
unzip EventQ-svn-adapter-1.0.zip
```

This will unpack the following into your hooks directory:

- `LICENSE.txt`
- `README.txt`
- `configdefaults.py`
- `configwrapper.py`
- `eventq_commit_gate_hook.py`
- `orchestrator_commit_gate_hook_snippet.sh`



- pika/
- post-commit.bat.sample
- post\_commit\_orc\_publish\_amqp.conf.sample
- post\_commit\_orc\_publish\_amqp.py

### Configuration

2. In a web browser, log into TeamForge EventQ as a privileged EventQ user and edit (or add) the relevant Commit source you wish to associate with this Subversion repository.

#### Manage Sources > Toggle Commit > Add or Edit a Source > External repository

3. On the Edit Source screen, locate the section titled "Adapter Configuration Information" and copy all values; these values are used later in the configuration process.
4. Switch back to your command line console and copy `post_commit_orc_publish_amqp.conf.sample` to `post_commit_orc_publish_amqp.conf`.
5. Edit `post_commit_orc_publish_amqp.conf` and modify options for your specific installation, for example:

#### Logging section

- `filename` — Full path to the desired log file, for example, `/tmp/post-commit.log`.
- `level` — The logging level, which is one of these: `debug`, `info`, `warning`, `error`, or `critical`.

#### svnlook section

- `svnlook` — The path to your `svnlook` executable, for example, `/usr/bin/svnlook` or `/opt/subversion/bin/svnlook`.

#### RabbitMQ (amqp) section

- `host` — The hostname of the `amqp` server. Copy and paste the hostname from the Queue Server value from the Edit Source screen in TeamForge EventQ.
- `port` — The port number of the `amqp` server (default is 5672). Copy and paste the port number the Queue Server value from the Edit Source screen in TeamForge EventQ.
- `username` — Username used to access the `amqp` server. Copy and paste the queue username value from the Edit Source screen in TeamForge EventQ.
- `password` — Password used to access the `amqp` server. Copy and paste the queue password value from the Edit Source screen in TeamForge EventQ.

#### TeamForge EventQ section

- `source_association_key` — the key provided by TeamForge EventQ. Copy and paste the source association key value from the Edit Source screen in TeamForge EventQ. For more on association keys, see [Source association keys](#) on page 634.

### Activate post-commit hook

6. Copy or create a `symlink` called "post-commit" linking to `post_commit_orc_publish_amqp.py`.

### Install a Subversion adapter on Windows

Install a Subversion adapter as a post-receive hook on Windows.

Minimum requirements:

- Subversion 1.7 or later.

- python 2.6. Python version 3.0 and later are not compatible. Note, python is bundled with Subversion Edge.

This procedure assumes that you have installed CollabNet Subversion Edge, version 4.0 or higher, although other Windows distributions can be similarly configured. Each external Subversion repository you wish to monitor with TeamForge EventQs requires the installation of a customized post-commit hook script as follows.

### Install Subversion post-commit hooks

1. Contact the [CollabNet Support](#) to download the `eventq-subversion-adapter-[version].zip` adapter file and unpack it inside your repository's hooks directory. For example, if your repository's hooks directory is located at `C:\csvn\repositories\my_repo\hooks`:

```
cd C:\csvn\repositories\my_repo\hooks
```

```
# Download the zip file to the current location
```

```
unzip EventQ-svn-adapter-1.0.zip
```

This will unpack the following into your hooks directory:

- LICENSE.txt
- README.txt
- configdefaults.py
- configwrapper.py
- eventq\_commit\_gate\_hook.py
- orchestrator\_commit\_gate\_hook\_snippet.sh
- pika/
- post-commit.bat.sample
- post\_commit\_orc\_publish\_amqp.conf.sample
- post\_commit\_orc\_publish\_amqp.py

### Configuration


2. In a web browser, log into TeamForge EventQ as a privileged EventQ user and edit (or add) the relevant Commit source you wish to associate with this Subversion repository.

#### Manage Sources > Toggle Commit > Add or Edit a Source > External repository

3. On the Edit Source screen, locate the section titled "Adapter Configuration Information" and copy all values; these values are used later in the configuration process.
4. Copy `post_commit_orc_publish_amqp.conf.sample` to `post_commit_orc_publish_amqp.conf`.
5. Edit `post_commit_orc_publish_amqp.conf` and modify options for your specific installation, for example:

#### Logging section

- filename — Full path to the desired log file, for example, `C:\Temp\post-commit.log`.

 **Note:** Ensure that the log file exists, or create it, before proceeding to the next step.

- level — The logging level, which is one of these: `debug`, `info`, `warning`, `error`, or `critical`.

#### svnlook section

- svnlook — The path to your svnlook executable, for example, `C:\csvn\bin\svnlook.exe`.

#### RabbitMQ (amqp) section

- host — The hostname of the amqp server. Copy and paste the hostname from the Queue Server value from the Edit Source screen in TeamForge EventQ.

- port — The port number of the amqp server (default is 5672). Copy and paste the port number the Queue Server value from the Edit Source screen in TeamForge EventQ.
- username— Username used to access the amqp server. Copy and paste the queue username value from the Edit Source screen in TeamForge EventQ.
- password — Password used to access the amqp server. Copy and paste the queue password value from the Edit Source screen in TeamForge EventQ.

A RabbitMQ server is optionally set up during the installation process. The installation process sets a default password for the "eventq" user and stores it for reference in the following file on the RabbitMQ host (root access required):

```
/opt/collabnet/  
rabbitmq/.rabbit_passwords
```

### TeamForge EventQ section

- source\_association\_key — key provided by TeamForge EventQ. Copy and paste the source association key value from the Edit Source screen in TeamForge EventQ. For more on association keys, see [Source association keys](#) on page 634.

### Activate post-commit hook

6. Copy `post-commit.bat.sample` to `post-commit.bat`.
7. Edit `post-commit.bat` for your specific installation: set `PYTHON_EXE` to the full path to `python.exe` executable. For CollabNet Subversion Edge, this is typically `C:\csvn\Python25\python.exe`

### Troubleshooting the Subversion adapter

You may encounter these problems when using the Subversion adapter:

#### How do I get the queue server hostname and credentials? I need it to set up my adapter.

Adapters use a message queue server to transmit data to TeamForge EventQ. Adapters must therefore be configured with the queue server hostname, username, and password.

When you create a new "Source" in TeamForge EventQ, a unique set of queue server credentials is created for use with your adapters. To find the queue server hostname and credentials, log into TeamForge as a privileged EventQ user and navigate to: **EventQ (from the project navigation bar) > Manage Sources**. The **Sources Overview** page appears.

Select one of the adapters on the left hand side. The 'Source List' shows up on the right hand side.

From the list of sources, click the **EDIT** button of the desired source.

Locate the section named "Adapter Configuration Information".

Copy the queue server's hostname, username, and password and use them to configure your adapter.

Note that the same queue server hostname, username and password may be shared by multiple sources; however, the source association must be unique for each source.

**Commits fail after installing the adapter hook script**

After installing the hook script, I see an error complaining about svnlook:

```
Warning: post-commit hook failed
(exit code 1) with output:
svnlook: Expected FS format between
'1' and '4'; found format '6'
svnlook: Expected FS format between
'1' and '4'; found format '6'
Traceback (most recent call last):
File "/home/user/SVN_Edge/
csvn-4.0.0/data/repositories/
svn_edge/hooks/post-commit", line
488, in <module>
raise Exception("post-commit failed:
%s" % e)
Exception: post-commit failed: this
constructor takes no arguments"
```

**Workaround:** This problem is caused by having different versions of svnlook installed: typically one in /usr/bin and another bundled with Subversion Edge. To remedy, edit the EventQ Subversion Adapter configuration file `post_commit_orc-publish_amqp.conf` and modify the "svnlook" value to point to the binary packaged with Subversion Edge (example: `/home/user/SVN_Edge/current/bin`).

**Crucible adapter overview**

The TeamForge EventQ Crucible adapter enables communication between your Crucible server and TeamForge EventQ.

The TeamForge EventQ Crucible adapter is packaged as a Crucible add-on. Once installed and configured, the add-on supplies code review metadata for selected Crucible projects to TeamForge EventQ. The code review metadata is associated to the proper source using source association keys (see [Source association keys](#) on page 634). Configuration therefore requires that the MQ server URL and credentials be supplied, along with the appropriate source association key.

**Install a Crucible adapter**

Install the Crucible adapter using the Crucible add-on Manager.

TeamForge EventQ currently supports Crucible 3.0.1. Use the Crucible adapter to notify TeamForge EventQ when reviews are created or modified, or when files are added or removed from reviews. The Crucible adapter must be installed once on each Crucible server you wish to connect to TeamForge EventQ. Install the add-on by using the Crucible add-on Manager.

- Contact the [CollabNet Support](#) to download the `eventQ-Crucible-adapter-[version].jar` Crucible plug-in file.
- As a privileged Crucible user, navigate to **Administration > Manage Add-ons > Upload Add-on**
- In the **From my computer** field, select your `EventQ-Crucible-adapter-[version].jar` file, then click **Upload**.

You are now returned to the "Manage Add-ons" view and the "TeamForge EventQ Adapter" now be visible under "User-installed Add-ons".



**Note:** In the future you can upload new copies of the adapter and they will automatically upgrade the old adapter while retaining your old settings.

## Configure a Crucible add-on

Configure the "TeamForge EventQ Adapter" add-on to notify TeamForge EventQ about review changes.

Before you start configuration, make sure that you have installed the Crucible add-on: verify that **TeamForge EventQ Adapter** is present in the **User-installed Add-ons** list.

1. As a privileged EventQ user, edit (or add) the relevant Review source you wish to associate with this Crucible project.

### **Manage Sources > Toggle Review > Edit the Source**

2. On the Edit Source screen, locate the section titled "Adapter Configuration Information" and copy all values; these values are used later in the configuration process.
3. As a privileged Crucible user, navigate to **Administration > Manage Add-ons**
4. Click **TeamForge EventQ Adapter** in the **User-installed Add-ons** list.  
The TeamForge EventQ Adapter section displays new options, including "Configure".
5. Click **Configure**.  
You see a configuration form for each of your projects.
6. Find the project to associate with TeamForge EventQ and populate the following information.
  - a) Populate the Server URL field.  
Copy and paste the Queue Server value from the Edit Source screen in TeamForge EventQ. Please ensure to include the "ampq" protocol and the port information.
  - b) Populate the Server Username and Server Password fields.  
Copy and paste the queue username and password value from the Edit Source screen in TeamForge EventQ.
  - c) Populate the Source Association Key field.  
Copy and paste the source association key value from the Edit Source screen in TeamForge EventQ. For more on association keys, see [Source association keys](#) on page 634.
7. Click **Save**.  
Crucible now reports code review events for the configured project to TeamForge EventQ.

## Troubleshooting the Crucible adapter

You may encounter these problems when using the Crucible adapter:

### **How do I get the queue server hostname and credentials? I need it to set up my adapter.**

Adapters use a message queue server to transmit data to TeamForge EventQ. Adapters must therefore be configured with the queue server hostname, username, and password.

When you create a new "Source" in TeamForge EventQ, a unique set of queue server credentials is created for use with your adapters. To find the queue server hostname and credentials, log into TeamForge as a privileged EventQ user and navigate to: **Manage Sources > Toggle Review > Edit the Source**

Locate the section named "Adapter Configuration Information".

Copy the queue server's hostname, username, and password and use them to configure your adapter.

Note that the same queue server hostname, username and password may be shared by multiple sources; however, the source association must be unique for each source.

### **Where are communication errors reported?**

Communication problems with the TeamForge EventQ message queue server — such as inability to access the server URL or incorrect authentication credentials — will be reported in the Fisheye-Crucible error log,

usually found in `[FISHEYE_INST]/var/log/atlassian-fisheye-YYYY-MM-DD.log`

## Review Board adapter overview

The TeamForge EventQ Review Board adapter enables communication between your Review Board server and TeamForge EventQ.

The Review Board adapter is packaged as a Review Board extension, and once installed provides web-based configuration inside Review Board. TeamForge EventQ supports Review Board versions 1.7 and higher.

The adapter supplies code review metadata for selected repositories to TeamForge EventQ. Code review metadata is associated to the proper source using source association keys (see [Source association keys](#) on page 634). Configuration therefore requires that the MQ server URL and credentials be supplied, along with the appropriate source association key.

### Install a Review Board adapter on Mac OS X/Linux

Install a Review Board adapter as a Review Board extension to retrieve code review data from instances of Review Board 1.7 and higher.

Minimum requirements:

- Review Board 1.7 or higher
- python 2.6 or 2.7. Version 3.0 and later are not compatible.
- setuptools 2.1 or higher
- pip 1.5 - 6.0.6 (6.0.7 has issues and is not supported)

Please refer to the installation steps below, then continue to configuration.

#### 1. Upgrade pip

Upgrade pip if you do not have the minimum required version.

```
sudo pip install --upgrade pip
```

#### 2. Upgrade setuptools

Upgrade setuptools if you do not have the minimum required version.

```
sudo pip install --upgrade setuptools
```

#### 3. Install the EventQ adapter as a Review Board extension

If your server has internet access, use the command below to install the EventQ adapter as an extension. If upgrading, skip to the next step.

```
sudo pip install --use-wheel --no-index --find-links=http://eventq.collab.net/RHEL/6/x86_64/wheelhouse eventq_rb
```

If your server cannot access the internet do not use the command above. Instead, please see [Install a Review Board adapter without internet access](#) on page 662

#### 4. Upgrading the Review Board adapter

If upgrading (instead of installing), add the `-U` flag to upgrade to the latest version of the Review Board adapter

```
sudo pip install -U --use-wheel --no-index --find-links=http://eventq.collab.net/RHEL/6/x86_64/wheelhouse eventq_rb
```

#### 5. Restart Review Board's web server to register the newly installed EventQ adapter.

Proceed to [Configure Review Board adapter \(TeamForge integrated\)](#) on page 663.

### Install a Review Board adapter without internet access

Install a Review Board adapter without internet access.

The following instructions pertain to situations where your Review Board server does not have outbound access to the internet. Follow the instructions in [Install a Review Board adapter on Mac OS X/Linux](#) on page 662 until you are directed here to continue.

1. Contact the [CollabNet Support](#) to download the `eventq_rb-[version].zip` archive file.
2. Extract the archive file on the target server.

```
unzip eventq_rb-[version].zip
```

3. Install the EventQ adapter as a Review Board extension

Use the command below to install the EventQ adapter as an extension. If upgrading the Review Board adapter, skip to the next step.

```
sudo pip install --use-wheel --no-index --find-links=PATH_TO_WHEELHOUSE
eventq_rb
```

where `PATH_TO_WHEELHOUSE` is the path to the directory where you unzipped the adapter file. Suppose you unzipped the adapter file to the `/tmp` directory. In such a case, the `PATH_TO_WHEELHOUSE` is `/tmp`.

Example: `sudo pip install --use-wheel --no-index --find-links=/home/myusername/wheelhouse/ eventq_rb`

4. Upgrading the Review Board adapter

If upgrading (instead of installing), add the `-U` flag to upgrade to the latest version of the Review Board adapter

```
sudo pip install -U --use-wheel --no-index --find-links=PATH_TO_WHEELHOUSE
eventq_rb
```

5. Continue with installation documentation

Please continue with [Install a Review Board adapter on Mac OS X/Linux](#) on page 662 picking up with step 5 on page 662.

### Configure Review Board adapter (TeamForge integrated)

Configure a TeamForge integrated Review Board adapter to retrieve code review metadata.

Once the adapter has been installed successfully, configure Review Board by following these steps. The instructions below pertain to Review Board instances integrated with CollabNet TeamForge. See [Configure Review Board adapter \(standalone\)](#) on page 664 for configuration instructions when using a standalone version of Review Board.

#### Navigate to Review Board extension administration

1. Log into TeamForge with a site-admin user.
2. Click the **Admin** link in the header area.
3. Click the **Integrated Apps** tab
4. Click the radio button pertaining to Review Board
5. With the radio button selected, click **Administer**

You should now see Review Board administrative page.

6. Click **Extensions** on the menu bar to list your currently installed extensions.  
`eventq-rb` should be listed.
7. Click **Enable** to activate the EventQ extension.
8. Click **Configure** on the `eventq-rb` extension.

The `eventq-rb` configuration page will be displayed. From here you can select and configure the repository whose reviews you would like to push to an EventQ review source. To create a review source in EventQ see: [Manage review sources](#) on page 577.

#### Configure a Review Board repository

9. Select the repository name from the **Repository** dropdown menu.
10. Copy and paste the fields from the EventQ review source that you wish to associate to this repository

## 11. Click **Save**

Repeat steps [9](#) on page 663 through [11](#) on page 664 for each Review Board repository you wish to monitor in TeamForge EventQ.

### Configure Review Board adapter (standalone)

Configure a standalone instance of Review Board adapter to retrieve code review metadata.

Once the adapter has been installed successfully, configure Review Board by following these steps. The instructions below pertain to standalone Review Board instances. See [Configure Review Board adapter \(TeamForge integrated\)](#) on page 663 for configuration instructions when using an instance of Review Board integrated with CollabNet TeamForge.

### Navigate to Review Board extension administration

1. Log into your Review Board instance as a privileged user.
2. Click the **Admin** link to reach the administrative page.
3. Click **Extensions** on the menu bar to list your currently installed extensions.  
**eventq-rb** should be listed.
4. Click **Enable** to activate the EventQ extension.
5. Click **Configure** on the eventq-rb extension.

The eventq-rb configuration page will be displayed. From here you can select and configure the repository whose reviews you would like to push to an EventQ review source. To create a review source in EventQ see: [Manage review sources](#) on page 577.

### Configure a Review Board repository

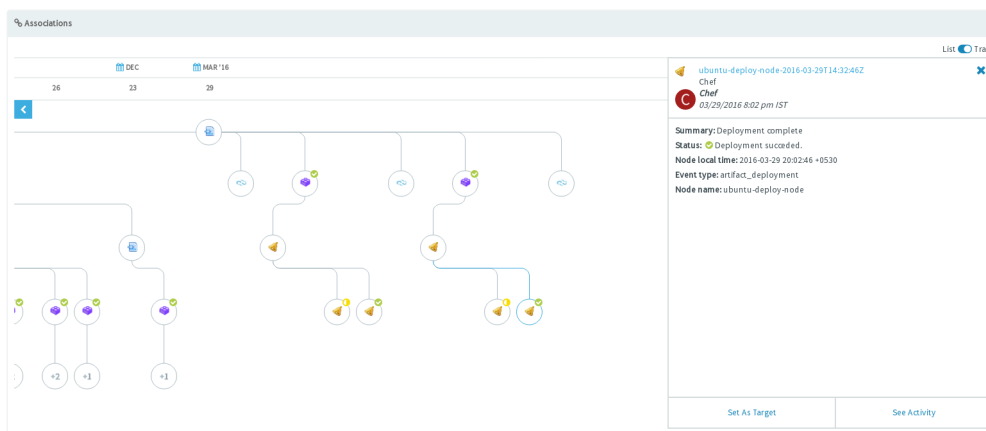
6. Select the repository name from the **Repository** dropdown menu.
7. Copy and paste the fields from the EventQ review source that you wish to associate to this repository
8. Click **Save**

Repeat steps [6](#) on page 664 through [8](#) on page 664 for each Review Board repository you wish to monitor in TeamForge EventQ.

## Chef adapter overview

The TeamForge EventQ Chef adapter supplies deployment data from your Chef environment to TeamForge EventQ.

This integration allows deployments done using Chef to be reported to EventQ and thus tracked for purposes of associations and traceability.




[\[Click to enlarge\]](#)


Chef deployments can be associated to either CI builds or binary artifact repository items, as illustrated below.

CI/Build to deployment work flow:





Binary artifact repository to deployment work flow:



This is achieved by linking the deployment done by Chef to either directly to the CI job that created the Build Artifact being deployed or to the build artifact residing in a Binary Repository. The possible use cases for Chef are nearly endless. The use cases covered in the instruction set below will be using Jenkins as the CI tool and Nexus as the Binary Repository.

Setting up Chef itself is outside the scope of this documentation and will not be covered. What is covered are only the steps involved in integrating an already existing Chef setup with a TeamForge EventQ instance. The setup works regardless of whether the Chef Server is hosted in the cloud or within your network, but you must consider the [firewall prerequisites](#).

### Install and configure Chef adapter (associate with Jenkins)

Install and configure the Chef adapter to associate with Jenkins.

Chef prerequisites	Jenkins host prerequisites
<ul style="list-style-type: none"> <li>• Chef Development Kit version: 0.10.0</li> <li>• chef-client version: 12.5.1</li> <li>• berks version: 4.0.1</li> <li>• kitchen version: 1.4.2</li> <li>• Server access to the host running Chef with sufficient privileges to modify Chef files</li> </ul>	<ul style="list-style-type: none"> <li>• Python version: 2.7</li> <li>• pychef version: 0.2.3, or later</li> <li>• "root" privileges on the host running Jenkins</li> </ul>

#### Assumptions

- Administrator with strong command of Chef
- A functioning Chef environment
- Chef is used in the actual deployment of software (for example, deploying a jar/war/ear or an rpm to the intended target)
- Jenkins CI job that creates the Binary artifact to be deployed by Chef is up and running, and *integrated with the CollabNet TeamForge instance*

### Add Chef deployment as a tool in TeamForge project

1. In TeamForge, go to the **Project Admin** page of the project for which you wish to integrate this Chef deployment scenario.
2. Select **Tools** from the left pane.
3. Click **Add Tool**.
4. Select Other for SELECT TOOL TYPE option.
5. Select the **Include Traceability** check box.

6. Select **Custom Activity Source** for "Source Type".
7. Select the desired Jenkins/Build source in the **Associated Source** field.
8. Finish any other confirmation, then click **SHOW CONFIGURATION** before saving the configuration. Copy all configuration parameters to a text file for use in later steps.

### Configuring Chef

9. Add the following code to your deployment cookbook's `metadata.rb` file.

```
depends 'teamforge-associations', '~> 0.1.0'
```

10. Add the following code to your `recipes/default.rb` file of your deployment cookbook.


```
include_recipe 'teamforge-associations::alldetails'
```

11. Edit your deployment recipe to notify the TeamForge EventQ reporting recipe when it is run. This can be done by adding the following code to your deployment resource.

```
notifies :enable, 'chef_handler[CustomHandler::AllDetails]'
```


12. Edit your `attributes/default.rb` file or any other attributes file that you are using for this recipe and add the following attributes.

```
default['teamforge-associations']['alldetails']['data_bag'] =
  'traceability_parameters'
```

 **Note:** This parameter denotes the data bag name on the Chef Server that has been written by Jenkins.

```
default['teamforge-associations']['alldetails']['ctf_url'] = <CTF URL>
```

```
default['teamforge-associations']['alldetails']['ctf_user'] = <CTF USER>
```

 **Note:** This user must have rights to get/post schema on EventQ.


```
default['teamforge-associations']['alldetails']['orc_server'] = '<ORC
QUEUE SERVER>'
```

```
default['teamforge-associations']['alldetails']['orc_user'] = '<QUEUE
USERNAME>'
```

```
default['teamforge-associations']['alldetails']['orc_src_assoc_key'] =
  '<SOURCE ASSOCIATION KEY>'
```

```
default['teamforge-associations']['alldetails']['orc_queue_name'] =
  'eventq.custom'
```

```
default['teamforge-associations']['passwords']['secret_path'] = '<PATH TO
SECRET KEY>'
```

 **Note:** This parameter contains the generic path on each node to the encryption key that has been used to encrypt the EventQ password that resides on the Chef server as an encrypted data bag.

13. Create a data bag named "passwords" on the Chef server and add two encrypted data bag items named 'ctf\_password' and 'orc\_password' containing the encrypted CTF password for the user set in `default['teamforge-associations']['alldetails']['ctf_user']` and the EventQ password for the user set up in `default['teamforge-associations']['alldetails']['orc_user']`

respectively. The passwords must be contained in a field named “password” in their respective data bags. Please see <http://docs.chef.io> for information on using data bags.

14. Once all changes are in place, upload the cookbook to the server. Any nodes that this cookbook is assigned to will automatically pick up these changes when they run `chef-client` next. If you wish to assign it to any new nodes, follow the usual procedure to bootstrap the node and assign this cookbook to its runlist.

### Configuring Jenkins

To create an association between a Jenkins build and a Chef deployment, the Chef deployment must be made aware of the “current” build number. To achieve this, Jenkins must be configured to create a data bag on the Chef server and write its build number to the data bag item. Please follow these instructions:

15. *Install Python* version 2 on the host running Jenkins, if it is not already present. Please note that the expected version is python2 (i.e., 2.7.x).
16. On the Jenkins server, *install the python library 'pychef'*, version 0.2.3 or later. Pychef is a tool needed for Jenkins to communicate with and operate on Chef. Installing Pychef can usually be achieved with the following commands:
  - `sudo easy_install pip`
  - `sudo pip install pychef`
17. Verify Pychef installation.
  - a) Enter `python` from the command line interface.
  - b) Enter `import chef` at the Python prompt.
  - c) If “import chef” is without any errors, Pychef installation is successful. Enter CTRL+D to exit the Python shell.
  - d) If there are errors about missing “libcrypto.so”, you must install the openssl developer package for your OS. Once installed, verify Pychef installation again.
18. Download the scripts to write/update from Jenkins to Chef DataBag from the [EventQ community site](#). The contents are a directory named “pychef” with two scripts inside it namely “data\_bag” and “set\_jenkins\_data\_bag.py”.
19. Login into the shell of the Jenkins server as the user “jenkins”.
20. Enter “cd” to make sure you are in the home directory.
21. Copy the “pychef” directory and it’s contents downloaded in [step 17](#) into the home directory.
22. Copy the “.chef” directory that is part of the starter kit into the pychef directory. Add a build step in Jenkins, after the actual build takes place, using option “execute command” and enter the following in the available text area:

```
cd ${HOME}/pychef
./data_bag -b ${BUILD_NUMBER}
```

That concludes the Chef integration with TeamForge via Jenkins.

Test the integration by running through an entire cycle to ensure EventQ receives the Chef deployment activity and also that the Chef deployment is associating properly to the configured Jenkins build.

### Install and configure Chef adapter (associate with Binary Artifact Repository)

Install and configure the Chef adapter to associate with Binary Artifact Repository.

#### Chef prerequisites

- Chef Development Kit version: 0.10.0
- chef-client version: 12.5.1
- berks version: 4.0.1
- kitchen version: 1.4.2
- Server access to the host running Chef with sufficient privileges to modify Chef files

#### Assumptions

- A functioning Chef environment

- Chef is used in the actual deployment of software (eg: deploying a jar/war/ear or an rpm to the intended target)
- The instance of Nexus hosting the Build Artifacts is up, running and integrated to the above said CollabNet TeamForge instance.

### Add custom source for Chef in TeamForge

1. In TeamForge, go to the **Project Admin** page of the project for which you wish to integrate this Chef deployment scenario.
2. Select **Tools** from the left pane.
3. Click **Add Tool**.
4. Select **Other** for SELECT TOOL TYPE option.
5. Select the **Include Traceability** check box.
6. Select **Custom Activity Source** for 'Source Type'.
7. Select the desired Nexus/Binary source in the **Associated Source** field.
8. Finish any other confirmation, then click **SHOW CONFIGURATION** before saving the configuration. Copy all configuration parameters to a text file for use in later steps.

### Configuring Chef

9. Add the following code to your deployment cookbook's `metadata.rb` file.

```
depends 'teamforge-associations', '~> 0.1.0'
```

10. Add the following code to your `recipes/default.rb` file of your deployment cookbook.

```
include_recipe 'teamforge-associations::alldetails'
```

11. Edit your deployment recipe to notify the TeamForge EventQ reporting recipe when it is run. This can be done by adding the following code to your deployment resource.

```
notifies :enable, 'chef_handler[CustomHandler::AllDetails]'
```

12. Edit your deployment recipe to notify the TeamForge EventQ reporting recipe when it is run. This can be done by adding the following code to your deployment resource.


```
default['teamforge-associations']['alldetails']['nexus_artifact_url'] =
  <NEXUS ARTIFACT URL>
```



**Note:** Where the URL is the REST API or the direct URL of the Nexus Artifact that is going to be deployed.

```
default['teamforge-associations']['alldetails']['ctf_url'] = <CTF URL>
```

```
default['teamforge-associations']['alldetails']['ctf_user'] = <CTF USER>
```

 **Note:** This user must have rights to get/post schema on EventQ.


```
default ['teamforge-associations'] ['alldetails'] ['orc_server'] = '<ORC
QUEUE SERVER>'
```

```
default ['teamforge-associations'] ['alldetails'] ['orc_user'] = '<QUEUE
USERNAME>'
```

```
default ['teamforge-associations'] ['alldetails'] ['orc_src_assoc_key'] =
'<SOURCE ASSOCIATION KEY>'
```

```
default ['teamforge-associations'] ['alldetails'] ['orc_queue_name'] =
'eventq.custom'
```

```
default ['teamforge-associations'] ['passwords'] ['secret_path'] = '<PATH TO
SECRET KEY>'
```

 **Note:** This parameter contains the generic path on each node to the encryption key that has been used to encrypt the EventQ password that resides on the Chef server as an encrypted data bag.

13. Create a data bag named “passwords” on the Chef server and add two encrypted data bag items named ‘ctf\_password’ and ‘orc\_password’ containing the encrypted CTF password for the user set in `default ['teamforge-associations'] ['alldetails'] ['ctf_user']` and the EventQ password for the user set up in `default ['teamforge-associations'] ['alldetails'] ['orc_user']` respectively. The passwords must be contained in a field named “password” in their respective data bags. Please refer to <http://docs.chef.io> for information on using data bags.
14. Once all changes are in place upload the cookbook to the server. Any nodes that this cookbook is assigned to will automatically pick up these changes when they run `chef-client` next. If you wish to assign it to any new nodes, follow the usual procedure to bootstrap the node and assign this cookbook to its runlist.

Test the integration by running through an entire cycle to ensure EventQ receives the Chef deployment activity and also that the Chef deployment is associating properly to the configured binary artifact.


## Install Nexus

---

The following instructions are for installing Nexus Open Source version 2.9.0 or later as a stand-alone server and integrating it with TeamForge 17.4.

Nexus comes bundled with a Jetty instance that listens to all configured IP addresses on a host (0.0.0.0) and runs on port 8081 by default.

Installing Nexus is straightforward. Unzip the Nexus bundle in a directory and start Nexus.

 **Note:** Though Nexus can be installed in Mac OS, CollabNet did not test Nexus integration in Mac OS.

1. Log on to the Nexus server.
2. Download the Nexus 2.9.0 or later zip file and unzip the content to a directory of your choice. See [Installing Nexus](#). You can find two directories, a directory that contains Nexus installation files and folders (hereinafter referred to as <nexus-install-directory>) and a Nexus work directory (hereinafter referred to as <nexus-work-directory>).

 **Important:**

- <sonatype-work> is the default Nexus work directory. As a notation, <nexus-work-directory> is used in place of <sonatype-work> in this document.
- Make sure you have full access permissions on all Nexus folders.

3. Open the command prompt and start Nexus.


**Linux:**

```
cd <nexus-install-directory>
Add the following token in the <nexus-install-directory>/conf/nexus.properties file
nexus.ui.keepAlive=false
./bin/nexus start
```

**Windows:**

```
cd <nexus-install-directory>
Add the following token in the <nexus-install-directory>\conf\nexus.properties file
nexus.ui.keepAlive=false
\bin\nexus start
```

Verify if Nexus is running by accessing the URL: <nexus host name>:port/nexus/index.html.

 **Tip:** Default port is 8081. In case you have multiple Nexus instances, modify the *application-port* token in /<nexus-install-directory>/conf/nexus.properties file.

4. Stop Nexus.

**Linux:**

```
./bin/nexus stop
```

**Windows:**

```
\bin\nexus stop
```

## Install the TeamForge-Nexus integration plugin

Once you have your Nexus server set up, install the TeamForge-Nexus integration plugin.

You must keep the following information handy before installing the TeamForge-Nexus integration plugin:

- **Nexus absolute path:** This is the path to the directory where you have Nexus installed. In other words, the path to the directory where you have your Nexus files unzipped. For example, /u1/nexus/nexus-2.11.4-01/.
- **TeamForge host URL:** The URL to access TeamForge. For example, http://ctf.cloud.collab.net.
- **TeamForge administrator user name and password.**
- **Nexus Application Name:** The name given to your Nexus integration. In other words, the name found in the Nexus integrated application configuration file. For example, see the pebble application's *pebble-app.xml* on page 776 for more information.
- **Nexus Application Prefix:** The prefix chosen for Nexus. In other words, the prefix found in the Nexus integrated application configuration file. For example, see the pebble application's *pebble-app.xml* on page 776 for more information.
- **Nexus URL:** The fully qualified Nexus URL.

1. Log on to the Nexus server.
2. Stop Nexus if it's running.

**Linux:**

```
./bin/nexus stop
```

**Windows:**

```
\bin\nexus stop
```

3. [Download](#) the CTF-Nexus-Integration-Plugin-2.1.1.zip file.
4. Unzip the CTF-Nexus-Integration-Plugin-2.1.1.zip file.

**Linux:**

```
cd <nexus-work-directory>/nexus/plugin-repository
unzip CTF-Nexus-Integration-Plugin-2.1.1.zip
```

**Windows:** Use a utility such as WinRAR.

5. Install the TeamForge-Nexus integration plug-in.

```
sudo java -jar <nexus-work-directory>/nexus/plugin-repository/CTF-Nexus-
Integration-Plugin-2.1.1/util/installer.jar -enable
```

Enter the Nexus absolute path and TeamForge host URL when prompted.

6. Start Nexus.

**Linux:**

```
cd <nexus-install-directory>
./bin/nexus start
```

**Windows:**

```
cd <nexus-install-directory>
\bin\nexus start
```

7. Once Nexus is up and running, upload the Nexus IAF descriptors to TeamForge.

```
java -jar <nexus-work-directory>/nexus/plugin-repository/CTF-Nexus-
Integration-Plugin-2.1.1/util/installer.jar -installxml
```

Enter the TeamForge Host URL, TeamForge admin user name and password, Nexus application name, Nexus application prefix and Nexus URL when prompted.

## Upgrade the TeamForge-Nexus integration plugin

CollabNet releases new versions of the TeamForge-Nexus integration plugin. It is recommended to upgrade your TeamForge-Nexus integration plugin whenever a new version is available.

1. Log on to the Nexus server.
2. Stop Nexus if it's running.

**Linux:**

```
./bin/nexus stop
```

**Windows:**


```
\bin\nexus stop
```

3. Modify the `ctf_nexus.properties` file.

```
vim sonatype-work/nexus/conf/ctf_nexus.properties
```

- a) Replace `TIME_TO_HOLD_CACHE=60` with the following two properties:

- `TIME_TO_HOLD_USER_CACHE=1800`
- `TIME_TO_HOLD_PERMISSION_CACHE=1200`

 **Note:** `TIME_TO_HOLD_USER_CACHE` and `TIME_TO_HOLD_PERMISSION_CACHE` are in seconds.

4. Back up the existing TeamForge-Nexus integration plugin directory: `<nexus-work-directory>/nexus/plugin-repository/CTF-Nexus-Integration-Plugin-1.0.2`.
5. [Download](#) the `CTF-Nexus-Integration-Plugin-2.1.1.zip` file.
6. Unzip the `CTF-Nexus-Integration-Plugin-2.1.1.zip` file you downloaded into the `<nexus-work-directory>/nexus/plugin-repository/ directory`.

**Linux:**

```
cd <nexus-work-directory>/nexus/plugin-repository
unzip CTF-Nexus-Integration-Plugin-2.1.1.zip
```

**Windows:** Use a utility such as WinRAR.

## 7. Start Nexus.

**Linux:**

```
cd <nexus-install-directory>
./bin/nexus start
```

**Windows:**

```
cd <nexus-install-directory>
\bin\nexus start
```

## TeamForge-Nexus release notes

Look here for information about the release highlights and bug fixes of various Nexus integration plugin versions supported by all the TeamForge releases starting from TeamForge 8.0.

Latest Nexus plugin: [CTF-Nexus-Integration-Plugin-2.1.1](#).

Download the latest Nexus plugin: [CTF-Nexus-Integration-Plugin-2.1.1](#)

To upgrade to Nexus plugin 2.1.1, check this page: [Upgrade the TeamForge-Nexus integration plugin](#) on page 671

.

For the feature enhancements and bug fixes in Nexus plugin 2.1.1, see [Nexus plugin 2.1.1 release notes](#).

### Previous releases

- [Nexus plugin 2.0.1 release notes](#)
- [Nexus plugin 2.0 release notes](#)

### Known Limitations with TeamForge-Nexus integration

Here's a list of known limitations of TeamForge-Nexus integration.

- While it is possible for users (all TeamForge users but site administrators) with 'create repository' permission to create binary repositories via the TeamForge **Binaries** application, such users cannot create repositories directly on the Nexus server (using the Nexus UI) as there is no TeamForge project mapping available in Nexus.
- A TeamForge user cannot be configured as an anonymous user in Nexus as TeamForge users are not available in the Nexus database.
- TeamForge broadcast messages and license notifications are not visible in Sonatype Nexus pages in TeamForge. This is due to a limitation with the TeamForge-Nexus integration plugin.
- When the TeamForge Project administrator changes the existing binary permissions for a user, the changes will not immediately take place due to the cache implementation in Nexus to improve the performance of session handling. Due to this limitation, when a user tries to create a new Nexus repository, he will get the permission denied error. Hence the user must wait for the changes to take effect. However, the user can view the changes immediately on a standalone Nexus application by restarting his session on it.

### Nexus plugin 2.0 release notes

This release notes provides the release highlights and bug fixes of Nexus 2.0 plugin.

Name of the plugin: CTF-Nexus-Integration-Plugin-2.0.



**Note:** You must upgrade to latest Nexus plugin version regardless of whichever TeamForge version you have installed or upgraded. To upgrade to latest Nexus version, check this page: [Upgrade the TeamForge-Nexus integration plugin](#) on page 671 .



## Release highlights of Nexus 2.0 plugin

- Usability of the Nexus installer (used to integrate Nexus with TeamForge) has been enhanced.
- Improved performance while loading Nexus UI.

## Bug fixes in Nexus 2.0 plugin

- When users log on to Nexus using CTF user credentials, the Nexus UI showed poor performance. This is fixed.
- Fixed the time delay due to the presence of Nexus plugin in Maven builds.

## Known Limitations with TeamForge-Nexus integration

Here's a list of known limitations of TeamForge-Nexus integration.

- While it is possible for users (all TeamForge users but site administrators) with 'create repository' permission to create binary repositories via the TeamForge **Binaries** application, such users cannot create repositories directly on the Nexus server (using the Nexus UI) as there is no TeamForge project mapping available in Nexus.
- A TeamForge user cannot be configured as an anonymous user in Nexus as TeamForge users are not available in the Nexus database.
- TeamForge broadcast messages and license notifications are not visible in Sonatype Nexus pages in TeamForge. This is due to a limitation with the TeamForge-Nexus integration plugin.
- When the TeamForge Project administrator changes the existing binary permissions for a user, the changes will not immediately take place due to the cache implementation in Nexus to improve the performance of session handling. Due to this limitation, when a user tries to create a new Nexus repository, he will get the permission denied error. Hence the user must wait for the changes to take effect. However, the user can view the changes immediately on a standalone Nexus application by restarting his session on it.

## Nexus plugin 2.0.1 release notes

This release notes provides the release highlights and bug fixes of Nexus 2.0.1 plugin.

Name of the plugin: CTF-Nexus-Integration-Plugin-2.0.1.

- 👉 **Note:** You must upgrade to latest Nexus plugin version regardless of whichever TeamForge version you have installed or upgraded. To upgrade to latest Nexus version, check this page: [Upgrade the TeamForge-Nexus integration plugin](#) on page 671 .

## Release highlights of Nexus 2.0.1 plugin

- Minor patch release.
- Enhanced the TeamForge top navigation bar supported by Angular JS.
- Available from TeamForge 16.7 release.

## Known Limitations with TeamForge-Nexus integration

Here's a list of known limitations of TeamForge-Nexus integration.

- While it is possible for users (all TeamForge users but site administrators) with 'create repository' permission to create binary repositories via the TeamForge **Binaries** application, such users cannot create repositories directly on the Nexus server (using the Nexus UI) as there is no TeamForge project mapping available in Nexus.
- A TeamForge user cannot be configured as an anonymous user in Nexus as TeamForge users are not available in the Nexus database.
- TeamForge broadcast messages and license notifications are not visible in Sonatype Nexus pages in TeamForge. This is due to a limitation with the TeamForge-Nexus integration plugin.
- When the TeamForge Project administrator changes the existing binary permissions for a user, the changes will not immediately take place due to the cache implementation in Nexus to improve the performance of session handling. Due to this limitation, when a user tries to create a new Nexus repository, he will get the permission denied error. Hence the user must wait for the changes to take effect. However, the user can view the changes immediately on a standalone Nexus application by restarting his session on it.

### Nexus plugin 2.1.1 release notes

This release notes provides the release highlights and bug fixes of Nexus 2.1.1 plugin.

Name of the plugin: CTF-Nexus-Integration-Plugin-2.1.1.

- 👉 **Note:** You must upgrade to latest Nexus plugin version regardless of whichever TeamForge version you have installed or upgraded. To upgrade to latest Nexus version, check this page: [Upgrade the TeamForge-Nexus integration plugin](#) on page 671 .

### Release highlights of Nexus 2.1.1 plugin

- User session management between Nexus application and TeamForge integrated with Nexus has been handled in a much better way so that communication between TeamForge and Nexus happens smoothly as expected. This enhancement is done for the Nexus versions 2.9 through 2.14.
- Improved performance while loading Nexus UI.
- Available from TeamForge 17.4 release.

### Bug fixes in Nexus 2.1.1 plugin

- Fixed the issue in which the '500: Internal Server Error' was shown instead of '401: Unauthorized Error' for any user activities in Nexus after a session timeout for the Nexus versions 2.9, 2.10, and 2.11. This fix is applicable for the Nexus versions 2.9 through 2.14.
- A '403: Forbidden' error was shown as Nexus authentication failed for users whose session has timed out due to the expiration of OAuth token. This issue is fixed now.

### Known Limitations with TeamForge-Nexus integration

Here's a list of known limitations of TeamForge-Nexus integration.

- While it is possible for users (all TeamForge users but site administrators) with 'create repository' permission to create binary repositories via the TeamForge **Binaries** application, such users cannot create repositories directly on the Nexus server (using the Nexus UI) as there is no TeamForge project mapping available in Nexus.
- A TeamForge user cannot be configured as an anonymous user in Nexus as TeamForge users are not available in the Nexus database.
- TeamForge broadcast messages and license notifications are not visible in Sonatype Nexus pages in TeamForge. This is due to a limitation with the TeamForge-Nexus integration plugin.
- When the TeamForge Project administrator changes the existing binary permissions for a user, the changes will not immediately take place due to the cache implementation in Nexus to improve the performance of session handling. Due to this limitation, when a user tries to create a new Nexus repository, he will get the permission denied error. Hence the user must wait for the changes to take effect. However, the user can view the changes immediately on a standalone Nexus application by restarting his session on it.

## Install Artifactory Pro

---

The following instructions are for installing Artifactory Pro 4.7 as a stand-alone server and integrating it with TeamForge 17.4.

Installing Artifactory Pro is straightforward. Unzip the Artifactory Pro bundle in a directory and start Artifactory Pro.

- 👉 **Note:** Though Artifactory Pro can be installed in Mac OS, CollabNet did not test Artifactory Pro integration in Mac OS.

1. Log on to the Artifactory Pro server.
2. [Download](#) the Artifactory Pro zip file.
3. To install Artifactory Pro manually, simply unzip the downloaded Artifactory Pro zip file to a location on your file system. This will be your \$ARTIFACTORY\_HOME location.

See [Installing Artifactory Pro](#).

#### 4. Run Artifactory Pro and verify the installation.

Artifactory Pro can be accessed using the following URL: `http://SERVER_DOMAIN:8081/artifactory`

For example, if you are testing on your local server you would use: `http://localhost:8081/artifactory`

## Install the TeamForge-Artifactory integration plugin

Once you have your Artifactory server set up, install the TeamForge-Artifactory integration plugin.

You must keep the following information handy before installing the TeamForge-Artifactory integration plugin:


- **TeamForge Host URL:** The URL to access TeamForge. For example, `http://ctf.cloud.collab.net`.
- **TeamForge administrator user name and password.**
- **Artifactory Application Name:** The name given to your Artifactory integration. In other words, the name found in the Artifactory integrated application configuration file. For example, see the pebble application's [pebble-app.xml](#) on page 776 for more information.
- **Artifactory Host URL:** The fully qualified Artifactory URL.

1. Log on to the Artifactory server.
2. [Download](#) the `Artifactory-teamforge-plugin.tar` file.
3. Unzip the `Artifactory-teamforge-plugin.tar` file.

#### Linux:

```
cd $ARTIFACTORY_HOME/etc/plugins/
tar -xf Artifactory-teamforge-plugin.tar
```

**Windows:** Use a utility such as WinRAR.

4.  **Note:** This step is required only on SSL-enabled sites (https sites).

Install the TeamForge-Artifactory integration plugin passing the `-m` flag to import TeamForge certificates.

```
sudo java -jar $ARTIFACTORY_HOME/etc/plugins/lib/installer-1.0.jar -m
```

Enter the TeamForge Host URL when prompted.

5. Start Artifactory.
6. Once Artifactory is up and running, upload the Artifactory IAF descriptors to TeamForge.

```
java -jar $ARTIFACTORY_HOME/etc/plugins/lib/installer-1.0.jar -i
```

Enter the TeamForge Host URL, TeamForge Administrator user name and password, Artifactory application name and Artifactory Host URL when prompted.

## TeamForge EventQ JIRA Integration: an overview

---

The TeamForge EventQ JIRA integration enhances further the extensible toolchain model of TeamForge by offering association and traceability between JIRA, TeamForge, and various orchestrated third-party tools. The TeamForge EventQ JIRA adapter enables JIRA as an alternative tracker for use with TeamForge by detecting and storing associations between JIRA issues and TeamForge managed SCM activities, like commits, builds and so on.

For example, the graphical representation of associations helps you perform a root cause analysis of an event, say, a build failure. Similarly, the traceability feature provides data useful for the purpose of audit or governance. For detailed information, see [JIRA adapter](#).

## TeamForge-TestLink integration

---

TestLink is a web-based Test Management system that supports all the various components and processes involved in a testing process. Using TestLink, you can create test specifications, execute test cases, create custom reports, generate test plan metrics and so on.



### Attention:

- CollabNet supports only the integration between TestLink and TeamForge. For TestLink support, contact TestLink directly. [Click here](#) for more information.
- TeamForge 16.7 and later releases support integration only with TestLink 1.9.15 and 1.9.16. TestLink 1.9.17 and later versions are not supported. If you are on earlier versions of TestLink, upgrade to one of the supported TestLink versions and integrate it with TeamForge. This integration does not provide backward compatibility (Data reliability and Migration) to older TeamForge-TestLink 1.9.11 integration that is based on TeamForge's Integrated Application Framework (IAF).

## TeamForge-TestLink integration: an overview

TeamForge 17.4 supports an integration with TestLink 1.9.15 and 1.9.16 to track and synchronize requirement and defect tracker artifacts between these two systems.



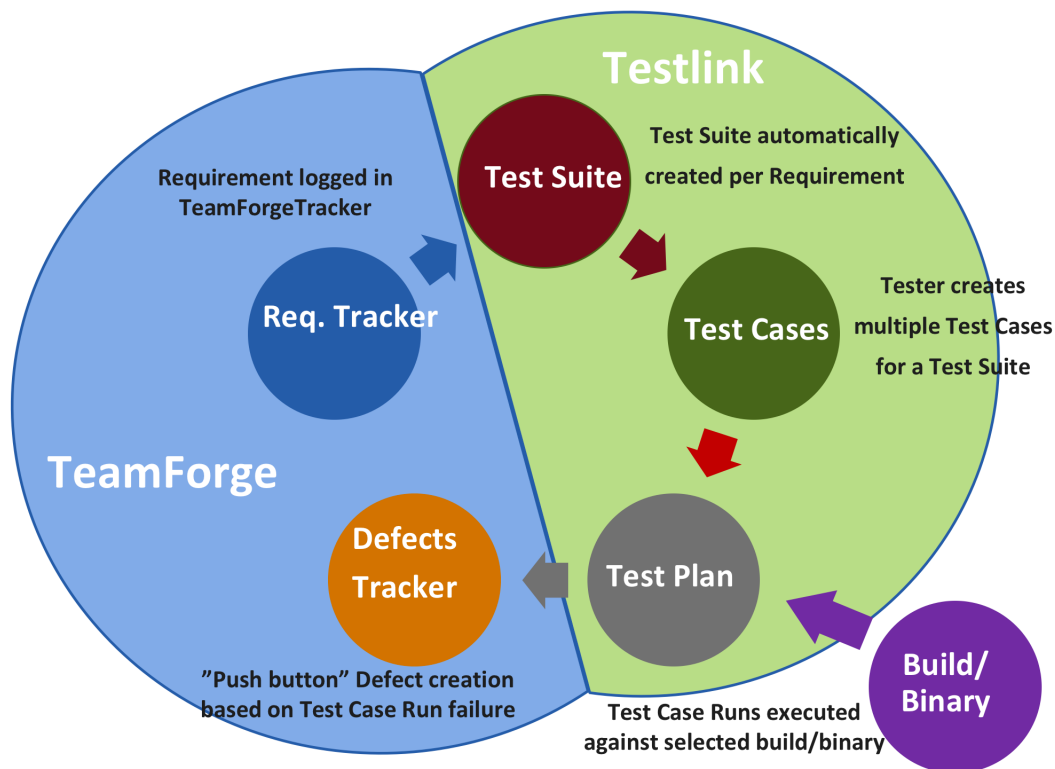
### Attention:

- CollabNet supports only the integration between TestLink and TeamForge. For TestLink support, contact TestLink directly. [Click here](#) for more information.
- TeamForge 16.7 and later releases support integration only with TestLink 1.9.15 and 1.9.16. TestLink 1.9.17 and later versions are not supported. If you are on earlier versions of TestLink, upgrade to one of the supported TestLink versions and integrate it with TeamForge. This integration does not provide backward compatibility (Data reliability and Migration) to older TeamForge-TestLink 1.9.11 integration that is based on TeamForge's Integrated Application Framework (IAF).

### Overview

The TeamForge-TestLink integration has two primary touch points:

- **Requirements** (TeamForge) to **Test Suites** (TestLink)
- **Test Case Run failures** (TestLink) to **Defects** (TeamForge)



Here is a sample workflow to illustrate these touch points:


#### Step 1: Map TeamForge and TestLink project

The first step is to establish a project mapping between the intended TeamForge and TestLink projects. During mapping, in TestLink you'll be able to specify the TeamForge "Requirements" and "Defects" trackers for your project. These trackers will be used throughout the workflow.

#### Step 2: Create Requirements and Test Suites

When a requirement artifact is created in the TeamForge tracker specified in Step 1, you will be given the opportunity to create a matching "Test Suite" in TestLink straight from the artifact creation screen in TeamForge.

**Important:** By default, only TeamForge users with "Site Administrator" privileges can create test suites. If you want other users to create test suites, you must create a site-wide role in TeamForge, grant "CREATE/VIEW" access to "All Projects" for the role and assign this role to users.



**Step 3: Populate Test Suite with Test Cases**


Within TestLink you can find and populate your Test Suite with Test Cases.

**Step 4: Run Test Cases against Builds**

Once the Test Cases have been defined, you can run them against desired builds, tracking results in TestLink. Test Case runs are associated to the build/binary being tested for traceability purposes.


**Step 5: Create Defects for failed Test Case runs**

If you encounter failing Test Case runs, you have the option to click a button inside TestLink to create a defect in TeamForge. This is a “push button” defect that opens a new TeamForge window pre-populated with Test Case run failure data. An association will be created between the failing Test Case run and the new Defect.



**Step 6: Visualize Traceability**

Once this cycle is complete, you can view the associations chain between Test Case runs and any associated defects, the Builds or Binaries they exercised, the Commits which originated the Build, and ultimately the Requirements that started it all.



## Installation and configuration

Instructions to integrate TeamForge 17.4 with TestLink 1.9.15 and 1.9.16.

Installation requirements for TeamForge-TestLink integration.

- [TeamForge 17.4 installation requirements](#)
- [TestLink installation requirements](#)



### Attention:

- CollabNet supports only the integration between TestLink and TeamForge. For TestLink support, contact TestLink directly. [Click here](#) for more information.
- TeamForge 16.7 and later releases support integration only with TestLink 1.9.15 and 1.9.16. TestLink 1.9.17 and later versions are not supported. If you are on earlier versions of TestLink, upgrade to one of the supported TestLink versions and integrate it with TeamForge. This integration does not provide backward compatibility (Data reliability and Migration) to older TeamForge-TestLink 1.9.11 integration that is based on TeamForge's Integrated Application Framework (IAF).



**Important:** By default, only TeamForge users with "Site Administrator" privileges can create test suites. If you want other users to create test suites, you must create a site-wide role in TeamForge, grant "CREATE/VIEW" access to "All Projects" for the role and assign this role to users.

### Do this on the TeamForge Application Server

1. Log on to the TeamForge Application Server.
2. A new version of TestLink integration jar (`collabnet-testlink-1.0.2.jar`) is available. If you have integrated TestLink with TeamForge using the `collabnet-testlink-1.0.jar` file in the past:
  - a) Go to **My Workspace > Admin**
  - b) Select **Projects > System Tools > Customizations**.
  - c) Select `collabnet-testlink-1.0.jar`.
  - d) Click **Delete**.

A confirmation message is displayed. Click **OK**.
3. [Download the `collabnet-testlink-1.0.2.jar` file](#).
4. Go to **My Workspace > Admin**
5. Select **Projects > System Tools > Customizations** and click **Create**.
6. Click **Choose File** and select the `collabnet-testlink-1.0.2.jar` file.
7. Click **Add**.

**Do this on the TestLink Server**

8. Log on to the TestLink Server, download and install TestLink.

 **Important:** Make sure that you have the following RPMs available during TestLink installation:


- php-xml
- php-mcrypt
- php-mbstring
- php-bcmath

9. Download the [collabnet-testlink-1.0.1.tar](#) file to the /tmp directory and untar the file to <testlink-installation-directory>/plugins/ directory.

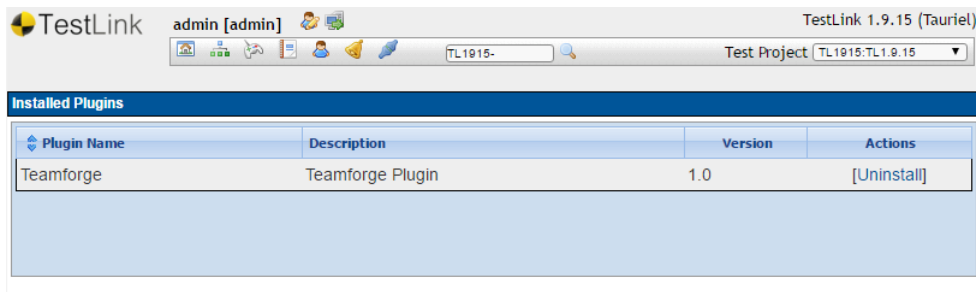
```
cd <testlink-installation-directory>/plugins/
tar -xvf /tmp/collabnet-testlink-1.0.1.tar
```

10. Log on to TestLink.

11. Click the **Plugins Management** icon from the toolbar.



12. Identify the TeamForge-TestLink integration plugin from the list of **Available Plugins** and click **Install**. The TeamForge-TestLink integration plugin is installed and shows up in the **Installed Plugins** section.



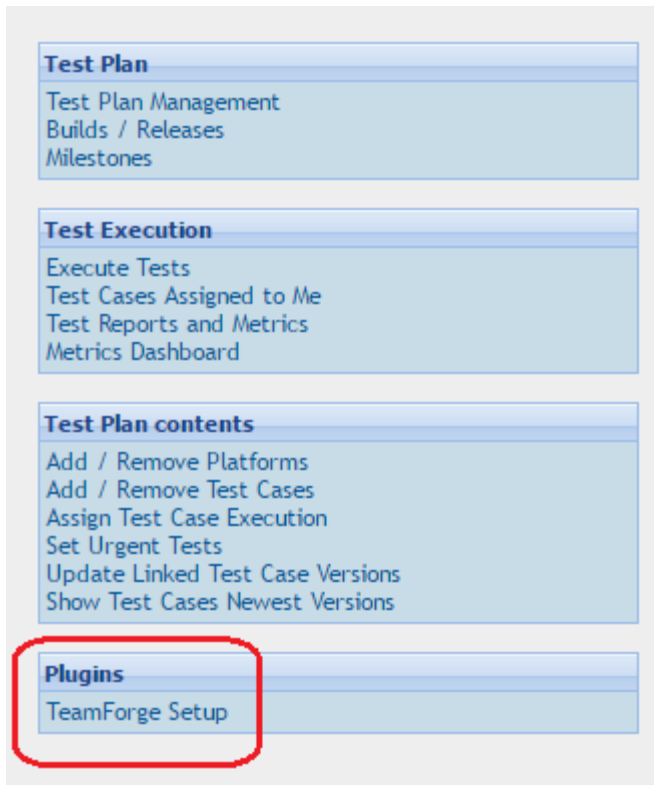
13. Click the **Home** icon from the toolbar.

14. Create a TestLink project.


- a) Click **Test Project Management** and click **Create**.
- b) Define project attributes such as the name, description, project prefix and so on. Click **Save**.

15. Go to TestLink **Home** and click the **TeamForge Setup** link.





16. Type the TeamForge **Project Home URL**, **Username**, **Password**, **Defect Tracker ID** and **Requirements Tracker ID** and click **Save**.

 **Important:** It is assumed that you have a TeamForge project, requirements tracker and defect tracker created already. If not, create them first and then perform this step of setting up TeamForge in TestLink. Have the requirements and defect tracker ID handy while setting up TeamForge in TestLink.

The screenshot shows the 'TeamForge Setup' configuration page in TestLink. The page header includes the TestLink logo and the user 'admin [admin]'. The main content area is titled 'TeamForge Setup' and contains the following fields:

Project Home URL	<input type="text" value="http://cu513.cloud.maa.collab.net/sf/projects/tl1_9_15"/>	
Username	<input type="text" value="admin"/>	
Password	<input type="password" value="....."/>	
Defect Tracker ID	<input type="text" value="tracker1033"/>	
Requirements Tracker ID	<input type="text" value="tracker1034"/>	

At the bottom of the form, there are three buttons: 'Save', 'Delete', and 'Cancel'. The 'Save' button is highlighted.


For more information on creating a TeamForge project and setting up trackers, see:

- [Create a new project](#) on page 343
- [Create a tracker](#) on page 27

Once you set up the TeamForge project mapping in TestLink, a new TestLink tool is created in TeamForge (**Project Home > Project Admin > Tools**).


**Do this on the TeamForge Application Server**

17. Go to **Project Home > Project Admin > Tools** and select the TestLink tool.



The **Edit Tool** page appears.

- Type a **Source Name**.
- Select Custom Activity Source from the **Source Type** drop-down list.
- If required, select an **Associated Source** from the drop-down list. For example, select Build.
- Change the icon, background color and add tags, if you want to.
- Click **Show Configuration** to view the configuration data.

 **Important:** Copy and keep the custom source information such as the Source Association Key, Queue Server, Queue Username and Queue Password handy to be used for Test Suite configuration later in the process.

- Click **Update**.


## Uninstall

Instructions to uninstall TeamForge-TestLink integration.

- Log on to TestLink.

### Delete the TeamForge project mapping


- Before uninstalling the TeamForge-TestLink integration plugin, it is a best practice but not mandatory to delete the TeamForge project mapping in TestLink.

 **Important:** While deleting the project mapping in TestLink turns TeamForge custom sources inactive, the TestLink tool in TeamForge should be removed manually, if required.


- Click **TeamForge Setup**.
- Click **Delete**. A confirmation message is displayed.
- Click **OK**.  
The TeamForge project mapping is deleted.

### Uninstall the TeamForge-TestLink integration plugin

- Click the **Plugins Management** icon from the toolbar.



- Identify the TeamForge-TestLink integration plugin from the list of **Installed Plugins** and click **Uninstall**.





The TeamForge-TestLink integration plugin is uninstalled.

### Define requirements, create test plans and test cases, and execute test cases

Define the requirements in TeamForge project's requirements tracker, create test plans and test cases, execute test cases and file defects for failed test runs.

- Log on to TeamForge, go to the TeamForge project and define the requirements in the requirements tracker. Create and submit an artifact.




 **Caution:** Test Suite titles in TestLink can be of 95 (or less) characters long. Therefore, you must limit requirement artifact titles to 95 characters or less while creating the requirement artifacts in TeamForge. Test Suite creation fails for TeamForge requirement artifacts with titles more than 95 characters.


- a) Select Create from the **TESTSUITE** drop-down list.
- b) The **TESTLINK URL** is auto-filled. Leave it as is.
- c) Type the **TESTLINK PROJECT NAME**.
- d) Click **Save**.

Log on to TestLink and verify if a new Test Suite has been created for the requirements you just created in TeamForge (go to **TestLink Home > Test Specification**).

- e) Go to **TestLink Home > Test Specification**, select the Test Suite from the tree and click **Actions**.

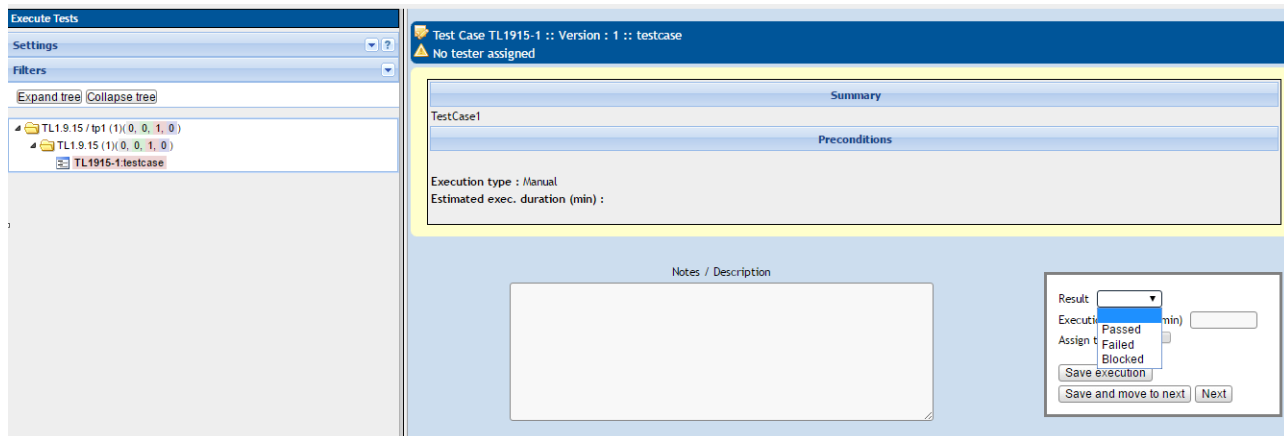


- f) Click the **Test Suite Operations > Edit** icon and update the custom source fields such as Source Association Key, Queue Server, Queue Username and Queue Password with the values you noted down earlier while *setting up the custom source in TeamForge*.




- g) Click **Save**.

2. Select the Test Suite you just configured from the tree, click **Actions** and create new test cases and builds. Once you have the test cases and builds created, add specific test cases to builds.
3. Go to **TestLink Home > Execute Tests**, select the test case from the tree, update the **Result** and **Save Execution**.




Notifications are added to the TestLink project's Activity Stream in TeamForge for your TestLink activities and events.




The Execution History section lists date-wise test case execution data with a **Submit Defect** link to file defects for "Failed" test runs. Clicking the **Submit Defect** link takes you to the TeamForge **Submit Artifact** page.

**Note:** Some of the test execution details such as the Tester's username, build name, time of execution and notes are automatically included in the artifact's "Description" field if and only if you have an active TeamForge session. Clicking **Submit Defect** just takes you to the TeamForge Login page otherwise.

You can also trace associations using the "TeamForge Associations" link for every test case execution.



4. Click **Submit Defect** and file a defect in the TeamForge.
5. Click the **TeamForge Associations** link and trace associations.



## Set up Pebble

Pebble is a blogging application that has been enhanced to support quick and easy integration with TeamForge.

### Install Pebble

To see what you can do with an integrated application in TeamForge, start by installing Pebble, a blogging application that you can configure to work as part of your TeamForge site.

1. Get the Pebble installer package from [open.collab.net](http://open.collab.net) and unzip it.
2. Modify these values in the `installer/install.conf` file to suit your installation environment.

#### Option

**pebble.base.dir**

#### Description

Path where you want Pebble to be installed on this host.

Option	Description
<b>ctf.baseurl</b>	Absolute URL of the TeamForge site that you want to associate to, such as <code>https://my.ctf.instance/</code>
<b>tomcat.port</b>	Pebble runs on Tomcat. This token indicates which port you want Tomcat to be running on. Make sure there are no other services running on that port.
<b>domain</b>	The base url to be used for Pebble, such as <code>my.ctf.instance</code> . (You don't need <code>http:</code> or <code>https:</code> here.)
<b>timezone</b>	The time zone Pebble will use to timestamp blog entries.
<b>java_home</b>	Path to a JDK 1.6.x instance.
<b>protocol</b>	<code>http</code> if SSL is not being used; <code>https</code> if SSL is being used.
<b>data.dir</b>	Path in the file system where Pebble blogs will be stored.


### 3. Run the installer.

```
sudo python install.py -i -r
```

### 4. Set up the initial blog data.

```
sudo python bootstrap-data.py
```

This is known as "bootstrapping" the application.

 **Tip:** You can bootstrap again if you want to start from scratch, but any existing blogs will be deleted if you do.

### 5. Restart the Pebble application.

```
service pebbled stop
service pebbled start
```

You should now have a working Pebble instance ready to work with TeamForge. The installer has created two configuration files: `installer/conf/pebble-app.xml` and `installer/conf/pebble-dep.xml`. See [Integrate Pebble into your TeamForge site](#) on page 687 for how to use them.

If you have installed Pebble with SSL, restart the TeamForge server before integrating Pebble with TeamForge. You can restart the TeamForge server using `/opt/collabnet/teamforge/bin/teamforge restart`

## Integrate Pebble into your TeamForge site

When the sample Pebble blogging application has been installed on your site, you can make it available for projects on your TeamForge site.

Pebble must be installed and configured before you can integrate it into your TeamForge site. See [Install Pebble](#) on page 686.

When you have integrated Pebble, projects on your site can add Pebble to their set of collaboration tools. The blogs they create will share many of the core TeamForge features, such as authorization, authentication, go-urls, association, linkification, templating, Project Pages components, and source code management support.

1. Log into TeamForge as an admin user.
2. Go to **My Workspace > Admin**.
3. Click **INTEGRATED APPS** from the **Projects** menu.
4. Click **Create**.

5. Use the **Browse** window to find the two configuration files that enable the Pebble application to work as a part of TeamForge:
  - `pebble-app.xml` (Application configuration file): Contains the text strings for the Pebble user interface.
  - `pebble-dep.xml` (Deployment configuration file): Contains the data that Pebble needs to interact with TeamForge.

Click **Next**.

6. On the **Preview** screen, review the parameters you set in the configuration files.



**Note:** You may have to revise one or more values to ensure they are valid.

7. Click **Save**.

The Pebble application is now available for all projects on your site. You can direct project administrators to [the project admin help](#) for instructions on adding it to their own project toolbars.



**Note:** You may need to adjust your site's look and feel to support your integrated application. See [the site admin help](#) for details.

## Add Pebble to a TeamForge project

When the site administrator has made the sample Pebble blogging application available, project administrators can add it as one of their project tools.

When you add Pebble to your project, it works just like the other TeamForge tools, with features such as authorization, authentication, go-urls, association, linkification, templating, Project Page components, and source code management support.

For this example, we'll call your project "testproject" and we'll assume you have Project Admin rights in that project.

1. Click **PROJECT ADMIN** from the **Project Home** menu.
2. Click **Tools** to see the list of integrated applications available in the site.
3. Click **Add Tool**.
4. Select **Pebble Blog** from the list of tools displayed.

Set the values that make sense for your TeamForge installation, then click **Save**.

<b>Option</b>	<b>Description</b>
<b>Prefix</b>	Specify a unique alphanumeric string that will identify this tool throughout the site. For example, suppose you set your prefix to ZZ. Entering <code>ZZ_20100612</code> in the <b>Jump to ID</b> box will get you all blog posts on June 12, 2010 in the "testproject" project. Another may also add Pebble as a tool, with a prefix of YY. Jumping to <code>YY_20100101</code> will get all the blog posts on January 1, 2010 from that project.
	<b>Note:</b> The tool prefix cannot be changed after you have set it. Your site administrator may have elected to make other parameters permanent too, such as the blog name.
<b>Blog Name</b>	This will appear as a banner in the blog for this project. You can enter any name.
<b>Blog Description</b>	This will appear under the Blog Name in the blog for this project. You can enter any description.
<b>Rich Text Editor Enabled</b>	Select this if your users want word-processor-like writing support when they create blog posts.



Option	Description
<b>Number of Entries in Blog Home</b>	You can choose to have 3, 5, 7 or 9 entries in the blog homepage.

If a **Pebble Blog** button appears along with the prefix as a tooltip when you mouse over the button, your job is done.

## Authenticate your integrated application with TeamForge

---

To help third party developers write integrated applications, CollabNet provides an SDK. At the heart of the SDK is the `IntegratedAppSupport` helper class. You use `IntegratedAppSupport` to authenticate integrated application requests with TeamForge, and provide context information for subsequent processing of forms and links within the application.

You can download the SDK from [here](#).

You must call the `IntegratedAppSupport` class for every request made by the integrated application. This class takes `HttpServletRequest` and `HttpServletResponse` for each request and determines whether the user is already authenticated. It also provides project- and user-related information that can be used throughout the request.

 **Tip:** It is a good idea to store this as a `ThreadLocal` so that it can be used from anywhere in the application.

Here's what you need to do to call `IntegratedAppSupport` from the application:

- To create an object, pass in these three parameters:
  - TeamForge Base URL -- this is where TeamForge is installed
  - Integrated application URL -- this is where the application is installed
  - Application Name -- the name of the application as defined in the integrated application's descriptor

Typically these parameters are stored as `web.xml` initialization parameters so that they are passed to `IntegratedAppSupport`'s constructor.

- When you have a validated `IntegratedAppSupport` object, store it in a `ThreadLocal` object so that it can be used in other places. In the example, when constructing the blog URLs in the Pebble blogging application, we get `IntegratedAppSupport` from `ThreadLocal` and retrieve the `IntegratedAppId` from there.

(Refer to `/src/net/sourceforge/pebble/domain/Blog.java`'s `getUrl` method)

Once the object is created, another method called `processRequest` is called on it. This takes an `HttpServletRequest` and `HttpServletResponse` object that the servlet or a filter gets.

- For each request that enters your system and needs to be validated, create an `IntegratedAppRequest` object.

Subsequently, you also need to call the `processRequest` method to set the required parameters internally.

Where you construct the `IntegratedAppSupport` and call `ProcessRequest` subsequently depends on the architecture of your application. We've provided two cases -- for servlet filters and `httpservlets`.

- If your application has servlet filters, use them as shown in the example Pebble application. (Refer to `src/net/sourceforge/pebble/web/filter/BlogLookupFilter.java` and `src/com/collabnet/pebble/filter/CTFAuthenticationFilter.java`.)
- If your application does not use servlet filters, but directly calls an `HttpServlet` as the point of entry, place the call to `IntegratedAppRequest` where the application enters the `HttpServlet` code.

`ProcessRequest` does the following tasks:

- If the URL is coming in for the first time from the TeamForge, it will have the `singlesignon` token. The method validates the token and stores a cookie which identifies the soap session for the user as well as the project for which this request is being made. This will also be the case for a Go URL coming in from TeamForge.
- If it is a subsequent request (from a form submission or a click from a previous page), then the cookie information is picked up and used in `IntegratedAppSupport`. For this, the integrated application expects that the `linkid` (the id

linking the integrated application and the project, obtained through `IntegratedAppSupport.getIntegratedAppId()` is provided either as a request attribute or as a request parameter or in the request path -- it can be any one of these:

```
http://my.integrated.application/reach/linkid/prplxxxx/me
http://my.integrated.application/reach/me?linkId=prplxxxx
```

a request parameter obtained through any other means and set as a Request Attribute.

It is advisable to set it so that all originating URLs (forms and other links to the application) have the `/linkid/prplxxxx` in their request URL. For example, in the Pebble application (refer to `./src/net/sourceforge/pebble/domain/Blog.java`'s `getUrl` method), the URL for each link that goes out of the blog is added a `/linkid/prplxxxx` to its request URL. This helps subsequent URLs to be validated correctly. You can retrieve the linkid using `IntegratedAppSupport.getIntegratedAppId`.

## Internationalize your integrated application

---

You can configure your integrated application to display languages based on the TeamForge site user's browser locale.

These parameters can be internationalized:

- The description of the integrated application.
- The title and description of the integrated application when it appears as a component on a project page.
- The name and description of any configuration parameter.

1. In the *XML application configuration file*, select the values for which you want to provide localized content.

2. Inside the `bundles` tag, create a `bundle` tag identical to the default English `bundle` tag.

(Keep the `en` bundle.)

3. Change the value of the `locale` attribute of the new `bundle` tag to the language you are going to provide.

The value for each of the internationalized tags must start with `l10n`.

For example:

```
<bundles>
  <bundle locale="en">
    <key name="l10n.application.description">My Test App</key>
    <key name="l10n.blogname.title">Blog Name</key>
    <key name="l10n.blogname.description">Please provide a name
for the Blog. This appears on all blog pages</key>
    .....
  </bundle>
  <bundle locale="ja">    <!-- Japanese Language Translations -->
    <key name="l10n.application.description">My Test App</key>
    <key name="l10n.blogname.title">Blog Name</key>
    <key name="l10n.blogname.description">Please provide a name
for the Blog. This appears on all blog pages</key>
    .....
  </bundle>
</bundles>
```

4. Save the application configuration file.

## Maintain your TeamForge 17.4 site

---

You've got TeamForge installed, up to date, and operating on the appropriate scale. Now you're in day-to-day maintenance mode. While you're waiting around for something to go wrong, there's a lot you can do to support users and project managers, monitor the site's performance, and generally keep things running smoothly.

### TeamForge System Administration

---

You can follow these useful best practices for your TeamForge install and upgrade procedures.

#### Install project templates manually


Provide sample projects to help users get started quickly.

TeamForge comes with a sample template useful for agile development projects. Site administrators and project managers can use this template to jump-start a project without a lot of manual setup steps.

In the TeamForge installation directory, run the `install-project-templates.py` script.

```
cd /opt/collabnet/teamforge/installer/
./install-project-templates.py -V
```

Use a site administrator user name and password. For a new site, these are `admin` and `admin`.

 **Tip:** On some servers, it may take a few seconds for the SOAP server to be ready after installation. If `install-project-templates.py` returns an error, try waiting briefly and then running it again.

#### Install project templates using the installer

With TeamForge 17.4, you can install project templates automatically.

TeamForge comes with a sample template useful for agile development projects. Site administrators and project managers can use this template to jump-start a project without a lot of manual setup steps.

- To install project templates automatically when the TeamForge startup script is executed, you must set these tokens:

- ```
INSTALL_TEMPLATES = true
```
- ```
REQUIRE_USER_PASSWORD_CHANGE = false
```

 **Note:**

- The TeamForge startup script installs the sample TeamForge project templates by default.
- If the project templates are already installed, you cannot re-install them using the TeamForge startup script.
- You may choose to delete the sample project templates. After deleting the sample project templates, you must set the `INSTALL_TEMPLATES` site options token to false. Otherwise, the project templates, if not found in the database, are installed automatically every time you restart the CollabNet services.

#### Set up an Oracle database

To use an Oracle database for your TeamForge data, set up the Oracle database and tell the installer how to handle it.

##### TeamForge database setup

1. Make sure your database uses UTF8 or AL32UTF8 encoding.

This is needed to support users in Asian languages.

For information about discovering and changing the database encoding, see [this Oracle knowledge base article](#).

**2.** Connect to your Oracle database.

```
SQL> connect <adminusername>@<db_name>/<adminpassword> as sysdba
```

**3.** Create the database user and password you will use to connect from TeamForge to Oracle.

```
SQL> create user <sf user> identified by <sf passwd> default tablespace
<your tablespace> temporary tablespace <temporary tablespace>;
```

**4.** Grant permissions to the user that you just created.

```
SQL> grant unlimited tablespace to <sf user>;
SQL> grant create snapshot to <sf user>;
SQL> grant create cluster to <sf user>;
SQL> grant create database link to <sf user>;
SQL> grant create procedure to <sf user>;
SQL> grant create sequence to <sf user>;
SQL> grant create trigger to <sf user>;
SQL> grant create type to <sf user>;
SQL> grant create view to <sf user>;
SQL> grant query rewrite to <sf user>;
SQL> grant alter session to <sf user>;
SQL> grant create table to <sf user>;
SQL> grant create session to <sf user>;
SQL> grant create any synonym to <sf user>;
SQL> exit
```

**5.** Create the database read-only user that you will use to connect from TeamForge.

```
SQL> create user <database_readonly_user> identified by
<database_readonly_password> default tablespace <your tablespace>
temporary tablespace <temporary tablespace>;
```

**6.** Grant the required permissions to the read-only user that you just created.

```
SQL> grant create session to <database_readonly_user>;
SQL> exit
```

**7.** Connect to your Oracle database as <sf user>.

```
SQL> connect <sf user>@<db_name>/<sf passwd>
```

**8.** Grant the 'create synonym' permission on TeamForge database to the read-only user that you just created.

```
SQL> begin
for i in (select table_name from user_tables) loop
execute immediate 'grant select on '|| i.table_name||' to
<database_readonly_user>';
execute immediate 'create synonym <database_readonly_user>.'||
i.table_name||' for '||i.table_name||'';
end loop;
end;

SQL> exit
```

**Datamart setup**

**9.** Make sure your database uses UTF8 or AL32UTF8 encoding.

This is needed to support users in Asian languages.

For information about discovering and changing the database encoding, see [this Oracle knowledge base article](#).

**10.** Connect to your Oracle database.


```
SQL> connect <adminusername>@<db_name>/<adminpassword> as sysdba
```

**11. Create the datamart user you will use to connect from TeamForge.**

```
SQL> create user <reports_database_user> identified by
<reports_database_password> default tablespace <your tablespace> temporary
tablespace <temporary tablespace>;
```

**12. Grant permissions to the user that you just created.**

```
SQL> grant unlimited tablespace to <reports_database_user>;
SQL> grant create snapshot to <reports_database_user>;
SQL> grant create cluster to <reports_database_user>;
SQL> grant create database link to <reports_database_user>;
SQL> grant create procedure to <reports_database_user>;
SQL> grant create sequence to <reports_database_user>;
SQL> grant create trigger to <reports_database_user>;
SQL> grant create type to <reports_database_user>;
SQL> grant create view to <reports_database_user>;
SQL> grant query rewrite to <reports_database_user>;
SQL> grant alter session to <reports_database_user>;
SQL> grant create table to <reports_database_user>;
SQL> grant create session to <reports_database_user>;
SQL> grant create any synonym to <reports_database_user>;
SQL> exit
```


 **Note:** Replace <reports\_database\_user> with the datamart username specified in the `site-options.conf` and <reports\_database\_password> with the database password specified in `site-options.conf`.

**13. Create the datamart read-only user that you will use to connect from TeamForge.**

```
SQL> create user <reports_readonly_user> identified by
<reports_readonly_password> default tablespace <your tablespace> temporary
tablespace <temporary tablespace>;
```

**14. Grant the required permissions to the read-only user that you just created.**

```
SQL> grant create session to <reports_readonly_user>;
SQL> exit
```

 **Note:** The TeamForge installer creates the tables and default values for you.

**15. Connect to your Oracle database as <reports\_database\_user>.**

```
SQL> connect <reports_database_user>@<db_name>/<reports_database_password>
```

**16. Grant the 'create synonym' permission on TeamForge datamart to the read-only user that you just created.**

```
SQL> begin
for i in (select table_name from user_tables) loop
execute immediate 'grant select on '|| i.table_name||' to
<reports_readonly_user>';
execute immediate 'create synonym <reports_readonly_user>.'||
i.table_name||' for '||i.table_name||'';
end loop;
end;

SQL> exit
```

## Change the scmviewer password

It is recommended to change the scmviewer password after installing TeamForge.

Follow these steps to change the scmviewer password:

1. Stop TeamForge.

```
/opt/collabnet/teamforge/bin/teamforge stop
```

2. Create an encrypted password using the [password\\_util.sh](#) on page 795

```
/opt/collabnet/teamforge/runtime/scripts/password_util.sh -encrypt
'<new_password_text>'
```

3. Set the encrypted password to the [SCM\\_USER\\_ENCRYPTED\\_PASSWORD](#) on page 770 token in the `/opt/collabnet/teamforge/etc/site-options.conf` file.

4. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

5. Start TeamForge.

```
/opt/collabnet/teamforge/bin/teamforge start
```

6. Run the following command to set the new password for Review Board.

```
cd /opt/collabnet/RBInstaller-17.4.7
python ./install.py --auth-scmuser
```

## Install Memcached

Memcached caches Subversion (SVN) authentication and authorization information and serves the `mod_authnz_ctf` module's authentication and authorization requests thereby reducing the number of SOAP calls, which in turn results in less load on the TeamForge Application Server.

- See this [wiki page](#) for more information about Memcached.
- Memcached can run on the TeamForge Application Server or on a separate server. This document assumes that you install Memcached on the TeamForge Application Server.
- Caching is disabled by default in TeamForge. Enable it by setting up the following `site-options.conf` tokens.
  - [ENABLE\\_CACHING\\_WITH\\_MEMCACHED](#)
  - [MEMCACHED\\_SERVER\\_HOST](#)
  - [MEMCACHED\\_SERVER\\_PORT](#)
  - [MEMCACHED\\_SERVER\\_TTL](#)

To set up caching with Memcached:

1. Install Memcached.
2. Configure Memcached settings.
3. Set up the TeamForge site options.
4. Recreate the TeamForge runtime environment.

### Do this on the TeamForge Application Server.

1. Install Memcached.


```
yum clean all
```

```
yum -y install memcached
```

2. Configure the `OPTIONS` key in the Memcached configuration file (`/etc/sysconfig/memcached`) and start Memcached.

- a) The `OPTIONS` key in the `memcached` configuration file is used to set additional options during Memcached startup. Add the `-l <ip-addr>` flag to have Memcached listen to `<ip_addr>`. This is an important option to consider as there is no other way to secure the installation. Binding to an internal or firewalled network interface is recommended.

```
vi /etc/sysconfig/memcached
```

-  **Important:** Remove the `-l` flag from the `OPTIONS` key to have Memcached listen to the server's default IP address or host name, including the 'localhost'.

- b) Start Memcached.

**Start Memcached during system boot up:** `sudo systemctl enable memcached`

**Start Memcached normally:** `sudo systemctl start memcached`

3. Set up the following tokens in the TeamForge `site-options.conf` file.

See:

- [ENABLE\\_CACHING\\_WITH\\_MEMCACHED](#)
- [MEMCACHED\\_SERVER\\_HOST](#)
- [MEMCACHED\\_SERVER\\_PORT](#)
- [MEMCACHED\\_SERVER\\_TTL](#)

4. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

-  **Note:**


The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

## Synchronize TeamForge source control integrations


Any time you upgrade your TeamForge site or a source control application, you must ensure that your users can still access their source code.

1. Click **Admin** in the TeamForge navigation bar.
2. On the site administration navigation bar, click **Integrations**.
3. For each source control service you are supporting, verify that the right paths are specified.
  - **SOAP service host** should be `localhost` or the host name of the server on which you just installed TeamForge.
  - **Repository base URL** should be the URL for the top level of your source code server (which may be the same as your application server). For example, `http://<myscmbbox>/svn/repos`
  - **SCM Viewer URL** should be the URL for the ViewVC application on your source control server. For example, `http://<myscmbbox>/integration/viewvc/viewvc.cgi`

-  **Note:** If you want to turn on TeamForge code browser, specify the appropriate URL.

4. Select all your CVS integrations and click **Synchronize Permissions**.

This updates the permissions on your code repositories so that users can access them from the new site.

-  **Note:** By default, the `DISABLE_CREATE_INTEGRATION_SERVERS` flag in the `site-options.conf` file is set to false, which allows users to create new external integrations. To suppress the ability to add

integrations, change this setting to true and recreate the runtime environment before making the site available to users.

## Upgrade CLI reports to the latest version on Red Hat or CentOS

Use these instructions to upgrade the CLI reports to the latest version.

Though life cycle metric reports are released as part of regular TeamForge product releases (including patch releases), new reports are, at times, made available to customers as and when they are ready. When a new report is available, you can upgrade your CLI reports to the latest version.

1. Log on to TeamForge application server as root.
2. Install the latest `ctf_clireports`.

```
yum install ctf_clireports
```

3. Run the `post-install.py` script.

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

4. Run the following command and verify the CLI reports installation. Observe the version and release information.

```
rpm -qi ctf_clireports
```

## Customize your Apache configuration

The following instructions illustrate how you can include custom configuration to Apache and disable the same if not required.

1. Create `conf.d/httpd/httpd.conf.d` under `/opt/collabnet/teamforge/etc/` directory.
2. Include `custom.conf` under `/opt/collabnet/teamforge/etc/conf.d/httpd/httpd.conf.d/`.
3. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

The following warning message is displayed, which you can ignore.

```
Custom configuration found in /opt/collabnet/teamforge/etc/conf.d/httpd/httpd.conf.d has been applied. Please be informed that such configuration may impact the reliability of TeamForge.
```

The following line is added to `/etc/httpd/conf/httpd.conf`:

```
Include /opt/collabnet/teamforge/etc/conf.d/httpd/httpd.conf.d/
```

4. Run the `httpd -e info` command to know the Apache configuration/syntax errors, if any.

### Remove custom configuration

5. To remove custom configuration:

```
cd /opt/collabnet/teamforge/etc/conf.d/
mv httpd/ httpd_old
```

6. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



### Note:



The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

## Customize your PostgreSQL configuration

The following instructions illustrate how you can include custom configuration to PostgreSQL and disable the same if not required.

1. Create `conf.d/pgsql/pg_hba.conf.d/` under `/opt/collabnet/teamforge/etc/` directory.

If the reporting service is running on a separate port, create `conf.d/reports-pgsql/pg_hba.conf.d/` under `/opt/collabnet/teamforge/etc/`.

2. Include `custom.conf` under `/opt/collabnet/teamforge/etc/conf.d/pgsql/pg_hba.conf.d/`.

If the reporting service is running on a separate port, include `custom.conf` under `/opt/collabnet/teamforge/etc/conf.d/reports-pgsql/pg_hba.conf.d/`

3. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

The following warning message is displayed, which you can ignore.

```
Custom configuration found in /opt/collabnet/teamforge/etc/conf.d/pgsql/pg_hba.conf.d has been applied. Please be aware of that such configuration may impact the reliability of TeamForge.
```

If the reporting service is running on a separate port:

```
Custom configuration found in /opt/collabnet/teamforge/etc/conf.d/reports-pgsql/pg_hba.conf.d has been applied. Please be aware of that such configuration may impact the reliability of TeamForge.
```

Configuration settings from `custom.conf` are included in `/var/lib/pgsql/9.6/data/pg_hba.conf`.

If the reporting service is running on a separate port, configuration settings from `custom.conf` are included in `/var/lib/pgsql/9.6/reports/pg_hba.conf`.

4. Check the `postgresql.log` file for any syntax errors: `/opt/collabnet/teamforge/log/pgsql/postgresql.log`.

### Remove custom configuration

5. To remove custom configuration:

```
cd /opt/collabnet/teamforge/etc/conf.d/
mv pgsql pgsql_old
```

If the reporting service is running on a separate port:

```
cd /opt/collabnet/teamforge/etc/conf.d/
mv reports-pgsql reports-pgsql_old
```

6. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



#### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

## EventQ System Administration

---

Manage the TeamForge EventQ server-side components

### System administration introduction

The scope of this guide is limited to key administrative features like starting/stopping services, log locations, etc. Administration and configuration of dependent services (e.g., nginx, MongoDB, RabbitMQ) is not covered; see respective product documentation. It is assumed that as TeamForge EventQ administrator you are familiar with basic Linux administration commands and terminology.

### Background services

TeamForge EventQ installs and starts six background services on the App server. Each service is managed using init scripts, which are installed and started as part of the installation process. Init scripts are all located in `/etc/init.d/` on the respective hosts. Once started, each service creates a pid file.

PID file	Host	Description
<code>/opt/collabnet/collabnet-mongodb/mongodb.pid</code> mongod	DB server	Database service
<code>/opt/collabnet/collabnet-nginx/nginx.pid</code> nginx	App server	Web and application server
<code>/opt/collabnet/collabnet-rabbitmq/rabbitmq.pid</code> rabbitmq-server	MQ server	AMQP message service
<code>/opt/collabnet/collabnet-redis/redis.pid</code> redis	App server	Internal job queue service
<code>/opt/eventq/eventq/ctf_job_runner.pid</code> ctf-jobs	App server	Background communication with TeamForge
<code>/opt/eventq/eventq/scm_queue.pid</code> queue-commit	App server	Processes message from commit queue
<code>/opt/eventq/eventq/build_queue.pid</code> queue-build	App server	Processes message from build queue
<code>/opt/eventq/eventq/review_queue.pid</code> queue-review	App server	Processes message from review queue


## Starting and stopping services

Use the init scripts to start, stop, restart, or get status on a service individually.

<b>status</b>	Is the process currently running? Does it have a PID (even if it's not running - e.g. in the case of a crash)
<b>start</b>	Starts the process
<b>stop</b>	Stop the process
<b>restart</b>	Stop and then start the process

Example usage:

```
sudo service eventq-queue-commit restart
```

 **Tip:** While each service can be managed individually, the `eventq` init script will start and stop all App server services: `/etc/init.d/eventq`

Example usage:

```
sudo service eventq start
```

## Service administration passwords

During the installation of TeamForge EventQ, if RabbitMQ and MongoDB were installed as part of the installation process then the administrative passwords have been written to disk for reference:

- DB server: `/opt/collabnet/mongodb/.mongo_passwords`
- MQ server: `/opt/collabnet/rabbitmq/.rabbit_passwords`

## Log files

The following log files are relevant to TeamForge EventQ services and may be useful for debugging. Log rotation on all log files is enabled by default.

### EventQ App server logs:

The `server.log` is the main log file for TeamForge EventQ, detailing the server's page serves and any errors that EventQ may encounter. It relates to the Phusion Passenger service, which is started with `nginx`.

`/var/log/eventq/`

- `build_queue.log`
- `ctf_job_runner.log`
- `nginx-error.log`
- `nginx.log`
- `review_queue.log`
- `scm_queue.log`
- `server.log`

`/var/log/collabnet-redis/`

- `redis.log`

### MongoDB server log:

- `/var/log/collabnet-mongodb/mongodb.log`

### RabbitMQ server log:

`/var/log/collabnet-rabbitmq/`

- `rabbit@localhost.log`
- `rabbit@localhost-sasl.log`
- `startup_err`

- `startup_log`


## Backup and restore

See [Back up and restore EventQ data](#)

## Monitoring TeamForge EventQ with New Relic

TeamForge EventQ may be configured for monitoring using New Relic, a third party service.

This article explains how to set up monitoring of TeamForge EventQ using a third party product, New Relic. Once set up, configure New Relic to alert you of outages, monitor application performance, observe errors, etc. It is assumed that you have a New Relic account and license key ready for use.

 **Note:** New Relic is not affiliated with CollabNet and may bill you for usage of their services.

### 1. Create a New Relic account and obtain your license key

If you haven't done so already, create an account on the New Relic web site. Obtain the license key associated with your account.

### 2. Contact [CollabNet Support](#) and request the TeamForge EventQ New Relic monitoring script.

### 3. Copy the New Relic script onto the TeamForge EventQ application server host (App server)

```
cp newrelic.sh /tmp cd /tmp chmod 755 newrelic.sh
```

### 4. Execute the New Relic script with root privileges, passing your license key as an argument

```
sudo ./newrelic.sh LICENSE_KEY
```

where `LICENSE_KEY` is replaced with your New Relic provided license key

After a few moments, your New Relic account will reflect the TeamForge EventQ server. Follow New Relic's instructions for setting up monitoring, alerts or other desired functionality.

## EventQ Troubleshooting

---

You may encounter these problems when using TeamForge EventQ:

### Certain activities appear grayed-out and I cannot get an activity's detail information

TeamForge EventQ sometimes learns about certain activities through an association reference before the details of the activity are officially reported by the source. Such activities are termed **Phantom Activities** and appear grayed-out temporarily. Once TeamForge EventQ receives information on this activity it creates a regular Details page and everything goes back to normal. Traceability graphs include these activities, but you cannot click on the commit nodes to get commit details. Activities may be perpetually grayed-out when TeamForge EventQ is unable to retrieve activity information such as:

- Activity details of events that occurred before TeamForge EventQ was started — TeamForge EventQ begins to gather data from configured sources as soon as those connections are established, but cannot retrieve historical activity data.
- Activities that occurred during system down-time — In unusual circumstances such as power outage or network failure, TeamForge EventQ may never

receive information about activities that occurred during such events.

### Activities in the Activity stream appear out of chronological sequence

Solve this problem by synchronizing the clocks on servers that report data to TeamForge EventQ. The Activity stream shows activities in reverse chronological sequence. EventQ obtains the time stamp for activities from the originating server. Failure to synchronize the clocks of the various servers reporting data to TeamForge EventQ may result in activities that appear out of sequence. On Unix systems, use NTP or similar services to synchronize server clocks.

## Authenticate users with LDAP

---

Use LDAP to facilitate managing users and groups in TeamForge.

### Set up LDAP integration for the TeamForge site

Follow these steps to convert your TeamForge installation to authenticate against your corporate OpenLDAP server.

 **Note:** Refer to the [Software requirements for TeamForge 17.4](#) topic for the supported OpenLDAP versions.

#### 1. Stop TeamForge.

```
sudo /opt/collabnet/teamforge/bin/teamforge stop
/opt/collabnet/teamforge/bin/teamforge stop
```

#### 2. Edit the `site-options.conf` file.

- a) Enable TeamForge to use LDAP authentication by editing the `site-options.conf` file, for example, edit `/opt/collabnet/teamforge/etc/site-options.conf` file.

Under "External User Authentication," uncomment this line:

```
USE_EXTERNAL_USER_AUTHENTICATION=false
```

and change its value to true.

- b) Configure the `site-options` tokens.

 **Note:** The values specified for the following tokens are only for illustration purpose.

```
EXTERNAL_AUTHENTICATION_TYPE=ldap
LDAP_DN_PREFIX=cn=
LDAP_DN_SUFFIX=,cn=Users,dc=testldap,dc=qa,dc=collab,dc=net
LDAP_SERVER_URL=ldap://testldap.qa.collab.net:3268
```

#### 3. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

#### 4. Start TeamForge.


```
sudo /opt/collabnet/teamforge/bin/teamforge start /opt/collabnet/
teamforge/bin/teamforge start
```

## Turn off LDAP authentication

During some maintenance operations, such as upgrades, you may need to turn off LDAP authentication.

1. Open the `site-options.conf` file, the master configuration file that controls your TeamForge site.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

 **Note:** `vi` is an example. Any \*nix text editor will work.

2. In the `site-options.conf` file, comment out these variables:

```
USE_EXTERNAL_USER_AUTHENTICATION
```

```
LOGIN_CONFIG_XML_FILE
```

```
MINIMUM_PASSWORD_LENGTH
```

3. Recreate the runtime environment.

```
./install.sh -V -r -d /opt/collabnet/teamforge
```

4. Review the variables you have changed, then save the `site-options.conf` file.

## Authenticate users with LDAP using Auth Manager

Use the Auth Manager to effectively manage and synchronize user profiles with LDAP.

There are a few limitations to the conventional method used to enable LDAP authentication. For example, when an external authentication source is configured using the conventional method, the template XML file needs to be copied and edited manually. The manual intervention may lead to ambiguities and it makes the process error-prone.

TeamForge makes the entire external authentication process easier with the Auth Manager add-on. It provides users the ability to create multiple profiles, manage profiles, and perform LDAP synchronization.

The Auth Manager allows you to:

- import `login-conf.xml` file
- maintain multiple profiles in LDAP servers
- manage each profile individually
- activate a profile without recreating runtime
- store files in the Integration Data Service (IDS) that survives an upgrade

### Install Auth Manager

The Auth Manager add-on provides customers with a central authentication service the ability to integrate TeamForge with external authentication services such as LDAP, Active Directory, and Kerberos.

The Auth Manager TeamForge add-on is available as an RPM file that you have to download and install. Contact CollabNet Support for more information.

1. Log on to TeamForge as a root user.
2. Extract the RPM file.  
Extracting creates the add-on directory at `/opt/collabnet/teamforge/add-ons`
3. Navigate to the new add-ons directory.

```
cd /opt/collabnet/teamforge/add-ons/ctf_authentication_manager
```

4. Install the Auth Manager:

```
./install
```

- Choose to synchronize with LDAP for user data or make use of the user provided data, as required.

For example, if you want to:

- create a user profile quickly, use the data available in LDAP by enabling LDAP sync and running `hide.sh` script. It displays only the **Re-type password** field to the user.
- create a user profile using the data provided by the user, disable LDAP and run the `show.sh` script. It displays all the fields that you expect the user to fill in and requires site administrator's approval. This is fairly a time-consuming process.

- Set up your site's master configuration file.

```
vim /opt/collabnet/teamforge/etc/site-options.conf
```

- Set `APPROVE_NEW_USER_ACCOUNTS` as false.

Hide fields: To skip the approval process and create an user profile with the data available in LDAP, you have to enable LDAP Sync and run the `hide.sh` script after installation. It conceals all the fields on the Create New User page except the Re-type password field.

Show fields: To get the data from the user and not through the LDAP sync, you have to disable LDAP Sync run the `show.sh` script after installation. It shows all fields including full name, email, locale string, and license type on the Create New User page.

- Set `REQUIRE_PASSWORD_SECURITY` as false.
- Set `PASSWORD_REQUIRES_NUMBER` as false.
- Set `PASSWORD_REQUIRES_NON_ALPHANUM` as false.
- Set `USE_EXTERNAL_USER_AUTHENTICATION` as true.
- Set `REQUIRE_USER_PASSWORD_CHANGE` as false.
- Set `MINIMUM_PASSWORD_LENGTH` as 0.
- Set `PASSWORD_REQUIRES_MIXED_CASE` as false.

- Protect Auth Manager with SSL, if preferred. [Click here](#) for more details.

- Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



#### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

- Start TeamForge.

```
/opt/collabnet/teamforge/bin/teamforge start
```



**Tip:** To ensure that the installation has been completed successfully and the external authentication functionality works do the following:

- Login to the TeamForge through UI as an admin user and check if the add-on is appearing as Auth Manager in the project navigation bar. Also, for fresh installation, an active **Default TeamForgeDatabase** profile appears under **Manage Existing profiles**, by default, with the green status indicator.
- Alternatively, in the CLI, scrutinize the log files, for example, `/opt/collabnet/teamforge/log/apps/server.log`.

## Configure AUTH MANAGER in TeamForge

You can configure Auth Manager as a linked application in TeamForge.

As a result of this configuration, the Auth Manager appears in TeamForge's project navigation bar as a linked application.

- Log on to the TeamForge as a site administrator and go to the **look** project.

2. Click **PROJECT ADMIN** from the **Project Home** menu.
3. On the **Project Admin Menu**, click **Project Toolbar**. Then click **LINKED APPLICATIONS**.  
A list of all currently linked applications in the project is displayed.
4. Click **Create**.
5. On the **Create Linked Application** page, enter the name for the linked application as **AUTH MANAGER**.
6. Enter the URL.  
`http://<host>/authenticationManager`
7. Enable Single Sign On (SSO) to allow TeamForge users to automatically log into the Auth Manager.
8. Click **Save**.  
The application is linked to the TeamForge and the it appears as **AUTH MANAGER** in the project navigation bar.

### Manage authentication profiles


TeamForge utilizes its own database to validate the user name and password and also supports external authentication sources such as LDAP, Active Directory, Kerberos and Master Password.

#### Create a profile


You can create a profile that determines the authentication method and settings in TeamForge.


You need to have at least one user profile or more in active status before configuring the external authentication.

1. Log on to the TeamForge as a site administrator and go to the **look** project.
2. From the project navigation bar, click **AUTH MANAGER**.
3. From the **Main Menu** pane on the left, click **Create Profile**.

 **Fastpath:** In the **Manage Existing Profiles** page, if you do not find the desired one in the list of existing authentication profiles, you can click **New Profile** and proceed.



4. From the drop-down menu, select the type of the new authentication profile.
  - LDAP: It uses the user name and password provided by the user to bind to the LDAP. If the bind is successful, the user is authenticated. This is a simple method of authentication. [Click here](#) for more information.
  - LdapExtended: It uses a service account to bind to the LDAP. Customizable filters are used to bind with the user and to validate authentication. [Click here](#) for more information.

 **Tip:** Use this module if the users are spread over LDAP or when a group membership is required to access TeamForge.

 **Important:** You can use only the **LdapExtended** profile as a source for LDAP Sync.

  - Active Directory: It uses the user data configured through Microsoft's Active Directory. This is a simple method of authentication.
  - Kerberos: It uses MIT KRB5 authentication. Contact your network admin for the host configuration settings.
5. Set the Jboss flag that determines the behaviour of the control flag with multiple login-modules.
  - Sufficient: The login-module is not required to succeed. If it does succeed, control immediately returns to the application (authentication does not proceed down the login-module list). If it fails, authentication continues down the login-module list.
  - Optional: The login-module is not required to succeed. If it succeeds or fails, authentication still continues to proceed down the login-module list.
  - Required: The login-module is required to succeed. If it succeeds or fails, authentication still continues to proceed down the login-module list.
6. Enter the value for each module property listed for the chosen profile type.
7. Click **Create**.  
The confirmation message, *The authentication profiles have been imported. Activate the profiles to apply to TeamForge authentication*, appears.



-  **Remember:** The newly created profile is listed under **Authentication Profiles** in the **Manage Existing Profile** page. It is now inactive and the status indicator is yellow. You must *activate* the newly created user profile.
-  **Tip:** Before you create any profiles using Auth Manager, you may see an inactive auto-imported **TeamForgeDatabase** profile appearing under **Authentication Profiles**. It is recommended to *delete* the **Auto-imported UsernamePasswordInDatabaseLoginModule** after creating and activating your first profile. Because the subsequent login and authentication request pass only through the active profile(s).

### Configure Master Password

With the Site Administrative privileges, you can generate a master password and use it along with any user name registered with the TeamForge.

You can create an authentication profile for master password users in Auth Manager. Impersonation is not supported in TeamForge but the master password feature facilitates any user in TeamForge, irrespective of the roles, to login as another user if they have a valid master password.

The Site Admin needs to generate the password through scripts and configure the authentication user profile using Auth Manager. It is important to have this configured for a web-based SSO system.


1. On the command-line interface, login to TeamForge as a root user to generate the master password.
2. Navigate to the new add-ons directory.

```
cd /opt/collabnet/teamforge/add-ons/ctf_authentication_manager
```


3. Set the master password.

```
./mpasswd.sh
```

4. Enter the password and re-enter when prompted for confirmation.

-  **Note:** Ensure that the password is masked while entering it. Also keep the master password confidential and share it with authenticated users on demand. It is a good practice to keep changing the master password frequently.


5. Log on to the TeamForge as a site administrator and go to the **look** project.
6. From the project navigation bar, click **AUTH MANAGER**.
7. From the **Main Menu** pane on the left, click **Create Profile**.

-  **Fastpath:** In the **Manage Existing Profiles** page, if you do not find the desired one in the list of existing authentication profiles, you can click **New Profile** and proceed.

8. On the **Create Authentication Profile** page, select **MasterPassword** from the drop-down list.
9. Enter an appropriate name for the new **MasterPassword** user profile.
10. Set the **Jboss** flag that determines the behaviour of the control flag with multiple login-modules.
  - **Sufficient:** The login-module is not required to succeed. If it does succeed, control immediately returns to the application (authentication does not proceed down the login-module list). If it fails, authentication continues down the login-module list.
  - **Optional:** The login-module is not required to succeed. If it succeeds or fails, authentication still continues to proceed down the login-module list.
  - **Required:** The login-module is required to succeed. If it succeeds or fails, authentication still continues to proceed down the login-module list.

11. Click **Create**.

The confirmation message, *The authentication profiles have been imported. Activate the profiles to apply to TeamForge authentication*, appears.

-  **Remember:** The newly created profile is listed under **Authentication Profiles** in the **Manage Existing Profile** page. It is now inactive and the status indicator is yellow. You must *activate* the newly created user profile.

-  **Tip:** Before you create any profiles using Auth Manager, you may see an inactive auto-imported **TeamForgeDatabase** profile appearing under **Authentication Profiles**. It is recommended to *delete* the

**Auto-imported UsernamePasswordInDatabaseLoginModule** after creating and activating your first profile. Because the subsequent login and authentication request pass only through the active profile(s).

### Upgrade a profile

You can upgrade a legacy authentication profile by importing the corresponding configuration file for authentication.

You can upgrade an authentication profile that was used in TeamForge 6.2 or older versions. For each profile, you have to import and save the respective `login-config.xml` file. It imports all the user profiles automatically into the `standalone.xml` file which is later converted into `standalone-full.xml`. The `login-config.xml` and `standalone-full.xml` files are placed at the same location.

1. Log on to the TeamForge as a site administrator and go to the **look** project.
2. From the project navigation bar, click **AUTH MANAGER**.
3. From the **Main Menu** pane on the left, click **Upgrade Legacy Config**.
4. Select an option to upload your existing authentication configuration.
  - a) Detected Configuration Files: The list of files detected in the default location.
  - b) Specify Server-based Configuration File: Specify the location of the `login-config.xml` file in the TeamForge Server.
  - c) Upload Existing Configuration File: Specify the legacy `login-config.xml` file saved in the local.
5. Click **Save**.

The confirmation message, *The authentication profiles have been imported. Activate the profiles to apply to TeamForge authentication*, appears.



**Remember:** The newly created profile is listed under **Authentication Profiles** in the **Manage Existing Profile** page. It is now inactive and the status indicator is yellow. You must *activate* the newly created user profile.

### Reorder profiles

You can set the order in which a user profile needs to be considered for authentication. On the **Manage Existing Profiles** page, you can manually reorder the list of authentication profiles.

For example, consider LDAP server 'A' that has 100 users profiles and LDAP server 'B' with just 15 users profiles. To reduce the network traffic effectively, you may arrange authentication profiles in such a way that the authentication requests pass through 'A' first and then through 'B'. When you enable reorder option, you can manually move the profiles back and forth within the list.

1. Log on to the TeamForge as a site administrator and go to the **look** project.
2. From the project navigation bar, click **AUTH MANAGER**.
3. From the **Main Menu** pane on the left, click **Manage Existing Profiles**.
4. Click **Enable Reorder**.  
You are now able to drag and drop profiles within the list.
5. Reorder the list as required and click **Save Order**.

### Activate a profile

To pass authentication requests through a profile, you must change the status of a profile to active. The status indicator is green for active profiles in .


By default, the status of a newly created profile is inactive. A profile creation process in is considered complete when you change the status from inactive to active. Activated profiles are automatically listed in the `standalone-full.xml` file.

You can activate a newly created profile or an existing profile that is currently inactive. The profiles that are listed in **Manage Existing Profiles** page with yellow status indicators are all inactive or deactivated. You can have multiple profiles in the active status and *reorder* them within the list, if required.




**Tip:** Before you create any profiles using Auth Manager, you may see an inactive auto-imported **TeamForgeDatabase** profile appearing under **Authentication Profiles**. It is recommended to *delete* the **Auto-imported UsernamePasswordInDatabaseLoginModule** after creating and activating your first profile. Because the subsequent login and authentication request pass only through the active profile(s).

1. Log on to the TeamForge as a site administrator and go to the **look** project.
2. From the project navigation bar, click **AUTH MANAGER**.
3. From the **Main Menu** pane on the left, click **Manage Existing Profiles**.
4. To activate all the profiles listed in the **Manage Existing Profiles** page atonce, click **Activate All**.
5. To activate a particular profile in the **Manage Existing Profiles** page, click the profile name that needs to be activated.  
The profile details are displayed.

 **Note:** Ensure that the current status of the profile is inactive and marked yellow.

6. Click **Activate**.

 **Note:** This button is visible only if the current status of the profile is inactive; scroll down to see this.


7. On the confirmation window, click **ok** to proceed.  
The profile is activated now and the status turns green.

### Deactivate a profile


At any point in time, you can deactivate an active profile in . The status indicator is yellow for inactive profiles in .

You may decide to deactivate a profile when you encounter issues with LDAP. If you are half way through the configuration setup and you do not want the profile to be acive for user authentication, you can deactivate it. Or If you want to troubleshoot and identify if the issue encounted prevails in the particuar LDAP server, you can deactivate the authentication profile.


The profiles that are listed in **Manage Existing Profiles** page with yellow status indicators are all inactive or deactivated. By default, the status of a newly created profile is inactive. You can have multiple profiles in the inactive status and *reorder* them within the list, if required. Deactivated profiles are automatically removed from the `standalone-full.xml` file.

 **Important:** The deactivated profile remains in the Integrated Data Space (IDS) and you can reterive the same profile by activating it. The deactivated profile is permanently removed from the application only when it is deleted.

1. Log on to the TeamForge as a site administrator and go to the **look** project.
2. From the project navigation bar, click **AUTH MANAGER**.
3. From the **Main Menu** pane on the left, click **Manage Existing Profiles**.
4. To deactivate all the profiles listed in the **Manage Existing Profiles** page, click **Deactivate All**.
5. To deactivate a particular profile in the **Manage Existing Profiles** page, click the profile name that needs to be deactivated.  
The profile details are displayed.

 **Note:** Ensure that the current status of the profile is active and marked green.

6. Click **Deactivate**.


 **Note:** This button is visible only if the current status of the profile is active; scroll down to see this.

7. On the confirmation window, click **ok** to proceed.  
The profile is deactivated now and the status turns yellow.

### Delete a profile


You can delete an inactive profile from the permanently.

You can delete only deactivated profiles. If you do not want to have a profile in your authentication profile list, you need to deactivate it first and then delete it totally from the system.


 **Important:** A deactivated profile remains in the Integrated Data space (IDS) until you perform deletion. To remove it completely from the system, delete the deactivated profile.

1. Log on to the TeamForge as a site administrator and go to the **look** project.

2. From the project navigation bar, click **AUTH MANAGER**.
3. From the **Main Menu** pane on the left, click **Manage Existing Profiles**.
4. Click the profile that you want to delete.  
The profile details are displayed.
5. Ensure that the current status of the profile is inactive.

 **Note:** The profile status is marked yellow if inactive.

6. Click **Delete**.

 **Note:** This button is visible only if the current status of the profile is inactive.


7. On the confirmation window, click **ok** to proceed.  
The profile has been deleted.

### Configure LDAP sync


Perform LDAP Sync to update TeamForge with the user data available in the LDAP server. You can use an extended LDAP profile as a source for LDAP sync.

LDAP Sync, basically searches in the LDAP server for the user data configured in a login-module. Then it fetches the user data to the TeamForge and performs synchronization. You have options to selectively include login-modules for the LDAP Sync. For example, you have two LDAP accounts, out of which only one needs to be considered for LDAP sync. You need to turn off the LDAP account that you do not want to include in LDAP Sync.

The extended LDAP profile needs to have *Bind DN*, *Bind Credentials*, *Base Filter*, and *Base DN* values for the synchronization. The Base DN, which is not available in other simple LDAP authentications, makes the synchronization possible in **ExtendedLDAP** profile.

 **Important:** Provide the site admin user name and obfuscated site admin password explicitly in `/opt/collabnet/teamforge/var/etc/soap-provider.properties` and do not provide ADMIN account credentials.

1. Log on to the TeamForge as a site administrator and go to the **look** project.
2. From the project navigation bar, click **AUTH MANAGER**.
3. From the **Main Menu** pane on the left, click **Manage Existing Profiles**.
4. Select your desired profile that needs to have the LDAP Sync enabled.
5. Click **Edit**.
6. Select true from the **Use As LDAP Sync Source** drop-down list to consider the selected profile for the synchronization.
7. From the **Main Menu** pane on the left, click **LDAP Sync**.
8. Click **Global Settings** and enable LDAP sync by selecting true from the drop-down list.
9. Click **Group Sync Settings** and do the following to synchronize groups:
  - a) Select true from the **Enable Group Sync** drop-down list to enable the LDAP/AD group synchronization.
  - b) Enter appropriate value in **Group Job Cron Interval** to set the time interval for running group synchronization.  
For example, [click here](#).
  - c) Enter the attribute in **Group Search Filter Expression**, to specify LDAP/AD expression for testing.
  - d) Enter the search text with '\*' at the end in **Group Search Filter Arguments** to synchronize with groups within the search results that have a particular prefix.

 **Tip:** To include all the groups in synchronization, just enter '\*'.

'CTF\*', 'TF\*' or '\*'


10. Click **User and User Data Sync Settings** and do the following to synchronize user status and attributes:
  - a) Select true from the **Enable User Data Sync** drop-down list to enable synchronization of user data.

- b) Enter appropriate value in the **User Data Sync Cron Interval** to set the time interval for group synchronizations. For examples, [click here](#).
- c) Enter the search text with '\*' at the beginning in **User Search Filter Arguments** to synchronize with groups within the search results that have a particular prefix.

 **Tip:** To include all the groups in synchronization, just enter '\*'.

'CTF\*', 'TF\*' or '\*'

- d) Select true from the **Enable LDAP Status Sync** drop-down. It enables the LDAP Sync for the user status. If the user account is disabled in LDAP/AD, it flags and then disables the user account in TeamForge.
- e) Enter the number of days in **User Grace Period** beyond which the user account is disabled in TeamForge, if the user is not existing in LDAP.
- f) Select false from the **Enable User Disable Action** drop-down list. It specifies if the user who has not logged into TeamForge for the specified period, needs to be disabled in TeamForge.
- g) Enter the number of days in **User Disable Interval (Days)** beyond which the user account is disabled in TeamForge, if the user has not logged into TeamForge for the specified period.
- h) Select true from the **Enable User Delete Action** drop-down list to enable the deletion of a disabled user account based on the delete interval.
- i) Enter the number of days in the **User Delete Interval (Days)** drop-down list beyond which the user account is deleted from TeamForge. It is mandatory to have **Enable LDAP Status Sync** or **Enable User Disable Action** enabled.
  - When **Enable LDAP Status Sync** is true, users that do not exist in Active Directory are deleted after the 'delete' and 'grace' intervals from the flagged date. However, users that are existing in Active Directory cannot be disabled or deleted.
  - When **Enable User Disable Action** is true, users are deleted, irrespective of their existence in Active Directory, after the 'disable' and 'delete' intervals. It is calculated from the last login date.
  - When both **Enable LDAP Status Sync** and **Enable User Disable Action** are set to true, users that are not existing in Active Directory are deleted after the 'grace' and 'delete' intervals from the flagged date. And the users that are existing in the Active Directory are deleted after the 'disable' and 'delete' intervals from the last login date.
- j) Select true from the **Enable User Membership Action** drop-down list to include user's LDAP/AD group membership in the synchronization. This needs to be enabled when the **Enable Group Sync** is set to true.
- k) Select true from the **Allow User Re-enabling** drop-down list to re-enable LDAP active users that have pending or disabled status in TeamForge.
- l) Enter the user names in **Excluded Usernames** to skip the respective users accounts during synchronization.
- m) Enter the email address in **Default Email Address** that needs to be associated with the TeamForge user account. It is used as the default email address when the email field is found null or empty.
- n) Enter the number of user batches in **Split Users for LDAP Sync** to perform the synchronization.

 **Note:** The user batch number entered splits the existing number of users into batches and then performs synchronization. Entering '0' or '1' considers all the existing users for the synchronization. Whereas entering '7' splits the existing users into seven batches and completes synchronization on the seventh run.

#### 11. Click to expand **Mail and Reporting Settings** and do the following:

- a) Select true in **Enable User Email Reports** drop-down list to enable the email notification and reporting.
- b) Enter the SMTP host to mail through. If you are using the TeamForge James mail server, enter localhost in **Mail Transport Host**.
- c) Enter the email address of the recipient in **Mail To** that is usually a TeamForge Discussion Forum or a Tracker.
- d) Enter the email address of the sender in **Mail From** that is usually a TeamForge Discussion Forum or Tracker. The sender should be a valid TeamForge user.
- e) Enter any optional email address in **Mail CC** that needs a carbon copy of the email.
- f) Enter the subject line for the email in **Mail Subject**.

- g) Enter the user name in **Mail Username** that authenticates connection to the SMTP servers, if required. Its optional to fill in this field.
- h) Enter the password in **Mail Password** that is required to connect to SMTP servers, if required. Its optional to fill in this field.

12. To save and apply all the changes you made to the profile, click **Save**.

13. To run the LDAP Sync once (on an ad hoc basis), click **Run Once**.

14. Click **Stop** and **Start** buttons to reinitiate the synchronization service.


### Configure selective sync

The Selective sync, otherwise known as single user sync, is similar to the LDAP sync that is performed only for a particular user and not for the users in the entire directory server. The Site Admin can perform synchronization for a selective user on the three selective attributes: full name, email, and organization.


The single user sync is helpful when you encounter inconsistencies or discrepancies in the LDAP behaviour. Some of the common scenarios include network delay, LDAP disconnectivity, and failure of authenticating a valid user.

In cases where the LDAP is not completely reliable, you can consider performing selective sync. For example:

- While performing LDAP sync for a huge site that has thousands of users, there are chances that a handful of users are not synchronized and found missing in the report. Just for users that are missed out, you can synchronize the full name, email, and organization of individual users.
- When the default (dummy) email is used to create a profile, the user may not receive any email notifications. In this case, the Site Admin can update only the email address of the user and perform single user sync. It saves the time spent on LDAP sync that scans through the entire directory server for a single user's email update.

 **Note:** You cannot perform selective sync on users that are disabled and deleted. Single user sync is not applicable for the 'ADMIN' user.

1. Log on to the TeamForge as a site administrator and go to the **look** project.
2. From the project navigation bar, click **AUTH MANAGER**.
3. From the **Main Menu** pane on the left, click **LDAP Sync**.
4. Select true from the **Enable Group Sync** drop-down list under **Global Settings>Group Sync Settings**. It enables the LDAP/AD group synchronization.
5. Select true from the **Enable User Data Sync** drop-down list under **Global Settings>User and User Data Sync Settings**. It enables synchronization of the user data.
6. To save and apply all the changes you made to the profile, click **Save**.
7. Go to **My Workspace > Admin**.
8. On the site administration navigation bar, click **USERS**.
9. On the **USERS** tab, click the name of the user whose account you want to edit.
10. On the **User Details** page, click **EDIT**.
11. On the **Edit User Information** page, make your changes to the full name of the user, email address and organization.
12. Click **Sync User Info** to synchronize the modified user data with the authentication source.

 **Important:** Selective Sync is limited to full name, email, and organization of the user. If you have issues with other attributes of a user profile, try LDAP sync through **Auth Manager**.

### Uninstall Auth Manager

You can remove the Auth Manager completely from the TeamForge. Unlike other linked applications in TeamForge, you need not run the default installer to uninstall the atthis add-on.

Before uninstalling, use the Auth Manager GUI to remove all the authentication profiles except the TeamForgeDatabase profile.

1. Login to TeamForge as a root user.
2. Navigate to the new add-ons directory.

```
cd /opt/collabnet/teamforge/add-ons/ctf_authentication_manager
```

### 3. Uninstall the **AUTH MANAGER**.

```
./uninstall.sh
```

### 4. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

#### **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

### 5. Start TeamForge.

```
/opt/collabnet/teamforge/bin/teamforge start
```

## Field description for Auth Manager

The credential store and identity manager properties that are required to create an authentication profile in the **Auth Manager** are described here.

Field	Description
allowEmptyPasswords	A flag indicating if empty (length 0) passwords should be passed to the LDAP server. An empty password is treated as an anonymous login by some LDAP servers and this may not be a desirable feature. Set this to false to reject empty passwords or true to have the LDAP server validate the empty password. The default is true.
baseCtxDN	It defines the fixed DN of the context to search for user roles. Consider that this is not the Distinguished Name of where the actual roles are located but the DN of where the objects containing the user roles are located (that is, for active directory, this is the DN with the user account).
baseFilter	It defines the search filter used to locate the context of the user to authenticate. The input username/userDN as obtained from the login module callback substitutes the {0} expression. This substitution behavior comes from the standard DirContext?.search(Name, String, Object[], SearchControls? cons) method. An common example search filter is "(uid={0})".
bindCredential	It defines the bindDN password. The password can be encrypted if the jaasSecurityDomain is specified.
bindDN	It defines the DN used to bind to the LDAP server. This is a DN with read/search permissions to the defined baseCtxDN and rolesCtxDN.
java.naming.factory.initial	The classname of the InitialContextFactory implementation. This defaults to the Sun LDAP provider implementation com.sun.jndi.Ldap.LdapCtxFactory.
java.naming.provider.url	This property specifies the host name and port of the DNS server used by the initial DNS context, as well the initial context's domain name.
java.naming.referral	It indicates the service providers how to handle referrals.
java.naming.security.authentication	Specifies the authentication mechanism and the security level to use. This defaults to simple
java.security.krb5.kdc	It defines the host name on which the Active Directory server runs.
java.security.krb5.realm	It defines the Microsoft domain in which the Active Directory server runs.

principalDNPrefix	A prefix to add to the username to form the user distinguished name.
principalDNSuffix	A suffix to add to the username when forming the user distinguished name. This is useful if you prompt a user for a username and you don't want the user to have to enter the fully distinguished name.
roleAttributeID	It defines the role attribute of the context that corresponds to the name of the role. If the roleAttributeIsDN property is set to true, this property is the DN of the context to query for the roleNameAttributeID attribute. If the roleAttributeIsDN property is set to false, this property is the attribute name of the role name.
roleAttributeIsDN	It defines if the role attribute contains the fully distinguished name of a role object or the role name. If false, the role name is taken from the value of the user's role attribute. If true, the role attribute represents the distinguished name of a role object. The role name is taken from the value of the roleNameAttributeId attribute of the corresponding object. In certain directory schemas (for example, Microsoft Active Directory), role (group)attributes in the user object are stored as DNs to role objects and not as simple names. In such case, set this property to true. The default value of this property is false.
roleFilter	It defines a search filter used to locate the roles associated with the authenticated user. The input username/userDN as obtained from the login module callback substitutes the {0} expression in the filter definition. The authenticated userDN substitutes the {1} in the filter definition. An example search filter that matches the input username is (member={0}). An alternative that matches the authenticated userDN is (member={1}). <b>If you omit the roleFilter attribute, the role search will use the UserDN as the DN to obtain the roleAttributeID value.</b>
roleNameAttributeID	It defines the role attribute of the context which corresponds to the name of the role. If the roleAttributeIsDN property is set to true, this property is used to find the name attribute of the role object. If the roleAttributeIsDN property is set to false, this property is ignored.
rolesCtxDN	The fixed DN of the context to search for user roles. Consider that this is not the Distinguished Name of where the actual roles are; rather, this is the DN of where the objects containing the user roles are (e.g. for active directory, this is the DN where the user account is)
searchScope	sets the search scope to one of the following (the default value is SUBTREE_SCOPE): OBJECT_SCOPE - searches the named roles context only. ONELEVEL_SCOPE - searches directly in the named roles context. SUBTREE_SCOPE - searches only the object if the role context is not a DirContext?. If the roles context is a DirContext?, the subtree rooted at the named object and the named object itself are searched.
searchTimeLimit	It defines the timeout for the user and role searches in milliseconds (defaults to 10000, that is 10 seconds).

## Maintain database

---

Follow these recommendations for maintaining your TeamForge database.

Context for the current task

### Schedule data extraction for reporting

Set the interval at which you want your TeamForge site's data extracted to the datamart from which reports are generated.


Each extract-transform-load (ETL) job consists of extracting the data from the production database, transforming it to support reporting, and loading it into the datamart.



By default, this is done every night at 2:30 a.m., by the host's local clock.

1. Open the `site-options.conf` file, the master configuration file that controls your TeamForge site.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

 **Note:** `vi` is an example. Any \*nix text editor will work.

2. Set the `ETL_JOB_TRIGGER_TIME` variable to the interval at which you want ETL jobs to run. For example, a value of `0 0/15 * * * ?` will run an ETL job every 15 minutes.
3. Review the variables you have changed, then save the `site-options.conf` file.

## Raise the logging visibility of selected database requests

For easier troubleshooting, you can dictate that certain database requests get logged in a handy central log file.

For example, database requests that run longer than 10 seconds are likely candidates for troubleshooting. You can have such requests automatically logged in the `vamessages.log` file for your inspection. The exact length of time after which a request becomes problematic depends on your environment.

How it works:


- All database queries are logged at DEBUG level by default.
- By default, the `vamessages.log` file is configured to include all events logged at the INFO level or higher.
- Database queries that run over a configurable time limit are logged at INFO rather than DEBUG, which causes them to appear in `vamessages.log`.

1. Stop TeamForge.

```
Stop Apache on RHEL/CentOS: /etc/init.d/httpd stop
Stop PostgreSQL on RHEL/CentOS: /etc/init.d/postgresql-9.6 stop
Stop TeamForge: /etc/init.d/collabnet stop
```

2. Open the `site-options.conf` file, the master configuration file that controls your TeamForge site.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

 **Note:** `vi` is an example. Any \*nix text editor will work.

3. In the `site-options.conf` file, change the value of the `LOG_QUERY_TIME_THRESHOLD` variable to a value, in milliseconds, that makes sense for your environment.
4. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

5. Start TeamForge.

```
/etc/init.d/httpd start
/etc/init.d/postgresql-9.2 start
/etc/init.d/collabnet start
```

## Optimize PostgreSQL with vacuum

To optimize your PostgreSQL database, run a built-in utility called "vacuum."

Normal use of database software often creates data overhead that needs to be cleaned periodically in order to ensure optimal speed and stability. This overhead is usually the result of temporary files and indexes that the database creates (analogous to a fragmented hard disk.)

The vacuum utility runs on a live database and, like the backup command, can be scripted to run nightly or at minimal server load times.

1. To vacuum the CollabNet TeamForge database, run the `vacuum` command as the PostgreSQL user.

```
vacuumdb -h `hostname` -U <DATABASE_USERNAME> -z <DATABASE_NAME>
```

For example, using the default values in the `site-options.conf` file:

```
vacuumdb -h `hostname` -U ctfuser -z ctfdb
```

2. To set up automatic vacuuming of the database based on activity statistics, set up auto-vacuuming according to these instructions:

<http://www.postgresql.org/docs/8.2/interactive/routine-vacuuming.html#AUTOVACUUM>.

## Move the datamart (dedicated database server mode)

The setup with all services on one server includes having the datamart in the same PostgreSQL instance as TeamForge as well as running it on a separate instance. In either case, you can move the PostgreSQL datamart from its own instance to a separate box in the dedicated database server mode.

1. Stop TeamForge on the app box.

If you have a distributed setup, stop TeamForge on all other servers as well.

```
/etc/init.d/collabnet stop
```


2. Create a dump of the PostgreSQL datamart.

```
su - postgres -c 'pg_dump -C -p <database-port > <reports-database-name > <path-to-dump-file>
```

3. Create a new datamart instance using the `datamart-pgsql-setup.sh` or use the following commands on the database server.

```
su - postgres
initdb -D /var/lib/pgsql/9.6/reports
```

4. Set the `REPORTS_DATABASE_PORT` token in `site-options.conf` file.

 **Note:** The port should use the same value as specified in `postgresql.conf` as specified in the previous step. The recommended value is 5632.

5. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.


6. Restore the datamart into the new instance.

```
/opt/collabnet/teamforge/runtime/scripts/db.py -a restore -t reporting -f <dump-location>
```

7. Copy the `postgresql-9.6` script from runtime scripts and replace `/etc/init.d/postgresql-9.6`

8. Restore the datamart from the database box.

```
su - postgres -c 'psql -p <reports-database-port> <path-to-dump-file>
```

 **Note:** Restart the PostgreSQL service. If any warning messages are displayed, kill the service and start again.

9. Start the PostgreSQL service on the database server.

```
/opt/collabnet/teamforge/bin/teamforge start -s postgres
```

10. Start the services on all the servers.

```
/opt/collabnet/teamforge/bin/teamforge start
```

11. Check your old reports for data accuracy.

12. Permanently remove the old datamart from the TeamForge instance.

```
su - postgres -c 'dropdb <datamart-name> -p <database-port>'
```

## Protect your CollabNet TeamForge site

---

You can take various measures to maximize the security of your CollabNet TeamForge users.

### Encrypt communication between the application and database servers

To prevent your data from being exposed in a readable format on the network, use Secure Socket Layer (SSL) to encrypt the network traffic between the application and the database servers.

If you have a dedicated server for your database (operational database or datamart database), encrypt the data traffic between the application and database servers and between the ETL and datamart servers.

 **Important:** The following steps are applicable only in a multi-box installation setup.


#### Log onto the server as root user always

1. Stop TeamForge on all the servers.

```
/opt/collabnet/teamforge/bin/teamforge stop
```

2. Add the following site option tokens in all the TeamForge servers.

- a) If the operational database is running on a separate server, include the token *DATABASE\_SSL=on*.
- b) If the datamart database is running on a separate server, include the token *REPORTS\_DATABASE\_SSL=on*

 **Note:** It is mandatory to include the tokens specified above in all the servers.

3. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

4. Start TeamForge.

```
/opt/collabnet/teamforge/bin/teamforge start
```

5. Verify that your PostgreSQL database is running in the SSL mode.

- a) Log in to the database server.
- b) Run the following command:

```
grep "ssl = " var/lib/pgsql/9.6/data/postgresql.conf
```

```
Observe:"ssl = on"
```

## Set up SELinux

If SELinux is active on the server that runs your TeamForge site, configure it to allow the services that TeamForge requires.

In case of same hardware upgrade using RHEL/CentOS, it is recommended to upgrade the OS to RHEL/CentOS 7.3 or later versions.



**Attention:** Installing or upgrading TeamForge needs root privileges. You must log on as root or use a root shell to install or upgrade TeamForge.

1. Verify SELinux mode using "getenforce" command. Do this if you have SELinux running in "disabled" mode:

a) Stop TeamForge.



**Note:** Stop TeamForge on all the servers in a distributed setup.

```
/opt/collabnet/teamforge/bin/teamforge stop
```

b) Edit the file `/etc/sysconfig/selinux` and set `SELINUX=enforcing`.

c) Turn off TeamForge startup on boot.

```
chkconfig collabnet off
```

d) Reboot the server and verify if SELinux is set to "enforcing" mode.

```
getenforce
```

e) Run the `grep` command to get the value of `rb_dir` and `rd_data_dir`.

```
grep rb_dir /etc/reviewboard.properties
grep rb_data_dir /etc/reviewboard.properties
```

f) Run the following commands if Review Board is integrated with TeamForge.

```
semanage fcontext -a -t httpd_sys_rw_content_t "<rb_data_dir>(/.*)?"
restorecon -R -v <rb_data_dir>
semanage fcontext -a -t httpd_sys_rw_content_t "<rb_dir>(/.*)?"
restorecon -R -v <rb_dir>
semanage fcontext -a -t httpd_sys_rw_content_t "/opt/collabnet/teamforge/var/home/apache(/.*)?"
restorecon -R -v /opt/collabnet/teamforge/var/home/apache
```

g) Before deploying services, remove `pdk-root` and `pdk-sf-admin` directories.

```
rm -rf /tmp/pdk-root
rm -rf /tmp/pdk-sf-admin
```

h) Turn on TeamForge startup on boot.

```
chkconfig collabnet on
```

i) Run the following command to switch SELinux to 'permissive' mode.

```
setenforce 0
```

j) Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



**Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

k) Start TeamForge.

```
/opt/collabnet/teamforge/bin/teamforge start
```

- l) Run the TeamForge post installation script.


```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

- m) Run the following command to switch SELinux to 'enforcing' mode.

```
setenforce 1
```

- n) Fix the SELinux data permissions.

```
cd /opt/collabnet/teamforge/runtime/scripts/
./fix_data_selinux_permissions.sh
```

 **Note:** If your data directory is on a NetApp volume, you may observe some warning messages which can be ignored safely.

2. Verify SELinux mode using "getenforce" command. Do this if you have SELinux running in "permissive" mode:

- a) Set SELinux to run in "enforcing" mode again.

```
setenforce 1
```

- b) Run the grep command to get the value of httpd\_user and httpd\_group.

```
grep rb_dir /etc/reviewboard.properties
grep rb_data_dir /etc/reviewboard.properties
```

- c) Run the following commands if Review Board is integrated with TeamForge.

```
semanage fcontext -a -t httpd_sys_rw_content_t "<rb_data_dir>(/*)?"
restorecon -R -v <rb_data_dir>
semanage fcontext -a -t httpd_sys_rw_content_t "<rb_dir>(/*)?"
restorecon -R -v <rb_dir>
semanage fcontext -a -t httpd_sys_rw_content_t "/opt/collabnet/
teamforge/var/home/apache(/*)?"
restorecon -R -v /opt/collabnet/teamforge/var/home/apache
```

- d) Restart TeamForge.

```
/opt/collabnet/teamforge/bin/teamforge restart
```

- e) Fix the SELinux data permissions if and only if your site has been upgraded to TeamForge 17.4. Skip this step otherwise for fresh installation of TeamForge 17.4.

```
cd /opt/collabnet/teamforge/runtime/scripts/
./fix_data_selinux_permissions.sh
```

## Protect your TeamForge site with SSL


Use Secure Socket Layer (SSL) to run your Web server securely.

### Set up SSL for your TeamForge site

SSL encryption (HTTPS) is enabled by default in TeamForge 17.1 and later.

1. Back up your existing /etc/httpd/conf/httpd.conf file.
2. Open the site-options.conf file, the master configuration file that controls your TeamForge site.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

 **Note:** vi is an example. Any \*nix text editor will work.


3. Set the options to enable SSL for the site.

- a) Set the `SSL_CERT_FILE` variable to the location of the file that contains your site's SSL certificates.

```
SSL_CERT_FILE=www.example.com.crt
```

- b) Set the `SSL_KEY_FILE` variable to the location of the file that contains your site's RSA private keys.

```
SSL_KEY_FILE=www.example.com.key
```

 **Important:** Select a location for your cert file and your key file that is permanent across restarts. Don't use a temp directory that can be wiped out.

4. In the `site-options.conf` file, make sure the value of the `<hostname>:PUBLIC_FQDN` token matches that of your SSL certificate.
5. Rename the `ssl.conf` file (RHEL/CentOS: `/etc/httpd/conf.d/ssl.conf`) as `ssl.conf.old`, if it exists.
6. If you are converting an existing site to use SSL (that is, if your site already has had users accessing it via HTTP and not HTTPS), you must update your site's publishing repository to use the new SSL settings.

To do this, ask your CollabNet support representative for the `fix-publishing-repos-to-ssl.py` script.

7. Stop TeamForge.

```
/opt/collabnet/teamforge/bin/teamforge stop
```

8. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```

 **Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

A new Apache configuration file is created with the information you provided in the `site-options.conf` file. The new file is named `httpd.conf.cn_new`. It contains `VirtualHost` sections for port 80 and port 443. All port 80 requests are redirected to port 443.

When you point your browser at CollabNet TeamForge, it should now automatically redirect to HTTPS.

As your site is configured to run in the SSL mode (`http` to `https`), there is a change in the URI scheme. Run the following post installation script to make your TeamForge integration components function seamlessly.

Run the TeamForge `post-install.py` script.

```
/opt/collabnet/teamforge/runtime/scripts/post-install.py
```

### Generate SSL certificates

To use HTTPS for web traffic, you will need to obtain a valid Apache SSL certificate.


When generating an Apache (`mod_ssl`) SSL certificate, you have two options:

- Purchase a SSL certificate from a certificate authority (CA). Searching the Web for "certificate authority" will present several choices.
- Generate a self-signed certificate. This option costs nothing and provides the same level of encryption as a certificate purchased from a certificate authority (CA). However, this option can be a mild annoyance to some users, because Internet Explorer (IE) issues a harmless warning each time a user visits a site that uses a self-signed certificate.

Regardless of which option you select, the process is almost identical.

1. Know the fully qualified domain name (FQDN) of the website for which you want to request a certificate.


If you want to access your site through `https://www.example.com`, then the FQDN of your website is `www.example.com`.

 **Note:** This is also known as your common name.

2. Generate the key with the SSL `genrsa` command.

```
openssl genrsa -out www.example.com.key 1024
```


This command generates a 1024 bit RSA private key and stores it in the file `www.example.com.key`.

 **Tip:** Back up your `www.example.com.key` file, because without this file your SSL certificate will not be valid.

### 3. Generate the CSR with `SSL req` command.

```
openssl req -new -key www.example.com.key -out www.example.com.csr
```

This command will prompt you for the X.509 attributes of your certificate. Give the fully qualified domain name, such as `www.example.com`, when prompted for `Common Name`.

 **Note:** Do not enter your personal name here. It is requesting a certificate for a webserver, so the `Common Name` has to match the FQDN of your website.

### 4. Generate a self-signed certificate.

```
openssl x509 -req -days 370 -in www.example.com.csr -signkey
www.example.com.key -out www.example.com.crt
```

This command will generate a self-signed certificate in `www.example.com.crt`.

You will now have an RSA private key in `www.example.com.key`, a Certificate Signing Request in `www.example.com.csr`, and an SSL certificate in `www.example.com.crt`. The self-signed SSL certificate that you generated will be valid for 370 days.

## Protect integrations with SSL


If you have registered Secure Socket Layer (SSL) certificates, your site's users can use SSL when they set up an SCM integration server.

If you use certificates that are generated in-house, self-signed, or signed by a non-established Certificate Authority, they must be registered with each client system that will connect to the TeamForge server. Registration consists of importing custom certificates into the Java runtime's global keystore on each server.

 **Important:** This will affect any other Java applications on the server that use the same Java runtime.

#### 1. Collect server certificates from all servers.

On RHEL, CentOS and other RedHat-based distributions, these are contained in `/etc/httpd/conf/ssl.crt/server.crt`.

 **Tip:** Be sure to use exactly this path, as there are other files with similar names, plus server certificates are not really secret, but some other files are. So, files must be copied (e.g., via `scp`) to the same directory, and renamed if necessary to avoid clashes. We recommend that you use the short server name of the corresponding server for this.

#### 2. Locate the Java keystore.

This is `PATH_TO_JAVA/jre/lib/security/cacerts`.

For example, this may be `/usr/local/j2sdk1.4.2_10/jre/lib/security/cacerts`.


#### 3. Locate the Java `keytool` utility.

This is `PATH_TO_JAVA/bin/keytool`

For example, `/usr/local/j2sdk1.4.2_10/bin/keytool`.

#### 4. Import each server certificate into the keystore.

```
PATH_TO_JAVA/bin/keytool -import -keystore PATH_TO_JAVA/jre/lib/security/
cacerts -file <server>.crt -alias <server>
```


 **Note:** Any value is accepted for `server` in `-alias <server>`.

#### 5. At the password prompt, use `changeit`.

Confirm that you trust the certificate by typing `yes`.

#### 6. Verify that all your certificates are added.

```
PATH_TO_JAVA/bin/keytool -list -keystore PATH_TO_JAVA/jre/lib/security/
cacerts |less
```

 **Note:** The list will contain many more certificates. These are top-level CA certificates, provided with Java.

7. If you are running more than one separate server, repeat these steps for each server.
8. Restart TeamForge

Now you can check the **Use SSL** checkbox when creating an SCM integration.

## Logs

---

Here is some additional information about maintaining your log files.

Context for the current task

### Change the location of a log file

To change where log files are written to, edit the `site-options.conf` file and recreate the runtime environment.

1. Stop the site.

```
sudo /opt/collabnet/teamforge/bin/teamforge stop
/opt/collabnet/teamforge/bin/teamforge stop
```

2. In the `/opt/collabnet/teamforge/etc/site-options.conf` file, set the value of the `LOG_DIR` token with the location where you want the log files to be written.

3. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



**Note:**

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge "provision" command fails otherwise.

4. Start the site.

```
/etc/init.d/httpd start
/etc/init.d/postgresql-9.6 start
/etc/init.d/collabnet start
```

All future Apache logs, mail logs, database logs, Java logs, and other logs will be written to the new location.

### Change the logging level on your site

Set the logging level appropriately to enable logging in `vamessages.log`.

1. Edit `$RUNTIME/jboss/bin/jboss-cli.sh` to enable logging in `vamessages.log`.



**Note:** You need **not** restart the site for JBoss to pick up these changes.

2. To change log levels, as a root user, perform the following:

- a) To enable debug logging, run the following command:

```
/subsystem=logging/root-logger=ROOT:change-root-log-level(level=DEBUG)
```

- b) To disable debug logging, run the following command:

```
/subsystem=logging/root-logger=ROOT:change-root-log-level(level=INFO)
```

3. To change log levels using your LDAP credentials, perform the following:



- a) To enable debug logging, run the following command:

```
/subsystem=logging/logger=com.vasoftware:write-attribute (name="level",
value="DEBUG")
```

- b) To disable debug logging, run the following command:


```
/subsystem=logging/logger=com.vasoftware:write-attribute (name="level",
value="INFO")
```

- c) To enable trace logging for LDAP, run the following command:


```
/subsystem=logging/
logger=org.jboss.security.auth.spi.LdapExtLoginModule:add (level=TRACE,handlers=["VAE
```

- d) To disable trace logging for LDAP, run the following command:

```
/subsystem=logging/
logger=org.jboss.security.auth.spi.LdapExtLoginModule:remove ()
```

 **Tip:** The LDAP debug output will be very limited unless you add `<module-option name="throwValidateErrors" value="true"></>` to the entry for the corresponding log-in module.

- Change in log levels will not need any site restart and these changes will not survive the Jboss restart.
- Changes made using CLI will survive a restart but not a runtime recreation.

 **Note:** Changes are immediately saved to `runtime/jboss/standalone/configuration/standalone-full.xml` and the file, `standalone-full.xml` is recreated every time the TeamForge runtime is rebuilt.

- To retain the configurations after the site restart, edit `$RUNTIME/jboss/standalone/configuration/standalone-full.xml`.
- To make the configurations survive the recreate runtime, edit `$SITE_DIR/dist/conf-snippets/jboss/standalone-full.xml.d/20-standalone-full.xml.py`.

## Search index optimization

---

For an optimal use of the search index feature, follow these tips.

Context for the current task

### Rebuild TeamForge search indexes


You can rebuild your site's search index without stopping TeamForge.

Any new objects created during this time will not be immediately indexed, but will be queued until after the re-indexing.

1. Make sure the TeamForge site is up.
2. Run the re-index script.

```
/opt/collabnet/teamforge/runtime/scripts/SearchReindex.py
```

After the script completes, everything is queued for re-indexing. It will take some time to process the re-index requests.

 **Note:** Re-indexing could take several hours, and the index data could double in size.

### Permit big file uploads

When many users store very large files on your site, you may sometimes notice a slowdown in your site's performance. You can reduce the impact of such a use pattern by telling TeamForge not to index files larger than a certain size.

- 👉 **Note:** It's also a good idea to let your users know that the Documents tool in TeamForge is not designed primarily as a storage device. As a best practice, upload documents to make them available for collaboration, not for backup or long-term storage.

1. Open the `site-options.conf` file, the master configuration file that controls your TeamForge site.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

- 👉 **Note:** `vi` is an example. Any \*nix text editor will work.

2. Add the `SEARCH_MAX_FILE_SIZE` parameter and give it a value equal to the maximum size (in bytes) of files to be indexed.

- 👉 **Note:** The default value is 10M. With this value, files larger than 10M are not indexed.

A value of zero or less specifies that there is no limit, which is the same as the default behavior without the variable.

3. Review the variables you have changed, then save the `site-options.conf` file.

## Allow searching of archive files

By default, users can't search the content of archive files uploaded to TeamForge, such as `zip`, `tar`, or `docx` files. If your users need it, you can provide this ability.

Microsoft Office 2007 files, such as files with the `.docx` extension, are archive files. By default their content is not indexed and does not show up in search results. However, information that TeamForge maintains about those documents, such as title, author, description and version, does appear in search results.

If you permit archive searching, watch for performance slowdowns associated with the larger volume of indexing that TeamForge is doing. Depending on your site members' use patterns, the performance cost may or may not be acceptable to your users.

- 👉 **Note:** It's also a good idea to let your users know that the Documents tool in TeamForge is not designed primarily as a storage device. As a best practice, upload documents to make them available for collaboration, not for backup or long-term storage.

1. Open the `site-options.conf` file, the master configuration file that controls your TeamForge site.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

- 👉 **Note:** `vi` is an example. Any \*nix text editor will work.

2. Add the `SEARCH_SUPPRESS_ARCHIVE_SUB_DOCS` parameter and set it to false. This enables files inside archives (such as `.zip`, `.gz`, or `.tar`) to be indexed for search.
3. Review the variables you have changed, then save the `site-options.conf` file.

## SCM

---

These are some additional details you need to know about SCM.

Context for the current task


Task step.

### Provide more than one source control server

To run more than one source control server of the same type on your site, each integration must have its own unique name and data directory.

- 👉 **Note:** No single server can host more than one source control integration of the same type. If you want to have more than one Subversion integration, they must run on separate machines. The same is true for CVS integrations.


1. Manually create the necessary directories on NetApp.

 **Important:** Each directory must have its own unique name. For example, if the first Subversion instance is named `svnroot`, you might name the second instance `svnroot_2`.

2. Set the permissions on the new directories to `ctf-admin:ctf-admin`
3. For each such directory on the NetApp, create a separate symlink in the local filesystem pointing to the NetApp mount folder.

For example, assuming the NetApp mount is mounted on the `/shared` mount point in the local filesystem:

```
sudo ln -s /svnroot /shared/svnroot_2
```

 **Note:** Only one source code integration of any one type can run on a machine.

## Email

---

These are additional details you can follow while configuring your email settings.


Context for the current task

### Removing users from monitoring objects

As a site or project administrator, if one or more users are no longer project members, you can remove them from monitoring selected TeamForge objects they once subscribed for monitoring.

However, you cannot remove a user from the monitoring list if the user is *monitoring applications* such as trackers, documents, tasks, and so on instead of individual TeamForge objects.

By default, this feature is disabled. To enable this feature, set the `USER_MONITORING_REMOVE_ENABLED` variable to true in the `site-options.conf` file.

 **Note:** Every user removal operation is being logged in the database for audit purposes.


1. Go to the item, from which you want to remove users from monitoring.
2. Select users to remove from monitoring list.
  - a) If you want to remove one or more users from monitoring one of the items, select the item, then click **Monitor > Users Monitoring Selected**. The **Users Monitoring This Item** window appears.
  - b) If you want to remove one or more users from monitoring more than one item, select all the items, then click **Monitor > Users Monitoring This Folder**.  
In the case of team monitoring, click **Monitor > Users Monitoring This Team**.
  - c) In the following window, select one or more check boxes corresponding to the users you want to remove from monitoring.
3. Click **Remove**.  
The `Are you sure you want to remove the selected user(s) from monitoring?` message appears.
4. Click **OK**.

The selected users are removed from monitoring the selected object. An e-mail notification is sent to all active users that are removed from monitoring selected objects.

### Limit the size of message attachments


To avoid overtaxing your mail server or your storage volume, you may want to set a ceiling on the size of the attachments that users can send to a forum via email.

When a user sends an attachment that is larger than the limit, the message is rejected and the user gets an email from the Site Administrator explaining that the attachment exceeded the limit.

-  **Tip:** Before imposing a file attachment size limit, it's a good idea to point your users to better ways of collaborating around large files. Consider suggesting source code repositories, backup systems, or other appropriate solutions.

1. Open the `site-options.conf` file, the master configuration file that controls your TeamForge site.


```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

-  **Note:** `vi` is an example. Any \*nix text editor will work.

2. Set the value of the `DISCUSSION_MAX_ATTACHMENT_SIZE` variable to a reasonable level. For example, if your users are given to using Microsoft Word documents on the site, you might set `DISCUSSION_MAX_ATTACHMENT_SIZE` to 10 MB, and increase the value by two or three MB at a time if users need more headroom.
3. Review the variables you have changed, then save the `site-options.conf` file.


## Limit the size of document attachments

When many users store very large documents on your site, you may sometimes notice a slowdown in your site's performance. You can reduce the impact of such a use pattern by telling TeamForge not to attach documents larger than a certain size.

-  **Note:** It's also a good idea to let your users know that the Documents tool in TeamForge is not designed primarily as a storage device. As a best practice, upload documents to make them available for collaboration, not for backup or long-term storage.

1. Open the `site-options.conf` file, the master configuration file that controls your TeamForge site.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

-  **Note:** `vi` is an example. Any \*nix text editor will work.

2. Add the `DOCUMENT_MAX_FILE_UPLOAD_SIZE` parameter and give it a value equal to the maximum size (in megabytes) of documents to be uploaded.
3. Review the variables you have changed, then save the `site-options.conf` file.

## Discussions

---

Here is some additional information on discussions.

Context for the current task


### Who can post to discussions by email?

To help reduce the risk of spam or other mischief, you may need to limit the users who can post to discussion forums by email.

To leverage the advantages of community collaboration, you should keep your forums as open as you can. However, some sites require tighter control over who can participate in discussions. TeamForge enables you to balance openness against privacy along a spectrum of choices.

1. Open the `site-options.conf` file, the master configuration file that controls your TeamForge site.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

-  **Note:** `vi` is an example. Any \*nix text editor will work.

2. Add the `DISCUSSION_EMAIL_POSTING` variable, and give it the value that reflects the degree of openness your site's discussion forums require.

Use one of these values:

Value	Description
0	Allow only forum admins.
1	Users with roles and permissions.
4	All logged in users.
5	Allow known email addresses only.
6	Allow all site users and guests.

- Review the variables you have changed, then save the `site-options.conf` file.

The value you set here determines the maximum degree of openness to email posting for all projects on your site. For example, consider a site where project members can post by email (level 3). For a project that requires extra security, the project administrator can choose to accept email only from users with the appropriate role (level 1). However, a project owner cannot accept email posts from a less restrictive category of users, such as all users who are logged in (level 4).


## Who can monitor discussions?

You may need to limit the users who can monitor discussion forums.

To leverage the advantages of community collaboration, you should keep your forums as open as you can. However, some sites require tighter control over how users keep track of discussions. TeamForge enables you to balance openness against privacy along a spectrum of choices.

- Open the `site-options.conf` file, the master configuration file that controls your TeamForge site.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

 **Note:** `vi` is an example. Any \*nix text editor will work.

- Add the `DISCUSSION_EMAIL_MONITORING` variable, and give it the value that reflects who you want to get updates on discussions via email.

Use one of these values:

Value	Description
0	Allow only forum admins.
1	Users with role permissions.
4	All logged in users.
5	Allow all site users and guests.

- Review the variables you have changed, then save the `site-options.conf` file.


The value you set here determines the maximum degree of openness to monitoring discussions for all projects on your site. For example, consider a site where project members can monitor discussions (level 3). For a project that requires extra security, the project administrator can dictate that only users with the appropriate role can monitor discussions (level 1). However, a project owner cannot allow monitoring for a less restrictive category of users, such as all users who are logged in (level 4).

## Reduce discussion spam

You can filter out some kinds of spam from your project's discussion forums.

- Open the `site-options.conf` file, the master configuration file that controls your TeamForge site.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```

 **Note:** `vi` is an example. Any \*nix text editor will work.

2. Add one or more MIME types to the Reject MIME types filter.

The presence of any of these MIME types in an incoming message (via email) will cause its rejection with appropriate notification to the posting user.

For example: `DISCUSSION_REJECT_MIME_TYPES=application/pdf,text/xml`

3. Add one or more MIME types to the Drop MIME types filter.

The presence of any of these MIME types in an incoming message (via email) will cause its attachment to be deleted with appropriate notification to the posting user.

For example: `DISCUSSION_DROP_MIME_TYPES=image/jpeg,image/jpg,text/xml`

4. Add one or more header names to the remove headers filter.

If an incoming email posting contains any of these headers, they will be quietly removed from the message before it is archived and before subscribers are notified.

For example: `DISCUSSION_REMOVE_HEADERS=precedence,x-no-archive,Return-Path`

5. Add one or more header names to reject headers filter to be rejected or moderated (if discussion is moderated).


Use regular expressions, each regular expression must match an entire header. The match of any of these headers in an incoming message (via email) will cause its rejection with appropriate notification to the posting user.

For example: `DISCUSSION_REJECT_HEADERS=(?s).*headername1:value2.*,(?s).*name2:value2.*`

6. Add one or more entries for Reject content filter.

Use regular expressions, each regular expression must match an entire entry. The match of any of these entries in discussion body and subject of an incoming message (via email) will cause its rejection with appropriate notification to the posting user.

For example: `DISCUSSION_REJECT_CONTENT=(?s).*word.*,(?s).*spam.*`

 **Note:** The content entry is a case sensitive.

7. Review the variables you have changed, then save the `site-options.conf` file.
8. Restart the site.

## General

---

You can do these things to help maximize the effectiveness and productivity of your site's users.

Context for the current task

Task step.

### Change your site's domain name

To change the domain name of your site, you must change the name in the database and any integrated applications such as the Review Board. TeamForge provides a script for each of these.

1. Make sure your site's database and file system are backed up.
2. In the `site-options.conf` file, update the domain name (and hostname if needed), then save the file.

```
DOMAIN_<Host_Name>=<Domain_Name>
JAMES_POSTMASTER_EMAIL=root@<Domain_Name>
JAMES_MTA_HOST=<Domain_Name>
```

For example:

```
DOMAIN_mybox.supervillain.net=www.smileyface.com
```

```
JAMES_POSTMASTER_EMAIL=root@www.smileyface.com
JAMES_MTA_HOST=www.smileyface.com
```

### 3. Deploy services.

```
/opt/collabnet/teamforge/bin/teamforge provision
```



#### Note:

The "provision" command prompts for response before it bootstraps or migrates data during TeamForge installation and upgrade respectively. Enter "y" or "N" to proceed. For more information, see the [TeamForge script](#).

TeamForge 17.4 (and later) installer expects the system locale to be LANG=en\_US.UTF-8. TeamForge "provision" command fails otherwise.

### 4. Go to the TeamForge scripts directory.

```
cd /opt/collabnet/teamforge/runtime/scripts/
```

### 5. Run the script to change the domain in the database.

```
./domain_change_db.py --old=www.myoldsitename.net --
new=www.mynewsitename.net > /tmp/domain_change_db.out 2>&1
```

### 6. If your site has Review Board integrated, run the following commands:

#### a) Start the TeamForge application services.

```
/opt/collabnet/teamforge/bin/teamforge start
```

#### b) Update the install.conf file with the new domain name.

```
vi /opt/collabnet/RBInstaller-17.4.7/installer/install.conf
```

#### c) Recreate the runtime on Review Board.

```
sudo python ./install.py -r
```

#### d) Restart the TeamForge application services.

```
setenforce 1
/opt/collabnet/teamforge/bin/teamforge restart
```

## Let users see what's in a project template

Help your site's project administrators choose a project template by enabling them to see the contents of the templates that are available.

By default, only site administrators can see project template detail, but project administrators normally create a project from a project template. To choose the right template, a project manager may want to know if tasks, documents, wiki pages or other kinds of content are included in a given template.

### 1. Open the site-options.conf file, the master configuration file that controls your TeamForge site.

```
vi /opt/collabnet/teamforge/etc/site-options.conf
```



**Note:** vi is an example. Any \*nix text editor will work.

### 2. Change the value of the SHOW\_PROJECT\_TEMPALTE\_DETAIL\_TO\_NON\_SITEADMINS variable to true.



**Note:** Use the parameter name as given, including the typo.

### 3. Review the variables you have changed, then save the site-options.conf file.

### 4. Recreate the runtime environment.

```
./install.sh -V -r -d /opt/collabnet/teamforge
```

# Reference information about TeamForge

---

Use this reference information to get deeper detail on configuration files, logs, scripts and other resources you use to administrate TeamForge.

## Configuration files in TeamForge

---

Edit these configuration files to get the behavior you want.

### site-options.conf

TeamForge is controlled by settings in a master configuration file called `site-options.conf`. Some of the most useful configuration settings you can specify in the `site-options.conf` file are described here.

The `site-options.conf` file resides in the `/opt/collabnet/teamforge/etc/` directory.

When TeamForge starts, it reads and implements the instructions provided as values to the variables in this file.


The basic procedure for configuring your TeamForge site, then, is to edit the `site-options.conf` file, supply a valid value for the token of interest, save the file, and rebuild the TeamForge runtime environment.

### site-options.conf changes in TeamForge 17.1

TeamForge 17.1 includes major changes to `site-options.conf` configuration. The syntax for writing the `HOST` and `DOMAIN` tokens and the names of certain services have been changed in TeamForge 17.1. This section walks you through the list of changes and what it takes to build your TeamForge 17.1 site's master configuration file.

#### List of services and their identifiers

The following table lists the names of, both old (TeamForge 16.10 and earlier) and new (TeamForge 17.1 and later), TeamForge services. The `ctfcore` service, including its dependent services such as `search`, `mail`, `ctfcore-database` and `codesearch`, are mandatory services. You can have all mandatory services run on a single server or distribute them across several servers. All other services such as `etl` and `subversion` are optional. However, if you install an optional component, for example `etl`, it is mandatory to install services such as `ctfcore-datamart` too, which is required by `etl`.

 **Notice:** To ensure backward compatibility, TeamForge 17.1 supports both old and new syntaxes for defining your `HOST` and `DOMAIN` tokens. Similarly, TeamForge 17.1 supports both old and new service names. However, going forward, it is recommended to adjust your `site-options.conf` in line with these changes as support for older syntax and naming convention would be dropped in due course.

Service	Also requires (either on the same or separate server)	Mandatory/Optional	Old name	Description
<code>ctfcore</code>		Mandatory	<code>app</code>	Main TeamForge application server
	<code>search</code>	Mandatory	<code>indexer</code>	Indexing and searching
	<code>mail</code>	Mandatory	NA (added in TeamForge 17.1)	Email server
	<code>ctfcore-database</code>	Mandatory	<code>database</code>	Operational database
	<code>codesearch</code>	Mandatory	<code>codesearch</code>	Code search



Service	Also requires (either on the same or separate server)	Mandatory/Optional	Old name	Description
etl		Optional	etl	ETL for datamart
	ctfcore-datamart	Mandatory if you have etl installed	datamart	Datamart database
subversion		Optional	subversion	SVN version control
cvs		Optional	cvs	CVS version control
gerrit		Optional	gerrit	Git/Gerrit version control
	gerrit-database	Mandatory if you have Git/Gerrit installed	NA (added in TeamForge 17.1)	Database for Git/Gerrit. In a distributed setup, add this identifier on the server where you want to run Gerrit database.
binary		Optional	binary	Artifact repository integration
	binary-database	Mandatory if you have binary installed	NA (added in TeamForge 17.1)	Database for artifact repository integration. Binary app (binary) and database (binary-database) have to be installed on the same server.
reviewboard		Optional	reviewboard	Review Board code review tool
	reviewboard-database	Mandatory if you have Review Board installed	NA (added in TeamForge 17.1)	Database for Review Board. In a distributed setup, add this identifier on the server where you want to run Review Board database.

## SERVICES and PUBLIC\_FQDN site-options.conf tokens

### Syntax to define the services running on a TeamForge host

The syntax for defining the services running on a TeamForge host is:

```
<hostname>:SERVICES = list of services separated by space
```

Where <hostname> can be "localhost" or the server name as returned by the `hostname` command on the console. The latter is recommended as this allows reuse of the same `site-options.conf` file across all servers in a distributed setup. Here are some examples:

Example 1: Default single-server setup

```
localhost:SERVICES = ctfc core ctfc core-database ctfc core-datamart etl mail
search codesearch subversion
```

#### Example 2: Single-server setup with Git integration

```
localhost:SERVICES = ctfc core ctfc core-database ctfc core-datamart etl mail
search codesearch subversion gerrit gerrit-database
```

#### Example 3: Single-server setup with Review Board integration

```
localhost:SERVICES = ctfc core ctfc core-database ctfc core-datamart etl mail
search codesearch subversion reviewboard reviewboard-database
```

#### Example 4: Three-server setup with Git and Binary integration

```
server01:SERVICES = ctfc core etl mail search codesearch binary
server02:SERVICES = subversion gerrit
server03:SERVICES = ctfc core-database ctfc core-datamart gerrit-database
```

### Syntax to define the domain names for a TeamForge site

Assign a public FQDN (optional, but strongly recommended). Make sure there is a DNS A or CNAME record for this FQDN.

```
server01:PUBLIC_FQDN = teamforge.example.com
server02:PUBLIC_FQDN = scm.example.com
```

### Service-specific FQDNs

Installing TeamForge with service-specific FQDNs (instead of machine-specific host/domain names) is highly recommended so that you will be able to change the system landscape at a later point in time without having any impact on the URLs (in other words, end users do not have to notice or change anything). For example, you can create FQDNs specifically for services such as Subversion, Git, mail, Codesearch and so on.

- All such service-specific FQDNs must be long to a single sub domain and it is recommended to create a new sub domain for TeamForge.
- A wildcard SSL cert is required if you are using service-specific FQDNs. SNI SSL cert cannot be used.
- When SSL is enabled and no custom SSL-certificates are provided, a self-signed wildcard cert is generated for the sub domain.
- When SSL is enabled and a custom SSL-certificate is provided, the CN of the certificate is verified to be a wildcard CN.

The following configuration shows the list of services and FQDNs for specific services such as `ctfc core`, `subversion`, `gerrit` and `mail` in a single server TeamForge installation:

```
localhost:SERVICES = ctfc core ctfc core-database ctfc core-datamart etl mail
search codesearch subversion
localhost:PUBLIC_FQDN = app.forge.collab.net
localhost:ctfc core:PUBLIC_FQDN = ctf.forge.collab.net
localhost:subversion:PUBLIC_FQDN = svn.forge.collab.net
localhost:gerrit:PUBLIC_FQDN = git.forge.collab.net
localhost:mail:PUBLIC_FQDN = mail.forge.collab.net
```

In a single server setup, all these domain names point to a single server. However, when services are later distributed across multiple servers, all it takes to avoid an end user impact is to adjust these domain names to point to different servers.

### Site options change log

Change log of `site-options.conf` tokens.

- 👉 **Important:** While upgrading to a latest TeamForge release, make sure that obsolete site option tokens, if any, are removed from the `site-options.conf` file of the TeamForge version you are upgrading to.

## TeamForge 17.4

### Obsolete tokens

- GERRIT\_FORCE\_HISTORY\_PROTECTION
- ORC\_HOSTNAME
- ORC\_SSL\_CA\_CERT\_FILE
- ORC\_PORT
- ORC\_PROXIED\_PATH
- ORC\_PROTOCOL
- ORCHESTRATE\_ENABLED
- POSTGRES\_INTERFACE\_IP
- POSTGRES\_INTERFACE

### New tokens

- [LISTEN\\_IP](#)

### Default values for site-options.conf tokens

Default values have been assigned to the following `site-options.conf` tokens and are therefore removed from the default TeamForge 17.4 `site-otptions-default.conf` file.



**Warning:** If you are upgrading from TeamForge 17.1 (or earlier) to TeamForge 17.4 (or later) and if you have been using your own values for the following tokens, you must make sure you use the same values in your `site-options.conf` file post upgrade to TeamForge 17.4.

```
# CTF Core
SCM_ADMIN_PASSWORD=$auto$
IAF_DBPASS=$auto$
SCM_DEFAULT_SHARED_SECRET=$auto$
SOAP_ANONYMOUS_SHARED_SECRET=$auto$

# Database
DATABASE_USERNAME=teamforge
DATABASE_NAME=teamforge
DATABASE_PASSWORD=$auto$
DATABASE_READ_ONLY_USER=teamforge_reader
DATABASE_READ_ONLY_PASSWORD=$auto$

# Datamart
REPORTS_DATABASE_USERNAME=teamforge_datamart
REPORTS_DATABASE_NAME=teamforge_datamart
REPORTS_DATABASE_PASSWORD=$auto$
REPORTS_DATABASE_READ_ONLY_USER=teamforge_datamart
REPORTS_DATABASE_READ_ONLY_PASSWORD=$auto$

# ETL
ETL_SOAP_SHARED_SECRET=$auto$

# Gerrit
GERRIT_DATABASE_PASSWORD=$auto$

# RabbitMQ
RABBITMQ_APP_ADMIN_USER=guest
RABBITMQ_APP_ADMIN_PASSWORD=$auto$
RABBITMQ_APP_CTF_USER=ctf
RABBITMQ_APP_CTF_PASSWORD=$auto$
```

```

RABBITMQ_APP_SERVICES_USER=eventq
RABBITMQ_APP_SERVICES_PASSWORD=$auto$
RABBITMQ_PERMISSION_PUBLISHER=permission_publish
RABBITMQ_PERMISSION_PUBLISHER_PASSWORD=$auto$

# Mongo
MONGODB_APP_DATABASE_NAME=eventq
MONGODB_ADMIN_DATABASE_NAME=admin
MONGODB_APP_ADMIN_USER=admin
MONGODB_APP_ADMIN_PASSWORD=$auto$
MONGODB_APP_BACKUP_USER=backup
MONGODB_APP_BACKUP_PASSWORD=$auto$
MONGODB_APP_SERVICES_USER=eventq
MONGODB_APP_SERVICES_PASSWORD=$auto$

# HAProxy
HAPROXY_STATS_PASSWORD=$auto$

# Binary
IAF_DBNAME=iafdb
IAF_DBUSER=iafdbusr

# James Admin (for management
  interface)
JAMES_ADMIN_USER=admin
JAMES_ADMIN_PASSWORD=$auto$

```

## TeamForge 17.1

### Obsolete tokens

- BDCS\_ADMIN\_PASSWORD
- BDCS\_ADMIN\_USERNAME
- BDCS\_HOST
- BDCS\_SSL
- BDCS\_TOMCAT\_PORT
- BDCS\_SDK\_SEARCH\_LIMIT\_MAX
- BDCS\_SSL\_CERT\_FILE
- BDCS\_SSL\_KEY\_FILE
- BDCS\_SSL\_CA\_CERT\_FILE
- BDCS\_SSL\_CHAIN\_FILE
- BDCS\_SCAN\_SOURCE\_DIR\_ROOT
- BDCS\_INSTALL\_PATH
- BDCS\_PGSQL\_HOME\_DIR\_ROOT
- BDCS\_PGSQL\_PORT
- BDCS\_TOMCAT\_MX\_IN\_MB
- BDCS\_TOMCAT\_SHUTDOWN\_PORT

### Newly added tokens

- [ELASTICSEARCH\\_JAVA\\_OPTS](#) on page 746

## TeamForge 16.10

### Obsolete tokens

- SSH\_TUNNEL\_ENABLED
- SELINUX\_SETUP

**Newly added tokens**

- SELINUX\_ENABLED
- [DISABLE\\_REMOTE\\_PUBLISHING](#) on page 740

**TeamForge 16.7****Obsolete tokens**

- TeamForge 16.7 installer automatically sets *JAVA\_HOME* during installation or upgrade. Therefore, the *JAVA\_HOME* site options token, if added to your *site-options.conf* file, must be removed while upgrading to TeamForge 16.7 and later.
- The *DEDICATED\_INSTALL* token is no longer supported. CollabNet recommends removing this token from the *site-options.conf* file. However, this token is ignored (has no effect whatsoever) if you continue to have it in your *site-options.conf* file post upgrade to TeamForge 16.7 and later.
- Debug settings can be done via the *JAVA\_OPTS* tokens (such as *JBOSS\_JAVA\_OPTS* and *ETL\_JAVA\_OPTS*). Hence, the following tokens are no longer supported:
  - JBOSS\_DEBUG
  - PHOENIX\_DEBUG
  - ETL\_DEBUG
  - INTEGRATION\_DEBUG
  - ETL\_DEBUG\_PORT
  - JBOSS\_DEBUG\_PORT
  - INTEGRATION\_DEBUG\_PORT
  - PHOENIX\_DEBUG\_PORT
- The *SITE\_DIR* and *DATA\_DIR* tokens are no longer supported. Starting from TeamForge 16.7 (and later) *SITE\_DIR* and *DATA\_DIR* are unconditionally set to */opt/collabnet/teamforge* and */opt/collabnet/teamforge/var* respectively during runtime creation.
- The *MODPAGESPEED\_ENABLED* token is no longer supported.
- SSL certificates are validated by default. Hence, *VALIDATE\_SSL\_CERTS* token is no longer supported. If in use, remove this token from the *site-options.conf* file post upgrade to TeamForge 16.7 and later.
- While the *SSL\_CA\_CERT\_FILE* token is still supported, it has been removed from the *site-options-default.conf* file as it is only needed for add-ons. References to this token have been removed from TeamForge documentation as well.

**Newly added tokens**

- [ALLOW\\_CASE\\_INSENSITIVE\\_LOGIN](#).
- [HTTP\\_MAX\\_PARAMETERS](#)

- [ENABLE\\_SITE\\_NEWS](#)
- [LOGROTATE\\_ARCHIVE\\_COUNT](#)
- [MAX\\_PASSWORD\\_LENGTH](#) on page 757

## TeamForge 8.2

### Obsolete tokens

PERFORCE\_CLIENT\_DIR  
 PERFORCE\_GROUP  
 PERFORCE\_LICENSE\_FILE  
 PERFORCE\_LOG\_DIR  
 PERFORCE\_PORT  
 PERFORCE\_REPOSITORY\_BASE  
 PERFORCE\_SERVICE\_CMD  
 PERFORCE\_USER  
 HELP\_AVAILABILITY  
 REMOTE\_HELP\_URL  
 EXTERNAL\_TOMCAT\_INSTALL\_DIR  
 INTEGRATION\_BUILTIN\_TOMCAT  
 ETL\_BUILTIN\_TOMCAT  
 CEE\_COMPATIBLE

### Newly added tokens

[ENABLE\\_CACHING\\_WITH\\_MEMCACHED](#)  
[MEMCACHED\\_SERVER\\_HOST](#)  
[MEMCACHED\\_SERVER\\_PORT](#)  
[MEMCACHED\\_SERVER\\_TTL](#)

## TeamForge 8.0

### Newly added tokens

- [NOTIFY\\_SITE\\_ADMINS\\_FOR\\_SITE\\_ACTIVITIES](#)
- [BINARY\\_SETUP\\_TYPE](#)

## ACTIVITY\_LINKS\_CUSTOMIZATION

When the *ACTIVITY\_LINKS\_CUSTOMIZATION* variable is set to true, the TeamForge Activity Chart and the Most Active Projects List do not appear on the main page before the user logs in.

### Values

true or false

### Default


false

## ADMIN\_EMAIL

The *ADMIN\_EMAIL* variable specifies a valid email address for the site administrator.

The mail account specified must be hosted on a separate server from the TeamForge site server.

The *SYSTEM\_EMAIL*, *ADMIN\_EMAIL*, and *JAMES\_POSTMASTER\_EMAIL* variables can specify the same address.

 **Important:** In TeamForge 6.x, the sender name and address for system-generated emails is taken from the value of the *SYSTEM\_EMAIL* variable. Therefore, changing the admin user's full name or email address does not affect the sender details of system-generated emails. This is different from TeamForge 5.x, in which the

sender name and address for system-generated emails is derived from the admin user's full name and email address.

### Values

Email address specification

### Default

root@{\_\_APPLICATION\_HOST\_\_}

### ALLOW\_CASE\_INSENSITIVE\_LOGIN

In general, TeamForge usernames are validated case-sensitively. Set this token to true so that username validations are done case-insensitively.

### Values

true or false

### Default

false

### ALLOW\_NO\_PASSWORD\_ON\_USER\_CREATION

The *ALLOW\_NO\_PASSWORD\_ON\_USER\_CREATION*, when set to true, allows the admin to create users with null password through SOAP when external authentication is enabled.

### Values

true, false

### Default

false

### ALLOW\_USERNAME\_IN\_PASSWORD

The *ALLOW\_USERNAME\_IN\_PASSWORD*, when set to true, allows users to set a password that includes the string that they use for their user name on the site.

### Values

true, false

### Default

true

### APPLICATION\_LOG\_DIR

The *APPLICATION\_LOG\_DIR* specifies the directory to which the application writes its log files.

### Values

Path specification

### Default

/opt/collabnet/teamforge/log/apps

**APPROVE\_NEW\_USER\_ACCOUNTS**

The *APPROVE\_NEW\_USER\_ACCOUNTS* variable specifies whether a site administrator must approve the requests to join the site.

**Values**

true or false

**Default**

true

**ARTIFACT\_DESC\_EDITOR**

The *ARTIFACT\_DESC\_EDITOR* variable allows you to choose the type of text that can be used for artifact description using the editor tool.

**Values**

Plain Text

**Default**

Plain Text

**ARTIFACT\_LIST\_LIMIT**

The *ARTIFACT\_LIST\_LIMIT* variable specifies the maximum number of artifacts displayed in and exported from the **Planned Tracker Artifacts** tab available in **File Releases**.

**Values**

Integer

**Default**

5000

**AUTO\_DATA**

TeamForge 7.1 and later support automatic password creation. Once you turn on automatic password creation, the *AUTO\_DATA* token is auto-generated by the installer during runtime recreation.



**Restriction:** The *AUTO\_DATA* token is not a user-editable site options token. Do not modify this token manually.



The following password-related `site-options.conf` tokens can have the passwords automatically created if set to **\$auto\$**. When set to **\$auto\$**, the passwords for the tokens are randomly generated, encrypted and stored in the `AUTO_DATA` token that is added automatically to the `site-options.conf` file during runtime recreation.

```
DATABASE_PASSWORD=$auto$
```

```
DATABASE_READ_ONLY_PASSWORD=$auto$
```

```
REPORTS_DATABASE_PASSWORD=$auto$
```

```
REPORTS_DATABASE_READ_ONLY_PASSWORD=$auto$
```

```
ETL_SOAP_SHARED_SECRET=$auto$
```

```
JAMES_ADMIN_PASSWORD=$auto$
```

```
BDCS_ADMIN_PASSWORD=$auto$
```

```
MIRROR_DATABASE_PASSWORD=$auto$ (applicable only if you are mirroring your database)
```



**Caution:**

As this `AUTO_DATA` token is auto-generated and managed by the TeamForge installer, do not modify this token manually during TeamForge upgrades. When you upgrade TeamForge, you must copy this token intact (along with its value) from the old `site-options.conf` file to the upgraded site's `site-options.conf` file, before recreating the runtime. This applies to all the servers in a distributed set up (copy the `AUTO_DATA` token and its value to all the servers).

This feature is enabled by default. You can, however, override any of the above password-related tokens with the password of your choice.

**BCC\_MAIL\_BATCH\_SIZE**

Set `ENABLE_BCC_MONITORING` token to true and then depending on your site's requirements, you can set the value of this `BCC_MAIL_BATCH_SIZE` token to configure the number of monitoring emails delivered in a single delivery.

When a work item such as an artifact or discussion is updated, instead of sending separate monitoring emails to all monitoring users, you can now choose to have just one monitoring email sent with all monitoring users added to the BCC. This reduces the load on the email server and results in faster email delivery. You must enable this feature by setting `ENABLE_BCC_MONITORING` to true and then, depending on your site's requirements, you can optimize the value of this `BCC_MAIL_BATCH_SIZE` token to increase or decrease the number of emails delivered in a single delivery.

**Values**

Integer

**Default**

100

This token was added in TeamForge 7.2.

**BINARY\_SETUP\_TYPE**

Set this token appropriately to have the binary application installed the way you want (as part of the TeamForge installation or upgrade).

TeamForge 8.0 and later support integration with Nexus OSS, an open source repository manager for binary artifacts. By default, the TeamForge installer installs a binary application (referred to as the 'binary app' hereinafter), which is essentially a launching pad for all binary integrated applications such as Nexus OSS. Once this binary app is installed as part of TeamForge 8.0 installation, you can integrate your Nexus servers and repositories with TeamForge.

Though the binary app is installed by default with TeamForge 8.0, you can change the value of this site option token, *BINARY\_SETUP\_TYPE*, to skip binary app installation altogether. You can also have this token configured to have the binary app installed and rolled out for all projects (both existing and new projects to be created) or only for new projects to be created or for select projects on a need basis.

**Values**

- all: Binary app is installed and available for all projects.
- new: Binary app is installed and available for new projects only.
- manual: Binary app is installed at a site level and project administrators can add it to select projects on a need basis.
- none: TeamForge installer skips the binary app installation altogether.

**Default**

**new**

This token was added in TeamForge 8.0.

**DATABASE\_NAME**

The *DATABASE\_NAME* variable specifies the name of the site's database.

**Values**

Alphanumeric string

**Default**

teamforge

**DATABASE\_PASSWORD**

The *DATABASE\_PASSWORD* variable is the password for the Unix user that is authorized to read from and write to the site's database.

**Values**

Alphanumeric string

**Default**

\$auto\$

**DATABASE\_TYPE**

The *DATABASE\_TYPE* token specifies the type of database in which the TeamForge site's data is stored.

**Values**

postgresql or oracle

**Default**

postgresql

**DATABASE\_USERNAME**

The *DATABASE\_USERNAME* variable specifies the Unix user that is authorized to read from and write to the site's database.

**Values**

Alphanumeric string

**Default**

teamforge

**Comments**

For some advanced operations, you may need to log into the database as the database user. However, under normal conditions only the TeamForge site process itself needs to access the database.

**DEFAULT\_LOCALE**

The *DEFAULT\_LOCALE* variable specifies the language in which automated email messages from the site are generated.

**Values****Default**

en

**DEFAULT\_PROJECT\_ACCESS**

The *DEFAULT\_PROJECT\_ACCESS* variable specifies the type of access that is assigned to a project when it is created. A project can be private, public, or gated.

**Values**

private, gated, public

**Default**

private

**DISABLE\_CREATE\_INTEGRATION\_SERVERS**

The *DISABLE\_CREATE\_INTEGRATION\_SERVERS* token specifies whether the creation of new SCM integrations is allowed.

**Values**

true or false

**Default**

false

## Comments

When this token is set to its default value of "false", you can add SCM integration servers to your TeamForge site. Also, the **Discover Subversion Edge Servers** option, which enables you to find and connect to Subversion Edge servers on your LAN, is available.

### DISABLE\_REMOTE\_PUBLISHING

Publishing repository, like the branding repository, is one of the default repositories that's created automatically when a TeamForge project is created and is intended to contain publicly-consumable files. However, site administrators can toggle access to Publishing Repositories and restrict access based on defined RBAC.

For more information on Publishing Repository, see *What is a Publishing repository? How does it work?* on page 1003. This token is set to false by default, meaning the Publishing Repository is publicly accessible.

If set to true:

- The Publishing Repository is stripped of its public access and behaves like any other Subversion repository with RBAC applied to it.
- Project Administrators can no longer view or modify the **PROJECT HOME OPTIONS** (see *Create a custom project home page* on page 141).

### Values

true or false

### Default

false

### DISABLE\_USER\_SELF\_CREATION

The *DISABLE\_USER\_SELF\_CREATION* variable restricts users from creating their own accounts on the TeamForge home page.

### Values

true or false

### Default

true

### DISCUSSION\_ADD\_HEADERS

The *DISCUSSION\_ADD\_HEADERS* variable allows you to add custom headers to the emails posted in the forum.

### Values

You can choose to add or remove headers by specifying the particular information you want to be added or dropped from the header. For example, if you add `<#d#>` in the Add header field, the URL of that discussion will be added to the header of all the available messages in that discussion.

### Default

None

### Example

`DISCUSSION_ADD_HEADERS=headername1:value1, name2: value2 , post-id:<#n#>, forum-url:<#d#>, message-url:<#m#>, domain:<#h#>, list-name:<#l#>, list-address:<#l#>@<#h#>`

## Comments

Add one or more header names. The match of any of these headers in an outgoing message (via email) causes its addition with appropriate notification to the posting user.

### DISCUSSION\_DROP\_MIME\_TYPES

The *DISCUSSION\_DROP\_MIME\_TYPES* variable allows you to delete the mime types submitted by email that contain arbitrary strings.

#### Values

image/jpeg,image/jpg,text/xml

#### Default


Regular expression

#### Example

DISCUSSION\_DROP\_MIME\_TYPES=image/jpeg,image/jpg,text/xml

## Comments

Add one or more mime types to the Drop mime types filter. The presence of any of these mime types in an incoming message (via email) causes its deletion with appropriate notification to the posting user.

 **Note:** If a mime type is specified in both the Reject and Drop mime filters, then the Reject mime type filter must take higher precedence than the Drop mime type filter.

### DISCUSSION\_EMAIL\_MONITORING

The *DISCUSSION\_EMAIL\_MONITORING* variable determines which users can monitor a forum on the site.

#### Values

Value	Description
0	Allow only forum admins.
1	Users with role permissions.
4	All logged in users.
5	Allow all site users and guests.

#### Default

1

#### Example

DISCUSSION\_EMAIL\_MONITORING=4

## Comments

This setting applies to the site as a whole. Project owners can choose to be more restrictive in their own project by selecting a lower value on the project administration page.

### DISCUSSION\_EMAIL\_POSTING

The *DISCUSSION\_EMAIL\_POSTING* variable determines which users on your site can post to forums by e-mail.

**Values**

Value	Description
0	Allow only forum admins.
1	Users with roles and permissions.
4	All logged in users.
5	Allow known email addresses only.
6	Allow all site users and guests.

**Default**

1

**Example**

DISCUSSION\_EMAIL\_POSTING=4

**Comments**

This setting applies to the site as a whole. Project owners can choose to be more restrictive in their own project by selecting a lower value on the project administration page.

**DISCUSSION\_FORUM\_EDITOR**

The *DISCUSSION\_FORUM\_EDITOR* variable allows you to choose the type of text that can be used in discussion forum description using the editor tool.

**Values**

Plain Text

**Default**

Plain Text

**DISCUSSION\_MAX\_ATTACHMENT\_SIZE**

The *DISCUSSION\_MAX\_ATTACHMENT\_SIZE* variable sets an upper limit to the size of files that users can attach to an email message sent to any discussion forum on the site.

**Values**

Integer (Megabytes)

**Default**

blank

**Comment**

A value of zero or less specifies that there is no limit, which is the same as the default behavior without the variable.

**DISCUSSION\_POST\_EDITOR**

The *DISCUSSION\_POST\_EDITOR* variable allows you to choose the type of text that can be used for posting in discussion forums using the editor tool.

**Values**

Plain Text

**Default**

Plain Text

**DISCUSSION\_REJECT\_CONTENT**

The *DISCUSSION\_REJECT\_CONTENT* variable allows you to block the discussion messages submitted by email that contain arbitrary strings.

**Values**

Regular expression

**Default**


None

**Example**

```
DISCUSSION_REJECT_CONTENT=(?s).*word.*(?s).*spam.*
```

**Comments**

Add one or more entries. Each regular expression must match an entire entry. The match of any of these entries in the body or subject of an incoming message (via email) causes its rejection, with appropriate notification to the posting user.

 **Note:** The content entry is case sensitive.

**DISCUSSION\_REJECT\_HEADERS**

The *DISCUSSION\_REJECT\_HEADERS* variable allows you to block different headers submitted by email that contain arbitrary strings.

**Values**

Regular expression

**Default**

None

**Example**

```
DISCUSSION_REJECT_HEADERS=(?s).*headername1:value2.*(?s).*name2:value2.*
```

**Comments**

Add one or more header names. Each regular expression must match an entire header name. The match of any of these headers in an incoming message (via email) causes its rejection, with appropriate notification to the posting user.

**DISCUSSION\_REJECT\_MIME\_TYPES**

The *DISCUSSION\_REJECT\_MIME\_TYPES* variable allows you to delete the mime types submitted by email that contain arbitrary strings.

**Values**

Application/PDF,text/xml

**Default**

Regular expression

**Example**

DISCUSSION\_REJECT\_MIME\_TYPES=application/pdf,text/xml

**Comments**

Add one or more mime types to the Reject MIME types filter. The presence of any of these mime types in an incoming message (via email) will cause its deletion with appropriate notification to the posting user.

**DISPLAY\_TIMEZONE**

The *DISPLAY\_TIMEZONE* token, if set with a preferred time zone, takes precedence over the physical TeamForge server location's time zone and will be the default time zone that's displayed throughout the application. In other words, use this token to set a preferred time zone to display across the TeamForge application in case the TeamForge physical server and users are not on the same time zone.

**Values**

The ID for a time zone can be either a full name such as *America/Los\_Angeles*, or a custom ID in the form GMT[+|-]hh[:mm] such as *GMT-08:00*. It can also be in the form of a three letter abbreviation such as *PST*.

**Default**

If set, this token overrides the default time zone of the TeamForge server. TeamForge uses the default time zone of the JVM otherwise.

**DOCUMENT\_MAX\_FILE\_UPLOAD\_SIZE**

The *DOCUMENT\_MAX\_FILE\_UPLOAD\_SIZE* variable sets an upper limit to the size of the documents that can be attached.

**Values**

Integer (Megabytes)

**Default**

blank

**Comment**

A value of zero specifies that there is no limit, which is the same as the default behavior without the variable.

**DOCUMENT\_TEXT\_EDITOR**

The *DOCUMENT\_TEXT\_EDITOR* variable allows you to choose the type of text that can be used for the document description using the editor tool.

**Values**

Plain Text

**Default**

Plain Text

**ENABLE\_BCC\_MONITORING**

When a work item such as an artifact or discussion is updated, instead of sending separate monitoring emails to all monitoring users, you can now choose to have just one monitoring email sent with all monitoring




users added to the BCC. This reduces the load on the email server and results in faster email delivery. Set `ENABLE_BCC_MONITORING` to true to enable this feature.

### Values

true or false

### Default

false

 **Tip:** Depending on your site's requirements, you can also optimize the value of `BCC_MAIL_BATCH_SIZE` to increase or decrease the number of emails delivered in a single delivery.

This token was added in TeamForge 7.2.

### ENABLE\_CACHING\_WITH\_MEMCACHED

Set the `ENABLE_CACHING_WITH_MEMCACHED` token to true if you want to cache SVN authentication and authorization information in Memcached.

### Values

true or false

### Default

Caching with Memcached is disabled by default in TeamForge.

You must also set up the following tokens if you set `ENABLE_CACHING_WITH_MEMCACHED` token to true:

- `MEMCACHED_SERVER_HOST` on page 756
- `MEMCACHED_SERVER_PORT` on page 756
- `MEMCACHED_SERVER_TTL` on page 756

See [Install Memcached](#) on page 694 for more information.

### ENABLE\_SITE\_NEWS

Site news is disabled by default. Set this token to true in `site-options.conf` file, customize the home page of your site and recreate runtime if you want to publish site news on your site's home page.

**To publish site news:** Set this token to true in `site-options.conf` file, [customize the home page of your site](#) (see "[siteNew](#)'s html block) and recreate runtime if you want to publish site news on your site's home page.

### Values

true or false

### Default

false

This token was added in TeamForge 16.7. Until TeamForge 16.3, regardless of whether you have site news enabled or not, site news were processed in the background. With this `ENABLE_SITE_NEWS` token, there is no site news processing in the background (by default `ENABLE_SITE_NEWS=false`) thereby improving the site's home page performance a bit.

### ENABLE\_UI\_FOR\_CUSTOM\_EVENT\_HANDLERS

To support branding and customization changes, set the `ENABLE_UI_FOR_CUSTOM_EVENT_HANDLERS` variable to true.

### Values

true or false

**Default**

true

**ENFORCE\_MINIMUM\_USERNAME\_LENGTH**

The *ENFORCE\_MINIMUM\_USERNAME\_LENGTH* variable determines the minimum length that can be set for usernames.

**Values**

0-31

**Default**

0

**ELASTICSEARCH\_JAVA\_OPTS**

The *ELASTICSEARCH\_JAVA\_OPTS* token specifies the memory settings for the Java virtual machine that supports Elasticsearch, used by TeamForge Code Search.

**Values**

Java specifications

**Default**

-Xms2g -Xmx2g



**Note:** By default, Elasticsearch JVM heap size is set to 2GB in TeamForge. You can increase this, if required.

**ETL\_JAVA\_OPTS**

The *ETL\_JAVA\_OPTS* token specifies the memory settings for the Java virtual machine that supports the ETL (Extract Transform and Load) job.

**Values**

Java specifications

**Default**

-Xms160m -Xmx512m -server -XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=/tmp -verbose:gc -XX:+PrintGCtimeStamps -XX:+PrintGCDetails -Dsun.rmi.dgc.client.gcInterval=600000 -Dsun.rmi.dgc.server.gcInterval=600000 -Djava.security.egd=file:/dev/urandom

**INTEGRATION\_JAVA\_OPTS**

The *INTEGRATION\_JAVA\_OPTS* token specifies the memory settings for the Java virtual machine that supports the site's integrated source control services.

**Values**

Java specifications

**Default**

-Xms160m -Xmx160m -server -XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=/tmp -verbose:gc -XX:+PrintGCtimeStamps -XX:+PrintGCDetails -Dsun.rmi.dgc.client.gcInterval=600000 -Dsun.rmi.dgc.server.gcInterval=600000 -Djava.security.egd=file:/dev/urandom

## JBOSS\_JAVA\_OPTS


The *JBOSS\_JAVA\_OPTS* variable specifies the memory settings for the JBoss Java virtual machine.

### Values

Java specifications

### Default

-Xms1024m -Xmx2048m

 **Important:** All JVM parameters but -Xms1024m and -Xmx2048m have been hard-coded in the TeamForge core application. You cannot manually configure any of the following default JVM parameters in the `site-options.conf` file.

- -XX:+UseParallelGC
- -XX:MaxMetaspaceSize=512m
- -XX:ReservedCodeCacheSize=128M
- -server
- -XX:+HeapDumpOnOutOfMemoryError
- -XX:HeapDumpPath=/tmp -verbose:gc
- -XX:+PrintCodeCache
- -Djsse.enableSNIExtension=false
- -Dsun.rmi.dgc.client.gcInterval=600000
- -Dsun.rmi.dgc.server.gcInterval=600000
- -Djava.security.egd=file:/dev/urandom
- -Djava.awt.headless=true.

When you change the default value of a JVM parameter such as "-XX:HeapDumpPath", the JBoss runtime parameters include both the user defined and default values for the JVM parameter. However, JBoss runs with the default value and ignores any user defined value.

## PHOENIX\_JAVA\_OPTS

The *PHOENIX\_JAVA\_OPTS* variable specifies the memory settings for the Java virtual machine that supports the site's ability to send and receive email and to index data for search.

### Values

Java specifications

### Default

-Xms256m -Xmx256m -server -XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=/tmp -verbose:gc -XX:+PrintGCTimeStamps -XX:+PrintGCDetails -Dsun.rmi.dgc.client.gcInterval=600000 -Dsun.rmi.dgc.server.gcInterval=600000 -Dsf.luceneOptimizeEvery=100000 -Djava.security.egd=file:/dev/urandom

## ETL\_JOB\_THREAD\_COUNT

The *ETL\_JOB\_THREAD\_COUNT* variable specifies the number of Extract, Transform and Load (ETL) jobs that can be run simultaneously.

### Values

1-100

### Default

2

**Comments**

If you only have a few jobs to be triggered few times a day, then 1 thread is sufficient. If you have tens of thousands of jobs, that needs to be triggered every minute, then you probably require thread counts like 50 or 100 (this depends on the nature of the work that your jobs perform, and your systems resources).

**ETL\_JOB\_TRIGGER\_TIME**

The *ETL\_JOB\_TRIGGER\_TIME* variable specifies the time and date for recurrent Extract, Transform and Load (ETL) jobs.

**Values**

Cron expression.

**Default**

0 30 2 \* \* ?

**Comments**

This variable takes a cron expression for a value, and not an absolute time value. The default value evaluates to 2.30 a.m. local time. For help with cron expressions, see [Cron Trigger Tutorial](#).

**ETL\_SOAP\_SHARED\_SECRET**

The *ETL\_SOAP\_SHARED\_SECRET* variable enables users to access site-wide reporting data via a SOAP client.

**Values**

String (possibly encrypted).

**Default**

mightyetlsoapsecret

**FILTER\_DROPDOWN\_MAX\_SELECTION**

By default, the drop-down lists with multi-select feature let you select up to 10 filter values. However, you can set any value that suits your requirement for this *FILTER\_DROPDOWN\_MAX\_SELECTION* token to increase or decrease the count.

**Values**

Any positive integer.

**Default**

10

This token was added in TeamForge 7.1

**FORBIDDEN\_PASSWORD**

The *FORBIDDEN\_PASSWORD* variable restricts specified words from being used as passwords.

**Values**

Comma-separated strings

**Default**

None

**GERRIT\_DATABASE\_HOST**

This is the Gerrit Postgres database host. The Gerrit configuration property, *database.hostname*, is derived from the value of *GERRIT\_DATABASE\_HOST* token in the *runtime-options.conf*. It can be overridden in *site-options.conf* and requires runtime creation with execution of the post installation script.

**Default**

127.0.0.1

**GERRIT\_DATABASE\_NAME**

This refers the Gerrit database schema. The Gerrit configuration property, *database.database*, is derived from the value of *GERRIT\_DATABASE\_NAME* in *runtime-options.conf*. It can be overridden in *site-options.conf* and requires runtime creation with execution of post installation script.

**Default**

reviewdb

**GERRIT\_DATABASE\_USER**

This is the PostgresDB role name that has access to the Gerrit database '*GERRIT\_DATABASE\_Name*'. The Gerrit configuration property, *database.username*, is derived from the value of *GERRIT\_DATABASE\_USER* in the *runtime-options.conf*. It can be overridden in *site-options.conf* and requires runtime creation with execution of post installation script.

**Default**

gerrit

**GERRIT\_GIT\_PUSH\_THRESHOLD**

The *GERRIT\_GIT\_PUSH\_THRESHOLD* token determines the maximum number of commits in a single Git push. If the limit exceeds, only a single commit object is created in the TeamForge.

**Values**

Any positive integer.

**Default**

30

**GERRIT\_GIT\_REFRESH\_PERIOD**

The *GERRIT\_GIT\_REFRESH\_PERIOD* token sets the interval in seconds after which Git Integration synchronizes all the repositories and all RBAC permission with TeamForge.

**Values**

Seconds.

**Default**

3600 Seconds.

**GERRIT\_REPLICATION\_MODE**

Use this site-options token to set the Git integration server as either 'master' or 'slave' server. In case you do not want replication (standalone mode) or you have only one primary source for repositories, set this token to master. On the other hand, if you have a master Git integration server and you want to replicate (mirror) its repositories on a secondary slave Git integration server, set this token to slave on the slave Git integration server.

**Values**

master or slave

**Default**

master

**Comments**

By default, in TeamForge 8.1 (and later), this token is set to master in the `runtime-options.conf` during runtime creation. As you cannot change the replication mode of a Git server after initial runtime creation, you have to set this to slave at the very beginning of your installation process in case you want to configure the server as a mirror of a 'master' Git server. It is not possible to have a Git master and slave configured on the same node, but you can have multiple masters and slaves in your TeamForge environment. Each slave belongs to exactly one master. Once a replica server is set up, it is not possible to reassign it to a different master Git integration server at a later point in time.

**GERRIT\_REPLICATION\_MASTER\_EXTERNAL\_SYSTEM\_ID**

If GERRIT\_REPLICATION\_MODE is set slave, this token specifies external system ID of the 'master' Git integration server.

**Values**

Alphanumeric string (exsy<number>) of a master Git integration server

**Comments**

This token is mandatory if GERRIT\_REPLICATION\_MODE is set to slave, without which the runtime creation shall fail. On the contrary, the runtime recreation will also fail if GERRIT\_REPLICATION\_MODE is set to master and this token is present in the `site-options.conf` file.

**GERRIT\_SMTP\_SERVER**

This is the hostname of the SMTP mail server for Gerrit. The Gerrit configuration property, `sendmail.smtpServer`, is derived from the value of `GERRIT_SMTP_SERVER` in the `runtime-options.conf`. It can be overridden in `site-options.conf` and requires runtime creation with execution of post installation script.

**Default**

localhost

**GERRIT\_SYNCH\_PORT**

The Port over which TeamForge communicates to Gerrit. The Gerrit configuration property, `teamforge.apiPort`, is derived from the value of `GERRIT_SYNCH_PORT` in the `runtime-options.conf`. It can be overridden in `site-options.conf` and requires runtime creation with execution of post installation script.

**Default**

9081

**HIGHCHARTS\_EXPORT\_REQUEST\_MAX\_WAIT**

This token is used to set the number of milliseconds that the Highcharts web application has to wait for response from the phantomjs server before it times out.

TeamForge chart export feature requires a pool of phantomjs servers to be running in the application server that is managed by a Highcharts web application. The phantomjs server pool runs as a blocking queue.

This token is used to specify the number of milliseconds that the Highcharts web application has to wait before it times out. In other words, the phantomjs server is expected to respond to the Highcharts web application within the time limit (in milliseconds) set in `HIGHCHARTS_EXPORT_REQUEST_MAX_WAIT`.

**Default value**

The default value is 500 milliseconds.

This token was added in TeamForge 7.2.

**HIGHCHARTS\_EXPORT\_REQUEST\_POOL\_SIZE**

TeamForge chart export feature requires a pool of phantomjs servers to be running in the application server that is managed by a Highcharts web application. This token is used to specify the number of phantomjs servers to be running in the pool.

**Default value**

The default pool size is 10.

This token was added in TeamForge 7.2.

**HTTPD\_LOG\_DIR**

The *HTTPD\_LOG\_DIR* variable specifies the path where information about the activity of the TeamForge site's Apache service is written.

**Values**

Path specification

**Default**

```
{__LOG_DIR__}/httpd
```

**INCLUDE\_ORGANIZATION\_USER\_FIELD**

The *INCLUDE\_ORGANIZATION\_USER\_FIELD* variable controls whether the organization entry is displayed while creating a user account.

**Values**

true or false

**Default**

true

**INDEXING\_TIMEOUT**

The *INDEXING\_TIMEOUT* token allows you to configure the time limit for indexing a file.

**Values**

Integer (in min)

**Default**

5 min

**INITIAL\_PASSWORD\_CHANGE\_ACTIVATION\_CODE\_TIMEOUT**

An administrator can optionally supply a password when creating a user. If the password is not specified while creating the user, the user is sent an email with a ticket to set the password. This *INITIAL\_PASSWORD\_CHANGE\_ACTIVATION\_CODE\_TIMEOUT* token sets the duration (in hours) for which the password ticket is valid.

**Values**

Integer (hours)

**Default**

72

**INTEGRATION\_LOG\_DIR**

The *INTEGRATION\_LOG\_DIR* variable specifies the path where information about the activity of the TeamForge site's source code integrations is written.

**Values**

Path specification

**Default**

```
{__LOG_DIR__}/integration
```

**JAMES\_ACCEPTED\_RELAYS**

Add the TeamForge domain name, Git domain name and Git Server's host name to the *JAMES\_ACCEPTED\_RELAYS* token in case you have a distributed setup where Git runs on a separate server and you want Git notification emails to be sent to the TeamForge Application Server.

**Values**

Set this token as illustrated:

```
JAMES_ACCEPTED_RELAYS=127.0.0.1,<APP_DOMAIN_NAME>,<GIT_DOMAIN_NAME>,<GIT_SERVER_HOSTNAME>
```

**Default**

```
JAMES_ACCEPTED_RELAYS=127.0.0.1,{__CEE_DOMAIN__}
```

**JAMES\_GATEWAY\_HOST**

The *JAMES\_GATEWAY\_HOST* variable specifies a mail server with Internet access, separate from the TeamForge server.

**Values**

Email address specification

**Default**


None

**Comments**

Specifying a gateway host assures delivery of site email to users if your TeamForge server cannot connect to a DNS server or cannot get outside connections over port 25.

The mail account specified must be hosted on a separate server from the TeamForge site server.

The *SYSTEM\_EMAIL*, *ADMIN\_EMAIL*, and *JAMES\_POSTMASTER\_EMAIL* variables can specify the same address.

 **Note:** Specify the gateway host by its fully qualified domain name, not a host name.



**JAMES\_GATEWAY\_\***

You can set up TeamForge to relay emails through an SMTP gateway (such as Amazon AES) that uses authentication. By default, James sends emails directly. However, you may prefer relaying emails through an enterprise relay server. Configuring the `JAMES_GATEWAY_*` tokens let you do that.

For more information, see [Relay Emails through SMTP gateway with authentication](#).

**JAMES\_LOG\_DIR**

The `JAMES_LOG_DIR` variable specifies the path where information about the activity of the TeamForge site's email component is written.

**Values**

Path specification

**Default**

```
{__LOG_DIR__}/james
```

**JAMES\_POSTMASTER\_EMAIL**

The `JAMES_POSTMASTER_EMAIL` variable specifies a valid email address for the person or machine that handles email for the domain, such as `postmaster@supervillain.org`.

**Values**

Email address specification

**Default**

```
root@{__APPLICATION_HOST__}
```

**Comments**

The mail account specified must be hosted on a separate server from the TeamForge site server.

The `SYSTEM_EMAIL`, `ADMIN_EMAIL`, and `JAMES_POSTMASTER_EMAIL` variables can specify the same address.

**JBOSS\_ALARM\_TIMEOUT**

The `JBOSS_ALARM_TIMEOUT` variable specifies the time duration within which the JBoss service is expected to respond to requests sent by `jboss_watchdog`.

**Values**

Integer

**Default**

20

**LINUX\_USERNAME\_MODE\_ENABLED**

The `LINUX_USERNAME_MODE_ENABLED` token, if set to true, overrides the default TeamForge user naming convention that bars user names with anything but an alphabet as the first character. It is recommended to use this token on sites with CVS integration.

**Values**

true or false


## Default

false

## Comment

The token is commented out (disabled) by default and is available in the `site-options.conf` file.

 **Important:** Do not set the `LINUX_USERNAME_MODE_ENABLED` to true on sites without CVS integration. Instead, you must set the `RELAXED_USERNAME_MODE_ENABLED` token to true in case your site has no CVS integration.

 **Caution:** While TeamForge can allow user names with anything but an alphabet as the first character, the same may not be true with all or some of the integrated applications you may have on your site. As a word of caution, consider the job at hand and understand the consequences before overriding TeamForge's default user naming convention.

## LISTEN\_BACKLOG

The `LISTEN_BACKLOG` token is used to specify the maximum length of the queue for the pending connections in the Apache server.

## Values

Integer

## Default

The default value is obtained from the system Kernel configuration.

```
/sbin/sysctl -n -e net.ipv4.tcp_max_syn_backlog
```

## LISTEN\_IP

In a distributed setup, you can use this `<host>:<service>:LISTEN_IP` token to control which IPs the services bind to so that you can make sure that services are not overexposed than necessary.

By default, services bind to the IP address corresponding to the `<host>:PUBLIC_FQDN` token. However, you can override this using the `<host>:<service>:LISTEN_IP` token.

A few use cases:

- In a distributed setup, you may want to bind a particular IP address of the TeamForge database server (PostgreSQL server) to the `ctfcore-database` service:

```
server-01:ctfcore-database:LISTEN_IP = 1.2.3.4
```

- To bind the mail service to a particular IP:


```
localhost:mail:LISTEN_IP = 1.2.3.4
```

- To make TeamForge listen to a specific IP of a particular server, say the SCM server.

```
myscmserver:LISTEN_IP = 1.2.3.4
```

- To bind all your services to one IP address (typically in a single server setup):

```
localhost:LISTEN_IP = 1.2.3.4
```

 **Note:** You cannot use more than one IP address with the `LISTEN_IP` token.

## LOG\_DIR

The `LOG_DIR` variable specifies the path where information about the TeamForge site's activity is written.

**Values**

Path specification

**Default**

{\_\_SITE\_DIR\_\_}/log

**LOG\_QUERY\_TIME\_THRESHOLD**

The *LOG\_QUERY\_TIME\_THRESHOLD* variable enables you to log database requests at INFO level if they run longer than a given period.

By default, database requests are logged at DEBUG level. Configuring a value for *LOG\_QUERY\_TIME\_THRESHOLD* causes requests that run for a period greater than that value to be logged at the INFO level, which makes them show up in `vamessages.log`.

Set the value to zero to log all database queries at INFO.

**Values**

Integer (in milliseconds)

**Default**


1000

**LOGIN\_ATTEMPT\_LOCK**

Use the *LOGIN\_ATTEMPT\_LOCK* token to set the permissible number of unsuccessful login attempts after which the user account is automatically locked.

**Values**

1-3

 **Important:** You can now selectively disable this feature even if you have set the *REQUIRE\_PASSWORD\_SECURITY* site options token to true. Set any negative value (such as '-1' or '-10000') in case you want to disable this feature altogether.

Note that with TeamForge 8.1 and earlier versions, setting *LOGIN\_ATTEMPT\_LOCK=-1* means the user account would be locked at the very first unsuccessful login attempt. This behavior has been removed from TeamForge 8.2 (and later).

**Default**

3

**LOGIN\_CONFIG\_XML\_FILE**

The *LOGIN\_CONFIG\_XML\_FILE* variable specifies the path to the LDAP configuration file.

**Values**

Path specification

**Default**

{\_\_DATA\_DIR\_\_}/etc/login-config.xml

**LOGROTATE\_ARCHIVE\_COUNT**


Use the *LOGROTATE\_ARCHIVE\_COUNT* token to set the number of most recent logs to be preserved at any give point in time.

**Values**

Any positive integer.

**Default**

The default value is "7". Meaning, logs for the last 7 days are preserved at any given point in time. Logs older than 7 days are removed from the log archive folder.

 **Note:** This token replaces the *LOGROTATE\_MAXAGE* token.

**MAX\_WWW\_CLIENT**

The *MAX\_WWW\_CLIENT* variable specifies the maximum number of Tomcat request processing threads to be created by the HTTP connector.

**Values**

Integer

**Default**

220

**MEMCACHED\_SERVER\_HOST**

This is the host name or IP address of the server that runs Memcached. Enter the host name or IP address of the Memcached server if and only if you have set the *ENABLE\_CACHING\_WITH\_MEMCACHED* token to true.

**Values**

Host name or IP address of the server that runs Memcached.

**Default**

localhost

See [Install Memcached](#) on page 694 for more information.

**MEMCACHED\_SERVER\_PORT**

This is the port on which Memcached is configured. Enter the port that TeamForge connects to in order to communicate with Memcached if and only if you have set the *ENABLE\_CACHING\_WITH\_MEMCACHED* token to true.

**Values**

Memcached port number.

**Default**

11211

See [Install Memcached](#) on page 694 for more information.

**MEMCACHED\_SERVER\_TTL**

The *MEMCACHED\_SERVER\_TTL* value, in seconds, is the duration for which the cached information is valid (or retained) in Memcached before it is invalidated or removed from Memcached. Enter the TTL (time to live) duration, in seconds, for the *MEMCACHED\_SERVER\_TTL* token if and only if you have set the *ENABLE\_CACHING\_WITH\_MEMCACHED* token to true.

**Values**

Time duration, in seconds.

**Default**

600 (seconds)

**Comments**

Updates to TeamForge authorization rules such as PBPs, roles and role assignments, user restrictions, enabling or disabling users and so on can take up to N seconds to take effect, where `MEMCACHED_SERVER_TTL=N` (seconds). In other words, the `mod_authnz_ctf` module uses the cached information available with Memcached irrespective of the most recent changes, if any, on the TeamForge side. This implies that having a long TTL duration, while improves the response time, can also result in cached information being out of sync with TeamForge, especially on large sites that are prone to frequent RBAC changes.

See [Install Memcached](#) on page 694 for more information.

**MIGRATION\_LOG\_DIR**

The `MIGRATION_LOG_DIR` variable specifies the path where information about the conversion of site data is written during an upgrade.

**Values**

Path specification

**Default**

```
{__LOG_DIR__}/runtime
```

**MAX\_PASSWORD\_LENGTH**

The `MAX_PASSWORD_LENGTH` variable sets the longest password that the system allows when a user account is created.

**Values**

Integer (number of characters)

**Default**

256

**MINIMUM\_PASSWORD\_LENGTH**

The `MINIMUM_PASSWORD_LENGTH` variable sets the shortest password that the system allows when a user account is created.

**Values**

Integer (number of characters)

**Default**

6

**MINIMUM\_USERNAME\_LENGTH**

The `MINIMUM_USERNAME_LENGTH` variable sets the shortest username that the system allows when a user account is created.

**Values**

Integer (number of characters)

**Default**

3

**MIRROR\_DATABASE\_HOST**

The *MIRROR\_DATABASE\_HOST* variable is a TeamForge database token that specifies the host of the database. This variable allows to extract the reporting data from the mirror TeamForge database through the Extract, Transform and Load (ETL) process.

**Values**

Alphanumeric string

**Default**

The *MIRROR\_* token takes the value of *DATABASE\_* token.

**Comment**

Example: Enter *MIRROR\_DATABASE\_HOST*=cu349.cloud.sp.collab.net (server name)

Add this token to the `site-options.conf` only if you setup a mirror database.

**MIRROR\_DATABASE\_NAME**

The *MIRROR\_DATABASE\_NAME* variable is a TeamForge database token that specifies the name of the TeamForge database. This variable allows to extract the reporting data from the mirror TeamForge database through the Extract, Transform and Load (ETL) process.

**Values**

Alphanumeric string

**Default**

The *MIRROR\_* token takes the value of *DATABASE\_* token.

**Comment**

Example: Enter *MIRROR\_DATABASE\_NAME*=ctfdb

Add this token to the `site-options.conf` only if you setup a mirror database.

**MIRROR\_DATABASE\_PASSWORD**

The *MIRROR\_DATABASE\_PASSWORD* variable is a TeamForge database token that specifies the password of the database. This variable allows to extract the reporting data from the mirror TeamForge database through the Extract, Transform and Load (ETL) process.

**Values**

Alphanumeric string

**Default**

The *MIRROR\_* token takes the value of *DATABASE\_* token.

**Comment**

Example: Enter *MIRROR\_DATABASE\_PASSWORD=ctfpwd*

Add this token to the *site-options.conf* only if you setup a mirror database.

**MIRROR\_DATABASE\_PORT**

The *MIRROR\_DATABASE\_PORT* variable is a TeamForge database token that specifies the port number of the database. This variable allows to extract the reporting data from the mirror TeamForge database through the Extract, Transform and Load (ETL) process.

**Values**

Port specification

**Default**

The *MIRROR\_* token takes the value of the *DATABASE\_* token.

**Comment**

Example: Enter *MIRROR\_DATABASE\_PORT=5432*.

Add this token to the *site-options.conf* only if you setup a mirror database.

**MIRROR\_DATABASE\_USERNAME**

The *MIRROR\_DATABASE\_USERNAME* variable is a TeamForge database token that specifies the database user's name. This variable allows to extract the reporting data from the mirror TeamForge database through the Extract, Transform and Load (ETL) process.

**Values**

Alphanumeric string

**Default**

The *MIRROR\_* token takes the value of the *DATABASE\_* token.

**Comment**

Example: Enter *MIRROR\_DATABASE\_USERNAME=ctfuser*

Add this token to the *site-options.conf* only if you setup a mirror database.

**NOTIFY\_SITE\_ADMINS\_FOR\_SITE\_ACTIVITIES**

The *NOTIFY\_SITE\_ADMINS\_FOR\_SITE\_ACTIVITIES* variable ensures that the activities at the site level are intimated to the site administrators through email notifications.

The site administrator can receive notifications on the following operations:

- User creation
- Project creation
- Blacklisted users
- SCM operations

**Values**

True or False

**Default**

True

This token was added in TeamForge 8.0.

**OBFUSCATION\_ENABLED**

The *OBFUSCATION\_ENABLED* variable is used to run the TeamForge application in the obfuscation mode for security purpose. Password obfuscation is enabled by default. As a result, all password-related tokens are encrypted in all the TeamForge configuration files.

**Values**

true or false

**Default**

true

**Comment**

When the TeamForge application is running in the obfuscation mode, the database login credentials, shared secrets etc., are encrypted and stored in the TeamForge configuration files for security reasons.

**OBFUSCATION\_KEY**

The *OBFUSCATION\_KEY* variable is used by the TeamForge obfuscation component as an input to the obfuscation algorithm for encryption and decryption purposes.

**Values**

AlphaNumeric (length greater than or equal to 8 bytes)

**Default**

XSJt43wN

**ONLY\_SITE\_ADMIN\_CAN\_EDIT\_SINGLE\_SIGN\_ON**

This site-options token, if set to true, ensures that only site administrators can turn on single sign on (SSO) for linked applications (including Build & Test). Set this token to false to have both site and project administrators turn SSO on and off.

**Values**

Either true or false.

**Default**

true

This token was added in TeamForge 7.2.

**ORGANIZATION\_EDITABLE**

The *ORGANIZATION\_EDITABLE* variable allows or prevents editing the organization value of a user account.

**Values**

true or false



**Default**

true

**PASSWORD\_CONTROL\_EFFECTIVE\_DATE**


The *PASSWORD\_CONTROL\_EFFECTIVE\_DATE* token is used to set the date from which the password security feature takes effect.

**Values**

Date (mm/dd/yyyy)

**Comments**

The [REQUIRE\\_PASSWORD\\_SECURITY](#) on page 768 is the master token that enables the password security feature.

 **Important:** Setting the *PASSWORD\_CONTROL\_EFFECTIVE\_DATE* token with a date is mandatory if [REQUIRE\\_PASSWORD\\_SECURITY](#) on page 768 is set to true.

**Example 1**

Consider a site with 130 users on which the password control kit (PCK) was not active. Of the 130 users, assume that:

- 100 users did not change password in last 100 days.
- 20 users did not change password in last 85 days.
- 10 users did not change password in last 75 days.

Assume that the following tokens are set on 01/01/2014 (current date):

- REQUIRE\_PASSWORD\_SECURITY=true
- PASSWORD\_WARNING\_PERIOD=20
- PASSWORD\_EXPIRY\_PERIOD=90
- PASSWORD\_DISABLE\_PERIOD=30
- PASSWORD\_DELETE\_PERIOD=60

PCK runs on 01/01/2014 and if you have *PASSWORD\_CONTROL\_EFFECTIVE\_DATE*=01/10/2014 (set to a future date):

- 100 users with no password change for the past 100 days would get a warning message that their passwords will expire in 10 days.
- 20 users with no password change for the past 85 days would get a warning message that their passwords will expire in 10 days.
- 10 users with no password change for the past 75 days would get a warning message that their passwords will expire in 15 days.

**Example 2**

Consider the following scenario in which:

- Current date = 01/01/2014
- *PASSWORD\_CONTROL\_EFFECTIVE\_DATE*=01/01/2013

In this scenario, the password control effective date is set to a date in the past. As a result, password control takes immediate effect and the PCK starts disabling, deleting or expiring user accounts right away.

**PASSWORD\_DELETE\_PERIOD**

The *PASSWORD\_DELETE\_PERIOD* variable specifies the time frame within which a disabled user account is automatically deleted.

**Values**

Integer (number of days)

**Default**

60

 **Note:** The *PASSWORD\_DELETE\_PERIOD* can be disabled by setting the value to zero.

**PASSWORD\_DISABLE\_PERIOD**

The *PASSWORD\_DISABLE\_PERIOD* variable specifies the time frame within which a user (soft-expired) is turned into a disabled user.

**Values**

Integer (number of days)

**Default**

30

**Comments**

A value of zero will disable this feature.

**PASSWORD\_EXPIRY\_PERIOD**

The *PASSWORD\_EXPIRY\_PERIOD* variable specifies the number of days after which the users' password expires.

**Values**

Integer (number of days)

**Default**

90

 **Note:** You cannot disable the *PASSWORD\_EXPIRY\_PERIOD* by setting the value to zero.

**PASSWORD\_REQUIRES\_MIXED\_CASE**

The *PASSWORD\_REQUIRES\_MIXED\_CASE* variable specifies that the user password must contain mixed case letters.

**Values**

true or false

**Default**

true

**PASSWORD\_REQUIRES\_NON\_ALPHANUM**

The *PASSWORD\_REQUIRES\_NON\_ALPHANUM* variable specifies that the user password must contain a non-alphanumeric character.

**Values**

true or false

**Default**

true

**PASSWORD\_REQUIRES\_NUMBER**

The *PASSWORD\_REQUIRES\_NUMBER* variable specifies that the user password must at least contain one number.

**Values**

true or false

**Default**

true

**PASSWORD\_WARNING\_PERIOD**

Set this token to alert users via emails about impending password expiration on a daily basis. Email alert starts "N" days before password expiration due date, where *PASSWORD\_WARNING\_PERIOD=N*, and ends only when the password is changed by the user.

**Values**

Positive integer (number of days).

**Default**

14

**PGSQL\_COMMIT\_DELAY**

The *PGSQL\_COMMIT\_DELAY* variable specifies the time delay between writing a commit record to the write ahead log (WAL) buffer and flushing the buffer out to disk.

**Values**

Integer (in microseconds)

**Default**

250

**Comments**

Together with the *PGSQL\_COMMIT\_SIBLINGS* option, this option allows a group of otherwise unrelated transactions to be flushed to disk at the same time, with possible significant performance gain.

**PGSQL\_COMMIT\_SIBLINGS**

The *PGSQL\_COMMIT\_SIBLINGS* variable sets the minimum number of concurrent open transactions to require before performing the delay specified by the *PGSQL\_COMMIT\_DELAY* option.

**Values**

Integer

**Default**

10

**Comments**

Together with the *PGSQL\_COMMIT\_DELAY* option, this option allows a group of otherwise unrelated transactions to be flushed to disk at the same time, with possible significant performance gain.

**PGSQL\_EFFECTIVE\_CACHE\_SIZE**

The *PGSQL\_EFFECTIVE\_CACHE\_SIZE* variable specifies the size of the OS data cache that is available to PostgreSQL. PostgreSQL can use that data to select the optimal way to execute requests.

**Comments**

The right value for this variable depends in part on the available RAM on the server where your site is running. Set this value at the highest amount of RAM that you expect to be always available to PostgreSQL.

See [What are the right PostgreSQL settings for my site?](#) on page 943 for values recommended by CollabNet.

**PGSQL\_LOG\_DIR**

The *PGSQL\_LOG\_DIR* variable specifies the path where information about the activity of the TeamForge site's PostgreSQL database is written.

**Values**

Path specification

**Default**

```
{__LOG_DIR__}/pgsql
```

**PGSQL\_MAINTENANCE\_WORK\_MEM**

The *PGSQL\_MAINTENANCE\_WORK\_MEM* variable specifies the maximum amount of memory to be used in maintenance operations such as `VACUUM`.

**Comments**

See [What are the right PostgreSQL settings for my site?](#) on page 943 for values recommended by CollabNet.

**PGSQL\_MAX\_CONNECTIONS**

The *PGSQL\_MAX\_CONNECTIONS* variable determines the number of concurrent connections available to the database server.

**Values**

Integer

**Default**

135

**PGSQL\_MAX\_FSM\_PAGES**

The *PGSQL\_MAX\_FSM\_PAGES* variable tells the `vacuum` process how many pages to look for in the shared free-space map.

**Values**

Integer

**Default**

500000

**Comments**

Each FSM page uses 6 bytes of RAM for administrative overhead, so increasing FSM substantially on systems low on RAM may be counter-productive.

**PGSQL\_MAX\_FSM\_RELATIONS**

The *PGSQL\_MAX\_FSM\_RELATIONS* variable specifies how many relations (tables) will be tracked in the free space map.

**Values****Default**

500

**PGSQL\_MAX\_STACK\_DEPTH**

The *PGSQL\_MAX\_STACK\_DEPTH* variable specifies the maximum safe depth of the server's execution stack.

**Values**

Integer

**Default**

5120

**PGSQL\_SHARED\_BUFFERS**

The *PGSQL\_SHARED\_BUFFERS* variable defines a block of memory that PostgreSQL will use to hold requests that are awaiting attention from the kernel buffer and CPU.

**Comments**

The right value for this variable depends in part on the available RAM on the server where your site is running.

See [What are the right PostgreSQL settings for my site?](#) on page 943 for values recommended by CollabNet.

**PGSQL\_STATEMENT\_TIMEOUT**

The *PGSQL\_STATEMENT\_TIMEOUT* variable is set to prevent the Postgres queries from running for a long period of time.

**Values**

Integer (Milliseconds)

**Default**

600000 (Milliseconds)

**Comments**

An error message is displayed for every timeout in the postgres.log file and the log message with the exid id is logged in the vessages.log and server.log files.

**PGSQL\_VACUUM\_COST\_DELAY**

The *PGSQL\_VACUUM\_COST\_DELAY* variable controls the length of time that an I/O process will sleep when the limit set by *vacuum\_cost\_limit* has been exceeded.

**Values**

Integer (milliseconds)

**Default**

50

**PGSQL\_WAL\_BUFFERS**

The *PGSQL\_WAL\_BUFFERS* variable specifies the number of buffers available for the Write Ahead Log.

**Comments**

If your database has many write transactions, setting this value bit higher than default may result better usage of disk space.

See [What are the right PostgreSQL settings for my site?](#) on page 943 for values recommended by CollabNet.

**PGSQL\_WORK\_MEM**

The *PGSQL\_WORK\_MEM* variable specifies the amount of memory to be used by internal sort operations and hash tables before switching to temporary disk files. .

**Comments**

The right value for this variable depends in part on the available RAM on the server where your site is running.

See [What are the right PostgreSQL settings for my site?](#) on page 943 for values recommended by CollabNet.

**PLANNING\_BOARD\_SWIM\_LANE\_LIMIT**

By default, not more than 250 cards are shown in a planning board swimlane. However, as a site administrator, you can increase or decrease the number of cards shown in the planning board swimlanes by configuring the site options token, *PLANNING\_BOARD\_SWIM\_LANE\_LIMIT*.

**Values**

A positive number.

**Default**

250

**Comments**

When you select a planning folder in one of the swimlanes and if X is greater than N, (where X = number of artifacts in the selected planning folder and N = *PLANNING\_BOARD\_SWIM\_LANE\_LIMIT*), the message, *Swimlanes in the Board View is currently configured to show N artifacts only, appears at the bottom of the swimlane.*

**PLANNING\_FOLDER\_DESC\_EDITOR**

The *PLANNING\_FOLDER\_DESC\_EDITOR* variable allows you to choose the type of text that can be used in the planning folder description using the editor tool.

**Values**

Plain Text

**Default**

Plain Text

**RELAXED\_USERNAME\_MODE\_ENABLED**

The *RELAXED\_USERNAME\_MODE\_ENABLED* token, if set to true, overrides the default TeamForge user naming convention that bars user names with anything but an alphabet as the first character. It is recommended to use this token on sites without CVS integration.

**Values**

true or false

**Default**

false

**Comment**

The token is commented out (disabled) by default and is available in the `site-options.conf` file.



**Important:** Do not set the `RELAXED_USERNAME_MODE_ENABLED` to true on sites with CVS integration. Instead, you must set the `LINUX_USERNAME_MODE_ENABLED` token to true in case your site has CVS integration.



**Caution:** While TeamForge can allow user names with anything but an alphabet as the first character, the same may not be true with all or some of the integrated applications you may have on your site. As a word of caution, consider the job at hand and understand the consequences before overriding TeamForge's default user naming convention.

**REPORTS\_DATABASE\_NAME**

The `REPORTS_DATABASE_NAME` variable specifies the name of the site's reporting database, also known as the datamart.

**Values**

Alphanumeric string

**Default**

teamforge\_datamart

**Comments**

It is OK for this variable to have the same value as `DATABASE_NAME`, because they are running in separate `pgsql` processes.

**REPORTS\_DATABASE\_PASSWORD**

The `REPORTS_DATABASE_PASSWORD` variable is the password for the Linux user that is authorized to read from and write to the site's reporting database.

**Values**

Alphanumeric string

**Default**

\$auto\$

**Comments**

It is OK for this variable to have the same value as `DATABASE_PASSWORD`, because they are running in separate PostgreSQL processes.

**REPORTS\_DATABASE\_PORT**

The `REPORTS_DATABASE_PORT` variable defines a separate port for the reporting database (aka datamart). Using a separate port can improve site performance when database utilization is high.

**Values**

Port specification

**Default**

5632

**Comments**

As of TeamForge 6.1, only port 5632 is supported.

**REPORTS\_DATABASE\_USERNAME**

The *REPORTS\_DATABASE\_USERNAME* variable specifies the Linux user that is authorized to read from and write to the site's reporting database.

**Values**

Alphanumeric string

**Default**

teamforge\_datamart

**Comments**

For some advanced operations, you may need to log into the database as the database user. However, under normal conditions only the TeamForge site process itself needs to access the database.

It is OK for this variable to have the same value as *DATABASE\_USERNAME*, because they are running in separate PostgreSQL processes.

**REPORTS\_ENABLE\_REPORT\_GENERATION**

The *REPORTS\_ENABLE\_REPORT\_GENERATION* variable is used to enable or disable the Reports tab in the UI.

**Values**

true or false

**Default**

true or false

**Comments**

Datamart is enabled by adding the 'datamart' service to the *HOST\_<hostname>* token. The service is disabled if datamart is not added. The default value of the *REPORTS\_ENABLE\_REPORT\_GENERATION* token is based on this service.

**REQUIRE\_PASSWORD\_SECURITY**

The *REQUIRE\_PASSWORD\_SECURITY* token, if set to true, enforces password security policy for the site.

**Values**

true or false

**Default**

true

**Comment**

This variable can be useful when an organization's security policy prohibits users from entering passwords without any restrictions. You can also set the *PASSWORD\_CONTROL\_EFFECTIVE\_DATE* token with a date from which



the password policy would be enforced. For more information, see [PASSWORD\\_CONTROL\\_EFFECTIVE\\_DATE](#) on page 761.

### **REQUIRE\_RANDOM\_ADMIN\_PASSWORD**

The *REQUIRE\_RANDOM\_ADMIN\_PASSWORD* variable restricts users from setting a random admin password.

#### **Values**

true or false.

#### **Default**

True (SaaS), false (BTF)

#### **Comment**

This variable when set to true, checks for a valid mail id in the ADMIN\_EMAIL token.

### **REQUIRE\_USER\_PASSWORD\_CHANGE**

The *REQUIRE\_USER\_PASSWORD\_CHANGE* variable determines if the user password needs to be changed during the first login instance.

#### **Values**

true or false.

#### **Default**

true

#### **Comment**

Setting a value true makes the new system force users to change password during first login and false otherwise.

### **RUNTIME\_LOG\_DIR**

The *RUNTIME\_LOG\_DIR* variable specifies the path where information about the activity of the TeamForge site's runtime environment is written.

#### **Values**

Path specification

#### **Default**

```
{__LOG_DIR__}/runtime
```

### **SAFE\_DOWNLOAD\_MODE**

Use this token to enforce downloading and saving of documents, attachments and files locally using the "Save" dialog box instead of inline views.

#### **Values**

true, false, none, all, html

You can set this token to "none" "all" or "html" that will force download of nothing, everything, or just html documents respectively.

#### **Default**

true

**SCM\_DEFAULT\_SHARED\_SECRET**

The *SCM\_DEFAULT\_SHARED\_SECRET* variable allows SCM Integrations to securely communicate with the TeamForge app server.

**Values**

- Alpha-numeric
- Special characters like '~!@#\$\$%^&\*'
- 16-24 byte length

**Default**

The default value is automatically generated during runtime.

**SCM\_SOAP\_TIMEOUT**

The *SCM\_SOAP\_TIMEOUT* variable is used to specify the connection timeout of the SCM soap requests between the APP and SCM servers.

**Values**

Integer (Milliseconds)

**Default**

300000 (Milliseconds)

**SCM\_USER\_ENCRYPTED\_PASSWORD**

The *SCM\_USER\_ENCRYPTED\_PASSWORD* variable is used to store the encrypted scmviewer password.

**Values**

- Alpha-numeric
- Special characters like '~!@#\$\$%^&\*'

**Default**

The default value will be in the encrypted format. See [password\\_util.sh](#) on page 795 for more information.

**SEARCH\_LOG\_DIR**

The *SEARCH\_LOG\_DIR* variable specifies the path where information about the activity of the TeamForge site's Lucene search component is written.

**Values**

Path specification

**Default**

```
{__LOG_DIR__}/james
```

**SEARCH\_MAX\_FILE\_SIZE**

The *SEARCH\_MAX\_FILE\_SIZE* variable sets an upper limit to the size of files that are indexed for search.

**Values**

Integer (bytes)

**Default**

10M

**Comment**

A value of zero or less specifies that there is no limit, which is the same as the default behavior without the variable.

**SEARCH\_SUPPRESS\_ARCHIVE\_SUB\_DOCS**

The *SEARCH\_SUPPRESS\_ARCHIVE\_SUB\_DOCS* variable prevents archive files from being indexed for search.

Archive files include zip, gzip, tar, and similar file types. They also include document files that are stored in archive format, such as docx files from Microsoft Word 2007.

**Values**

true, false

**Default**

true

**SESSION\_COOKIES\_ONLY**

the *SESSION\_COOKIES\_ONLY* variable restricts the persistence of all cookies to the user's current session.

If *SESSION\_COOKIES\_ONLY*=true, then all cookies created during the user session expire automatically when the user closes their browser. If it is false, the cookie expires according to the system logic for that particular cookie.

**Values**

true or false

**Default**

false

**Comment**

This variable can be useful when an organization's security policy prohibits cookies that persist across user sessions.

**SESSION\_TIMEOUT**

Use this token to set the user session timeout duration for newly created Application Server/Integration Server sessions.

**Values**

A positive value in the range of 1-1440.

**Default**

The default value of the *SESSION\_TIMEOUT* token is 30 minutes (for security reasons). You may change this to any value in the range of 1-1440 (minutes). However, you must create runtime for the changes to take effect.

**SOAP\_ANONYMOUS\_SHARED\_SECRET**

The *SOAP\_ANONYMOUS\_SHARED\_SECRET* variable allows users to have an anonymous login to the TeamForge site through SOAP.

**Values**

String (possibly encrypted)

**Default**

None

**Comment**

The variable must be configured to a non-empty value if users need to have an anonymous login to the site through SOAP. A value must be provided if site-wide reporting is enabled.

**SOAP\_ARTIFACT\_LIST\_LIMIT**

The *SOAP\_ARTIFACT\_LIST\_LIMIT* variable is used to limit the number of artifacts returned via SOAP calls.

**Values**

Integer

**Default**

-1: this means that the artifact list retrieved via SOAP is unlimited

**Comments**

In TeamForge releases earlier than 6.1.1, SOAP calls returned everything that was asked for, and that is the default behavior in TeamForge 6.1.1 as well. However, sites with performance and stability issues (OutOfMemory errors) in returning a large number of artifacts can now limit the number using this token. Changing this value requires a recreate-runtime and thus a site restart.



**Important:** Increasing the number of artifacts beyond the optimal 20,000 - 25,000 range might cause a heap dump.

**SSL**

The *SSL* variable activates Secure Socket Layer encryption for the TeamForge site.

**Values**

on or off

**Default**

on

**SSL\_CERT\_FILE**

The *SSL\_CERT\_FILE* specifies the path to the file where the TeamForge site's Secure Socket Layer certificate is stored.

**Values**

Path specification

**Default**

None

**SSL\_CHAIN\_FILE**

The *SSL\_CHAIN\_FILE* variable specifies the path to the file where the TeamForge site's SSL certificate chain file is stored.

**Values**

Path specification

**Default**

None

**SSL\_CIPHER\_SUITE**The *SSL\_CIPHER\_SUITE* token disables some of the less secure methods.**Values**

SSLCipherSuite method

**Default**

ECDHE-RSA-AES128-GCM-SHA256:ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES256-GCM-SHA384:ECDHE-ECDSA-AES256-GCM-SHA384:DHE-RSA-AES128-GCM-SHA256:DHE-DSS-AES128-GCM-SHA256:kEDH+AESGCM:ECDHE-RSA-AES128-SHA256:ECDHE-ECDSA-AES128-SHA256:ECDHE-RSA-AES128-SHA:ECDHE-ECDSA-AES128-SHA:ECDHE-RSA-AES256-SHA384:ECDHE-ECDSA-AES256-SHA384:ECDHE-RSA-AES256-SHA:ECDHE-ECDSA-AES256-SHA:DHE-RSA-AES128-SHA256:DHE-RSA-AES128-SHA:DHE-DSS-AES128-SHA256:DHE-RSA-AES256-SHA256:DHE-DSS-AES256-SHA:DHE-RSA-AES256-SHA:AES128-GCM-SHA256:AES256-GCM-SHA384:AES128-SHA256:AES256-SHA256:AES128-SHA:AES256-SHA:AES:CAMELLIA:DES-CBC3-SHA:!aNULL:!eNULL:!EXPORT:!DES:!RC4:!MD5:!PSK:!aECDH:!EDH-DSS-DES-CBC3-SHA:!EDH-RSA-DES-CBC3-SHA:!KRB5-DES-CBC3-SHA

**SSL\_PROTOCOL**The *SSL\_PROTOCOL* token disables some of the less secure methods.**Values**

SSLProtocol method

**Default**

all -SSLv3 -SSLv2

**SSL\_KEY\_FILE**The *SSL\_KEY\_FILE* specifies the path to the file where the TeamForge site's RSA private key is stored when Secure Socket Layer encryption is in effect.**Values**

Path specification

**Default**

None

**SUBVERSION\_BRANDING\_URI**The *SUBVERSION\_BRANDING\_URI* variable specifies the path component of the data repository URL.**Values**

BDB or FSFS

**Default**

BDB

**SVN\_AUTHNZ\_TIMEOUT**

The *SVN\_AUTHNZ\_TIMEOUT* token allows you to set the timeout value (in seconds) for the `mod_authnz_ctf` module.

**Values**

Timeout value in number of seconds.

**Default**

60 seconds


**SYSTEM\_EMAIL**

The *SYSTEM\_EMAIL* variable specifies a valid email address for the system administrator responsible for this site.

System administrators can use this email address to set up outage alerts and other notifications.

The mail account specified must be hosted on a separate server from the TeamForge site server.

The *SYSTEM\_EMAIL*, *ADMIN\_EMAIL*, and *JAMES\_POSTMASTER\_EMAIL* variables can specify the same address.

 **Important:** In TeamForge 6.x, the sender name and address for system-generated emails is taken from the value of the *SYSTEM\_EMAIL* variable. Therefore, changing the admin user's full name or email address does not affect the sender details of system-generated emails. This is different from TeamForge 5.x, in which the sender name and address for system-generated emails is derived from the admin user's full name and email address.

**Values**

Email address specification

**Default**

`root@{__APPLICATION_HOST__}`

**USE\_BROWSER\_CACHE\_PASSWORD**

The *USE\_BROWSER\_CACHE\_PASSWORD* variable restricts the storage of password in the browser when you login to the site.

**Values**

true/false

**Default**

true

**USE\_EXTERNAL\_USER\_AUTHENTICATION**

The *USE\_EXTERNAL\_USER\_AUTHENTICATION* variable specifies whether users can be authenticated through a separate system, such as OpenLDAP.

**Values**

true or false

**Default**

false

**USE\_STATIC\_SENDER\_EMAIL**

The *USE\_STATIC\_SENDER\_EMAIL* variable is used to prevent Out of Office emails from being posted to artifacts and discussion forums.

- The *USE\_STATIC\_SENDER\_EMAIL* token, if set to true, assigns the "Return-Path" parameter in the TeamForge notification email header with a "noreply@<user's email domain>" email ID. This prevents Out of Office replies from being posted to artifacts and discussion forums.
- On the other hand, the *USE\_STATIC\_SENDER\_EMAIL* token, if set to false, assigns the "Return-Path" parameter in the TeamForge notification email header with the email ID of the user (for example, tom@forge.collab.net) whose action triggers the notification email. In this case, Out of Office replies are posted to artifacts and discussion forums. If the user is a Site Administrator, the "Return-Path" parameter in the TeamForge notification email header is assigned with the email ID "root@<TeamForge domain name>".

**Values**

true/false

**Default**

false

**USER\_ACCOUNT\_RESTRICTED**

The *USER\_ACCOUNT\_RESTRICTED* variable determines whether newly created users are "restricted" or "unrestricted" users by default.

- Restricted users can access only public projects and projects of which they are members.
- Unrestricted users can access all projects except private projects of which they are not members.

**Values**

true or false

**Default**

true

**USER\_MONITORING\_REMOVE\_ENABLED**

Set the *USER\_MONITORING\_REMOVE\_ENABLED* variable to true, if you want to enable the feature that lets you remove one or more users from monitoring selected TeamForge objects.

**Values**

true or false

**Default**

false

**USER\_NEED\_PERMISSION\_TO\_VIEW\_FULL\_USER\_DETAILS**

The *USER\_NEED\_PERMISSION\_TO\_VIEW\_FULL\_USER\_DETAILS* variable restricts users from viewing other users' organization information.

**Values**

true or false

**Default**

false

## USERS\_WITH\_NO\_EXPIRY\_PASSWORD

The `USERS_WITH_NO_EXPIRY_PASSWORD` token specifies the users for whom there is no expiry of password.

### Values

Specify the usernames (for the user accounts) for which there is no expiry of password.

### Default

`USERS_WITH_NO_EXPIRY_PASSWORD=admin,nobody,system,scmviewer,scmadmin`

### Comment

The variable is enabled by default.

### Using multi-line blocks for site options

The multi-line block configuration is generally used by old SFEE sites. To define a `site-options.conf` token with a multi-line block value, you need to follow a certain syntax.

- Declare the token name with the value "START\_MULTILINE\_BLOCK". Syntax:  
`<TOKEN_NAME>=START_MULTILINE_BLOCK`
- Specify the multi-line values beneath the token.
- Complete the multi-line block with "END\_MULTILINE\_BLOCK" after all the multi-line values are specified.  
 Syntax: `END_MULTILINE_BLOCK`

### Example

```
SOURCEFORGE_CONFIGURATION_PROPERTIES_APPEND=START_MULTILINE_BLOCK
email.suppress.project_member_added=true
email.suppress.scm_user_password_synchronized=true
END_MULTILINE_BLOCK
```

## pebble-app.xml

The `pebble-app.xml` file, also known as the Pebble application configuration file, contains the text that the Pebble application displays in the TeamForge user interface.

This is an example of a default (unedited) `pebble-app.xml` file. To create your own integrated application config file, copy this one into a new file and replace the values with the values appropriate for the application you are integrating.

```
<?xml version="1.0" encoding="utf-8"?>
<!DOCTYPE integrated-application
  PUBLIC "-//CollabNet, Inc.//DTD Integrated Application Descriptor 1.0//EN"
  "http://schema.open.collab.net/sfee50/dtd/sf-pluggable-application-
  descriptor_1_0.dtd">
<integrated-application>
  <name>Pebble Blog</name>
  <description>l10n.application.description</description>
  <permissions>
    <permission dapMappedTo="View">Blog Reader</permission>
    <permission>Blog Contributor</permission>
    <permission>Blog Publisher</permission>
    <permission>Blog Owner</permission>
  </permissions>
  <prefix>PB</prefix>
  <id-pattern></id-pattern>
  <require-per-project-prefix>true</require-per-project-prefix>
  <require-scm-integration>true</require-scm-integration>
  <is-search-supported>>false</is-search-supported>
```



```

<!-- Page components for Integrated apps is not implemented for Alpha -->
<page-component>
  <require-page-component>true</require-page-component>
  <page-component-details>
    <inputtype>text</inputtype>
    <resultformat>html</resultformat>
    <description>l10n.pce.description</description>
    <title>l10n.pce.title</title>
  </page-component-details>
</page-component>
<config-parameters>
  <!-- Pebble Configuration Parameters -->
  <param>
    <title>l10n.blogname.title</title>
    <name>blogName</name>
    <description>l10n.blogname.description</description>
    <defaultvalue>My Blog</defaultvalue>
    <displaytype valuetype="String" maxlength="25">TEXT</displaytype>
    <editable>>false</editable>
  </param>
  <param>
    <title>l10n.blogdescription.title</title>
    <name>blogDescription</name>
    <description>l10n.blogdescription.description</description>
    <defaultvalue>My Awesome Blog</defaultvalue>
    <displaytype valuetype="String" maxlength="40">TEXT</displaytype>
    <editable>>true</editable>
  </param>
  <param>
    <title>l10n.richtexteditor.title</title>
    <name>richTextEditorEnabled</name>
    <description>l10n.richtexteditor.description</description>
    <defaultvalue>checked</defaultvalue>
    <displaytype valuetype="String">CHECKBOX</displaytype>
    <editable>>true</editable>
  </param>
  <param>
    <title>l10n.noofrecentblogentries.title</title>
    <name>recentBlogEntries</name>
    <description>l10n.noofrecentblogentries.description</description>
    <defaultvalue>3</defaultvalue>
    <displaytype valuetype="String">SELECT</displaytype>
    <option name="3">l10n.three.value</option>
    <option name="5">l10n.five.value</option>
    <option name="7">l10n.seven.value</option>
    <option name="9">l10n.nine.value</option>
    <editable>>true</editable>
  </param>
</config-parameters>
<bundles>
  <bundle locale="en">
    <key name="l10n.application.description">Pebble Blog App</key>
    <key name="l10n.pce.description">Display Blog Title for Given
Date.</key>
    <key name="l10n.pce.title">Enter Blog Date (in yyyy-mm-dd)</key>
    <key name="l10n.blogname.title">Blog Name</key>
    <key name="l10n.blogname.description">Please provide a name for
the Blog. This appears on all blog pages</key>
    <key name="l10n.blogdescription.title">Blog Description</key>
    <key name="l10n.blogdescription.description">Please provide a
description for the Blog. This appears below blog name on all pages</key>
    <key name="l10n.richtexteditor.title">Rich Text Editor</key>
    <key name="l10n.richtexteditor.description">Enable Rich Text
Editor for comments and Blog entries?</key>
  </bundle>
</bundles>

```

```

        <key name="110n.noofrecentblogentries.title">Recent Blog Entries</
key>
        <key name="110n.noofrecentblogentries.description">How many recent
blog entries do you want to see in the home page?</key>
        <key name="110n.three.value">3</key>
        <key name="110n.five.value">5</key>
        <key name="110n.seven.value">7</key>
        <key name="110n.nine.value">9</key>
    </bundle>
</bundles>
</integrated-application>

```

## pebble-dep.xml

The `pebble-dep.xml` file, also known as the Pebble deployment configuration file, contains the data that Pebble needs to interact with the TeamForge site.

### A sample `pebble-dep.xml` file for the REST service type

```

<?xml version="1.0" encoding="utf-8"?>
<!DOCTYPE integrated-application
PUBLIC "-//CollabNet, Inc.//DTD Integrated Application Descriptor 1.0//EN"
"http://schema.open.collab.net/sfee50/dtd/sf-pluggable-deploy-
descriptor_1_0.dtd">
<integrated-application>
    <name>Pebble Blog</name>
    <baseurl>https://cu064.cloud.sp.collab.net:13001/pebble/index.jsp</
baseurl>
    <gourl>https://cu064.cloud.sp.collab.net:13001/pebble/gourl/%p/%o</gourl>
    <endpoint>https://cu064.cloud.sp.collab.net:13001/pebble/services/rest/
ctfapi</endpoint>
    <servicetype>REST</servicetype>
</integrated-application>

```

### A sample `pebble-dep.xml` file for the SOAP service type

```

<?xml version="1.0" encoding="utf-8"?>
<!DOCTYPE integrated-application
PUBLIC "-//CollabNet, Inc.//DTD Integrated Application Descriptor 1.0//EN"
"http://schema.open.collab.net/sfee50/dtd/sf-pluggable-deploy-
descriptor_1_0.dtd">
<integrated-application>
    <name>Pebble Blog</name>
    <baseurl>https://cu177.cloud.maa.collab.net:13001/pebble/index.jsp</
baseurl>
    <gourl>https://cu177.cloud.maa.collab.net:13001/pebble/gourl/%p/%o</gourl>
    <endpoint>https://cu177.cloud.maa.collab.net:13001/pebble/services/
PebbleIntegrationService</endpoint>
    <servicetype>SOAP</servicetype>
</integrated-application>

```

## How is an integrated application described?

An integrated application is described using two XML files - a deployment configuration file and an application configuration file - that provide information to TeamForge about the configuration options exposed by the application.

In TeamForge version 6.1.1 and later, you have the ability to configure some integrated application settings using the user interface. You can also export these settings in XML format and make changes. To edit configuration settings, you would upload the XML file containing the updates.

## Integrated application settings



**Note:** Some of the tags are internationalized so that the application will display languages based on the browser locale. See [Internationalize your integrated application](#) on page 690 for more information.

### <name>

This is the title of the integrated application. When the integrated application is added to a project, the button that appears on the project pages has this name. This name must be unique -- you cannot use it for any other integrated application on the same TeamForge server.

This tag is used in both deployment and application configuration files.

### <adminurl>

When an application has an administration screen for configuring its parameters, this field contains that URL. It is optional.

This tag is used in the deployment configuration file.

### <baseurl>

This is the URL to which a user will be directed on clicking the integrated application button in a project.

This tag is used in the deployment configuration file.

### <endpoint>

This is the SOAP endpoint for the integrated application. The endpoint contains the various methods exposed by the integrated application that are called during the lifecycle of TeamForge.

This tag is used in the deployment configuration file.

### <gourl>

This indicates which URL must be used when an object id for an integrated application is specified (either via `Jump_to_id` or on the URL as `/sf/go/<objectid>`). This URL can support a couple of dynamic parameters.

- `%o` -- The object id entered by the user will be dynamically replaced here.
- `%p` -- The project id for the object entered will be dynamically replaced here.

For example, if the Go URL is `http://go.tourl.com/tracking?id=%o` and the object ID entered is `XYZ123`, then the URL will be replaced and redirected to `http://go.tourl.com/tracking?id=XYZ123`.

This tag is used in the deployment configuration file.

### <category>

This is used to display the integrated application under a separate **Category** tab. It is an optional parameter that is manually inserted based on the requirement. If the category is not defined, the integrated application appears under the **Generic** tab.

This tag is used in the deployment configuration file.

### <isbaseapp>

This validates whether the integrated application is a base application configured for the specified category. It is an optional parameter that is manually inserted based on the requirement.

This tag is used in the deployment configuration file.

### <is-search-supported>

This validates whether the integrated application needs to be displayed as an object category in the drop-down list for search options in TeamForge such as **Jump to ID** and **Advanced Search**. It is an optional parameter that is manually inserted based on the requirement.

This tag is used in the application configuration file.

### **<config-parameters>**

There can be any number of configuration parameters for an integrated application and they are displayed when associating the application to a project. These parameters are filled in by the project administrator and are available in the integrated application SOAP interface as configuration parameters. The integrated application gets a chance to validate these parameters and indicate back to TeamForge whether this project association is successful by passing in a "TRUE". It can return a "FALSE" if it doesn't want this project association to succeed. Each configuration parameter is placed inside the "param" tag, which can contain multiple elements to describe the parameter.

<b>&lt;title&gt;</b>	The internationalized title that appears for a project administrator to fill in while associating the integrated application to a project.
<b>&lt;name&gt;</b>	The Java variable under which the value for this parameter will be available on the integrated application.
<b>&lt;description&gt;</b>	The internationalized description that appears when a project administrator fills in or enters a configuration parameter.
<b>&lt;default value&gt;</b>	The default value for the parameter that will appear in the user interface during the association of an integrated application to a project.
<b>&lt;display type&gt;</b>	This is the type of display control used for the configuration parameter. We support "TEXT" for text fields, "CHECKBOX" for checkbox type controls, "RADIO" for radio buttons, and "SELECT" for select dropdowns. This field can also take an attribute that says what the value type for the field should be -- whether it should be an "Integer", "String" and so on. So if the field is expecting numbers, then entering "foo" as a value will throw a validation failure.
<b>&lt;option&gt;</b>	If the display type is "RADIO" or "SELECT", then these fields contain the individual options available for the display controls. This will contain a "name" attribute that will be sent to the integrated application when that option is selected from the UI. The value of this option should be an internationalized field as it is the value visible to the user.
<b>&lt;editable&gt;</b>	This specifies whether the configuration parameter should be editable once the integrated application is associated to a project. These configuration parameters are available when you add or edit an integrated project. If a parameter should not be "edited" post association, setting this to "false" will make it non-editable.

This tag is used in the application configuration file.

### **<description>**

This is an internationalized string for the integrated application's description. It contains information for TeamForge project and site administrators to know what the application does.

This tag is used in the application configuration file.

**<id-pattern>**

When trying to link to an integrated application id, this regular expression gets used for mapping. By default (if no value is provided), it looks for alphanumerical characters; in case you need specific characters to be matched (for example, JIRA, which has hyphens in ids), this value is used.

This tag is used in the application configuration file.

**<page-component>**

These settings are used for Project Content Editors. The integrated application content can become part of the standard Page Component data that appears in project home pages. The settings indicate the type of information that will be available from the integration application.

**<input-type>**

This is the input type control for an integrated app Page Component. We only support 2 types now. Either "select" so that the inputs can be shown from a "SELECT" dropdown and the users will be able to pick a value from there. Else, it can be a "text" where a simple "text" field will be entered for taking the user input.

**<result-format>**

This is the format in which the output of Page Component is returned. This can be a "list" which indicates that it will be a Table like output. The integrated app will send the results in an XML format and the Integrated app framework converts this into a list of records. The other option is "html", where the output from the Integrated application is just displayed on the screen.

**<page-component-description>**

The description that will appear when you add an Integrated application Page Component (Link to the page where " Add component" is available)

**<page-component-title>**

The title that will appear when adding an Integrated application Page Component (Link to the page where "PCE Add component" is available)

This tag is used in the application configuration file.

**<permissions>**

This is a collection of permissions that are exposed by the integrated application. There could be any number of such permissions. These permissions will appear as a part of the project's roles (existing ones, as well as ones newly added) and can be assigned along with other tool permissions. You can map one of these permissions with a "dapMappedTo" attribute -- this indicates the permission to be used when a user logs in without authentication (for example, for public projects). Typically, this is the permission to read data so that it doesn't need a login name; it varies from one application to another.

This tag is used in the application configuration file.

**<prefix>**

If the "require-per-project-prefix" attribute is false, the value of this tag is used for identifying the integrated application in Go URLs, associations, and linkifications. If the "require-per-project-prefix" attribute is true, the value is used only for the "Host" project. Each project must fill in its value as part of adding the integrated application. Click here for [steps to add integrated applications to a project](#).

This tag is used in the application configuration file. The prefix can contain alpha-numeric characters and cannot be more than six characters in length. For more information about prefix, see [How does an integrated application interact with other TeamForge tools?](#)

**<require-per-project-prefix>**

An integrated application can indicate to TeamForge whether the object ids that it generates are uniquely identifiable across the entire application (if yes, the value for the attribute is "false") or whether they need to be project-specific (in this case, the value for the attribute is "true"). If an integrated application needs per-project prefix, you must enter the prefix value when the integrated application is added to a project.

This tag is used in the application configuration file.

**<require-scm-integration>**

This indicates whether SCM commits need to be validated. Some applications might have business rules which indicate that a commit can be made only if certain conditions are met. If the integrated application has any such rules, the value for the attribute should be "true". There are also a couple of methods to be implemented in the SOAP endpoint.

This tag is used in the application configuration file.

**<require-page-component>**

Some integrated applications choose not to expose details as Page Components. For those that don't, set this tag to "false" and for those that do, set it to "true". If the value is "true", you must provide the "page-component-details" tags as well.

This tag is used in the application configuration file.

**<servicetype>**

TeamForge 6.1.0 and earlier releases supported only SOAP as the mechanism to talk from TeamForge to the integrated application. TeamForge 6.1.1 and later support REST calls. The servicetype tag indicates whether the protocol used for communication is REST or SOAP.

This tag is used in the deployment configuration file.

For examples of how these tags are used in the integration of the Pebble blogging application, see [pebble-dep.xml](#) and [pebble-app.xml](#).

**install.conf - Pebble**

The `install.conf` file contains the data needed to manage the Pebble installer.

This is an example of a default (unedited) `install.conf` file. To create your own Pebble installer config file, copy this one into a new file and replace the values with the values appropriate for the application you are integrating.

```
#Location of the Pebble Blog
pebble.base.dir=/u1/pebble
#CTF Information
ctf.baseurl=http://cu073.cubit.maa.collab.net/
#Tomcat Information
tomcat.port=13000
domain=cu073.cubit.maa.collab.net
timezone=Asia/Calcutta
java_home=/usr/java/jdk1.6.0_26/
protocol=http
java_opts=-Xms512m -Xmx512m
# This where the Pebble Blogs are stored
data.dir=/u1/pebble-data
secretkey=mistywasacat
```

**install.conf - Review Board**

The `install.conf` file contains the data needed to manage the Review Board installer.

This is an example of a default (unedited) `install.conf` file. To create your own Review Board installer config file, copy this one into a new file and replace the values with the values appropriate for the application you are integrating.

```
#Location of the Review Board installation directory
rb_dir=/ul/reviewboard

#Location of the Review Board data directory
rb_data_dir=/opt/collabnet/reviewboard/data

#Review Board site information
domain= myapp.collab.net
rb_database_type=postgresql
rb_database_name=ctfrbdb
rb_database_user=ctfrbuser
rb_database_password=ctfrbpwd
rb_database_host= <database_hostname>
rb_database_port=5432

#CTF Information
ctf_base_url=https://myapp.collab.net
ctf_site_var_dir=/opt/collabnet/teamforge/var
```

## iptables

This is the `/etc/sysconfig/iptables` output that will enforce the recommended security configuration.

```
# Firewall configuration written by redhat-config-securitylevel
# Manual customization of this file is not recommended.
*filter
:INPUT ACCEPT [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [0:0]
:RH-Firewall-1-INPUT - [0:0]
-A INPUT -j RH-Firewall-1-INPUT
-A FORWARD -j RH-Firewall-1-INPUT
-A RH-Firewall-1-INPUT -i lo -j ACCEPT
-A RH-Firewall-1-INPUT -p icmp --icmp-type any -j ACCEPT
-A RH-Firewall-1-INPUT -p 50 -j ACCEPT
-A RH-Firewall-1-INPUT -p 51 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state ESTABLISHED,RELATED -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 443 -j
ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 25 -j
ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 80 -j
ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 22 -j
ACCEPT
-A RH-Firewall-1-INPUT -j REJECT --reject-with icmp-host-prohibited
COMMIT
```

## login-config.xml

This is the sample `application-policy` block that you can copy into your `login-config.xml` file to support LDAP authentication.

## Notes

Replace the default `application-policy` block of the `login-config.xml` file with this code, then make the modifications specified in [Set up LDAP integration for the TeamForge site](#) on page 701. Option values that must be modified are highlighted in bold.

- When the username is passed to the login module from TeamForge, it is translated into a DN for lookup on the LDAP server. The DN that is sent to the LDAP server is `<principalDNPrefix><username><principalDNSuffix>`.
- In this example `application-policy` block, the username is stored in the People organizational unit in the `dev.sf.net` domain. This is represented as `,ou=People,dc=dev,dc=sf,dc=net`
- This example contains a single `login-module` section. If you are authenticating against multiple LDAP servers, include one `login-module` section per LDAP server, with the required option values modified appropriately for each one. If the same username exists in more than one LDAP server, the instance on the first LDAP server will be used.

## Sample code

```
<application-policy name="SourceForge">
  <authentication>
    <login-module code="org.jboss.security.auth.spi.LdapLoginModule"
flag="sufficient" >
      <module-option name="allowEmptyPasswords">>false</module-option>
      <module-option name="principalDNPrefix">uid</module-option>
      <module-option
name="principalDNSuffix">,ou=People,dc=dev,dc=sf,dc=net</module-option>
      <module-option
name="java.naming.factory.initial">com.sun.jndi.ldap.LdapCtxFactory</module-
option>
      <module-option name="java.naming.provider.url">ldap://
util.dev.sf.net:389/</module-option>
      <module-option
name="java.naming.security.authentication">simple</module-option>
    </login-module>
  </authentication>
</application-policy>
```

## Sample code for Active Directory integration

Active Directory is not supported. However, these sample lines in the `login-config.xml` file may help you make it work for a simple AD setup, without complex directory structures requiring additional search parameters.

Set the values of `java.naming.provider.url`, `principalDNSuffix` and `rolesCtxDN` as appropriate to your site.

For more detailed instructions, see <http://www.jboss.org/community/wiki/LdapLoginModule>.

```
<login-module code="org.jboss.security.auth.spi.LdapLoginModule"
flag="required" >
  <module-option name="java.naming.provider.url">ldaps://
<server_name>:636/</module-option>
  <module-option name="allowEmptyPasswords">>false</module-option>
  <module-option name="principalDNSuffix">@foo.bar.com</module-
option>
  <module-option name="rolesCtxDN">dc=Foo,dc=Bar,dc=Com</module-
option>
  <module-option name="matchOnUserDN">>true</module-option>
  <module-option name="uidAttributeID">sAMAccountName</module-
option>
  <module-option name="roleAttributeID">memberOf</module-option>
```



```

    <module-option name="roleAttributeIsDN">true</module-option>
    <module-option name="roleNameAttributeID">name</module-option>
</login-module>

```

## Scripts installed with TeamForge

---

System administrators can use these utilities to control the behavior of the application.

### TeamForge script

Use the `teamforge` script to deploy and undeploy services, start and stop services, verify the status of services, verify the application environment, bootstrap or migrate data and do much more.

#### Overview

Use the `teamforge` script wherever applicable as it subsumes the functions of the following legacy TeamForge scripts:

- `bootstrap-data.sh`
- `bootstrap-data.py`
- `bootstrap-reporting-data.sh`
- `bootstrap-reporting-data.py`
- `create-runtime.py`
- `collabnet`
- `migrate.py`

Run `teamforge` script as follows:

```

/opt/collabnet/teamforge/bin/teamforge [command] [-s serviceName]
[parameter]

```

For example, the following command displays the status of all the services:

```

/opt/collabnet/teamforge/bin/teamforge status

```

PID	Service	
6431	apache	[available]
7591	binary	[available]
	binary-database-postgres	[available]
	cron	[available]
	ctfcore	[available]
	ctfcore-app	[available]
	ctfcore-database	[available]
	ctfcore-database-admin	[available]
	ctfcore-database-admin-postgres	[available]
	ctfcore-database-postgres	[available]
	ctfcore-datamart	[available]
	ctfcore-datamart-admin	[available]
	ctfcore-datamart-admin-postgres	[available]
	ctfcore-datamart-postgres-cohosted	[available]
	cvs	[available]
	daemon-base	[available]
7688	etl	[available]
13723	gerrit	[available]
	gerrit-database-postgres	[available]
7618	integration	[available]
	integration-base	[available]
8103	jboss	[available]
	legacy-runtime	[available]
	logrotate	[available]
	mail	[available]
10286	phoenix	[available]
6833	postgres	[available]
	runtime-base	[available]
	search	[available]
	subversion	[available]

## TeamForge components and services

### Components

TeamForge comprises of a set of components such as ctfcore, subversion, james, and so on. Some components are always required to be installed, while others are optional. By spreading these components over multiple RPMs, we make sure that users do not have to install everything all the time. Though this is valuable, RPMs alone prove insufficient to manage the components and their inter-dependencies.

- Some components do not have a physical representation but are configuration-only.
- RPM dependencies are restricted to the local machine only; however, in a distributed installation, dependencies between components must be tracked across servers.

In addition to the physical componentization (using RPMs), there is also a need for a logical componentization of TeamForge.

### Services

Services represent a logical component of TeamForge. Such a logical component may either be a feature, which the user explicitly opts to install (for example, CVS) or a technical component (for example, Apache and Logrotate), which other services depend on. Services

come with additional metadata which makes it possible to track and manage dependencies to a more fine-grained level.

- **Deployment dependencies** specify which other services need to be deployed locally.
- **Provided Endpoints** specify which network endpoints the service offers.
- **Required Endpoints** specify which network endpoints the service depends on.

As such services are more fine-grained than RPMs and it is common to have a single RPM containing multiple services.

## Service life cycle

A service can be in one of the following states:

- **Uninstalled:** A service is uninstalled if the RPM that contains it is not installed. Uninstalled services do not exist as far as TeamForge is concerned.
- **Undeployed:** The RPM containing the service is installed, but the service has not been deployed yet. Deployment is also referred to as "creating the runtime", but is specific to one service.
- **Dependencies unavailable:** The service itself might be available, but at least one of its deployment dependencies is not in "Available" state.
- **Available:** Service is fully functional.

Services that manage data have the following additional states:

- **Not bootstrapped:** Data structures have not been initialized yet.
- **Not migrated:** Data structures are initialized, but data needs to be migrated.

Services that have a daemon have the following additional states:

- **Dependencies unavailable:** The service itself might be available, but at least one of its deployment dependencies is not in "Available" state.
- **Ports Blocked:** The service is impeded from starting up because at least one of the ports it needs is in use by a different process.
- **Stopped:** Service is an auto start-type service, yet is stopped.
- **Inactive:** Service is a demand start-type service and is stopped.
- **Starting:** Service is in the process of starting up.
- **Available:** Service is running and healthy according to its health check.
- **Unhealthy:** Service is running but unhealthy according to its health check.
- **Dead:** Service is supposed to be running, but the process disappeared.

- **Doomed:** Service is technically running, but it will never work properly because some part of it failed to initialize properly.
- **Stopping:** Service is in the process of stopping.

## Parameters

`teamforge` script accepts the following command line parameters:

**[-s | --service]**

Use the `-s` parameter to selectively start, stop, and bootstrap a specific service or component. For example, the following command gets you the status of Jboss:

```
/opt/collabnet/teamforge/bin/
teamforge status -s jboss
```

**[-f | --site\_options\_file]**

Use the `-f` parameter to pass the `site-options.conf` file's path as a command line parameter (default is `/opt/collabnet/teamforge/etc/site-options.conf`).

**[--skip-verification]**

Pass this parameter if you want to skip environment verification.

**[-y | --yes]**

Pass this parameter if you want to skip confirmation prompts.

## Commands

`teamforge` script can perform the following actions:

**status**

Show status of all services. Use the `-s` parameter to know the status of a specific service.

**bootstrap**

Bootstrap data (re-create data structures). Use the `-s` parameter to selectively bootstrap a specific component. Suppose, you did not have SVN on your TeamForge site and if you add SVN while upgrading to TeamForge 16.10. You can now selectively bootstrap SVN alone.

**deploy**

Deploy service(s).

**migrate**

Migrate data to latest schema.

**provision**

Provision/reprovision machine. The "provision" command performs tasks such as creating the runtime, starting and stopping services, bootstrapping (fresh install) or migrating (upgrade) data, initializing services and so on.

If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "Permissive" mode (recommended) before running the `/opt/collabnet/teamforge/bin/teamforge provision` command. TeamForge create runtime fails otherwise.



**Important:** If you have SELinux in "enforcing" mode, you must either disable SELinux or switch to "Permissive" mode (recommended) before running the /

```
opt/collabnet/teamforge/bin/
teamforge provision command.
TeamForge create runtime fails otherwise.
```

<b>undeploy</b>	Undeploy service(s).
<b>start</b>	Start service(s).
<b>stop</b>	Stop service(s).
<b>verify</b>	Verify environment.
<b>show-endpoints</b>	Show endpoints
<b>show-dependencies</b>	Show deployment dependencies
<b>reinitialize</b>	Reinitializes all the TeamForge services
<b>snapshot</b>	Dumps relevant diagnostic information to the console (stdout) for each deployed service.

## Troubleshooting

### Deploy error while upgrading TeamForge

The following error message might show up when you run the `/opt/collabnet/teamforge/bin/teamforge provision` command during TeamForge upgrades.

```
/opt/collabnet/teamforge/service/
legacy-runtime/hooks/runtime-option-
defaults/tomcat/00-scm-server-check.py
failed: No module named constants
```

This happens if the runtime folder is not found in `/opt/collabnet/teamforge/`. In such cases, run the following command:

```
find /opt/collabnet/teamforge/dist/
lib/multiplatform/python/ -name
*.pyc -exec rm -f {} +
```

## Logging

teamforge script writes entries to `/opt/collabnet/teamforge/log/runtime/runtime.log` file.

## add\_eventq\_to\_projects.py

Run the `add_eventq_to_projects.py` script to add EventQ to existing projects.

## Usage

Use the following command to run the this script.

```
[RUNTIME_DIR]/scripts/add_eventq_to_projects.py [--tfuser=<username>] [--
tfpasswd=<password>]]
```

## Parameters

Paramater	Description
-u   --tfuser	The TeamForge site administrator user name.
-p   --tfpasswd	The password for TeamForge user.

Parameter	Description
-s   --silent	To prevent printing progress message on standard-out.
-h   --help	To view help information for the script.

## backup-data.py

The `backup-data.py` script makes a compressed copy of the site's data. The script deletes any backup existing in the destination directory. By default, the TeamForge and the reporting database are backed up to the destination directory. If reporting is disabled, only the TeamForge database is backed up.

### Overview

`backup-data.py` makes compressed copies of these data resources:

- Subversion repositories
- CVS repositories
- The data directory ( `/var` )
- The TeamForge database

### Usage

```
./backup-data.py --destination_dir=<directory name>
```

### Location

```
/opt/collabnet/teamforge/runtime/scripts/
```

### Options

#### **--destination\_dir**

Directory where the compressed copy of the site data is to be stored.

#### **--help | -h**

Print this usage message and exit.

## backup-rb-data.py

The `backup-rb-data.py` script is used to back up the Review Board application data.

### Overview

The Review Board application data includes the database and files. If there are any files in the backup directory, the script overwrites these files.

### Usage

Run this script as follows:

```
python ./backup-rb-data.py --backupdir={dir}
```

### Options

The following options are available for the `backup-rb-data.py` script:

- |                        |   |
|------------------------|---|
| <b>-b   --both</b>     | Back up both the database and the filesystem. This is the default option. |
| <b>-d   --database</b> | Back up the database.   |
| <b>-f   --files</b>    | Back up the filesystem.   |

**-h | --help**

Provides a list of all available options for this script.

## db.py

The `db.py` script can be used to dump and restore a PostgreSQL database.

### Overview

This script can be used only for the PostgreSQL service. Don't run this script on a remote database. Execute the script only when the database is up and running.

### Usage

Run this script as follows:

```
./db.py --action=<action> --path=<destination directory>
```

### Options

Required options:

**-a | --action**

Values: dump or restore

**-f | --path**

Path where the database backup file will be created. Must be a directory owned by the postgresql user (usually `/var/lib/pgsql/9.0/`). Can be a new directory.


Optional options:

**-t | --type**

Specifies the type of database (ctf or reporting).


**-h | --help**

Print this usage message and exit.

 **Note:** The options `-a` and `-f` are mandatory.

## domain\_change\_db.py


The `domain_change_db.py` script handles all the steps required to change the domain name in the site database. It does not change anything in the file system.

 **Note:** Changing the domain through any other mechanism may cause problems.

### Usage

Execute this script with a command like this:

```
[RUNTIME_DIR]/domain_change_db.py [--debug] [--dir] --old={domain_name} --new={domain_name}
```

 **Note:** The new domain name must match the value defined for the DOMAIN token in the `site-options.conf` file.

### Options

The `domain_change_db.py` script provides the following parameters:

**--help**


Show this message

**--debug**

Include debugging output

**--old**

Old domain

<b>--new</b>	New domain
<b>--dir</b>	Run domain change in this directory only. You must specify the full path.  Use this feature to do a subset of the data directory. This instructs the script to do a recurse in the specified directory looking for the old <code>domain_name</code> and replacing it with the new <code>domain_name</code> .   <b>Note:</b> Without this option, only HTML, text, and VM files are modified.
<b>--threadlimit</b>	Defines the maximum number of simultaneous threads that can be invoked by this program. The default value is '50'

## etl-client.py

The `etl-client.py` script allows you to access the Extract, Transform and Load (ETL) scheduler and check the status of the jobs configured. The script also supports triggering jobs manually.


### Parameters

The following parameters are available for the `etl-client.py` script:

<b>-s   --status</b>	Prints the status of all the jobs configured in the ETL service.
<b>-a   --status-all</b>	Prints the status of incremental and historical jobs configured in the ETL service.
<b>-v   --verbose</b>	Chronicles the process of requested operation a bit more.
<b>-r   --run=</b>	Triggers a job manually for a given job.

### Data collection (ETL) jobs supported by the ETL service

While some ETL jobs are scheduled to run automatically, some must be triggered manually. The following table lists the available ETL jobs in TeamForge.

Job Category	Job Name	Description
<b>Historical Data Collection</b>  Historical data collection jobs must be triggered manually. As a best practice, these jobs are run as part of post migration activities. Refer to the "Related Links" for more information.	SCMCommitInitialJob	Collects the historical commit data from TeamForge.
	TrackerInitialJob	Collects the historical data of artifacts from TeamForge.
	LoadFlexFields	Collects the historical data of custom-defined artifacts from TeamForge.
	 <b>Note:</b> This job must be executed if and only if you are upgrading from TeamForge 8.0 and earlier versions.	
<b>Incremental Data Collection</b>	SCMCommitActivityJob	Collects the SCM commit data incrementally on an ETL run-to-run basis.



Job Category	Job Name	Description
Incremental data collection jobs collect data that are added or modified incrementally on an ETL run-to-run basis. These jobs are scheduled to run automatically on a regular basis.	TrackerIncrementalJob	Collects the tracker artifacts data incrementally on an ETL run-to-run basis.
	UserLoginActivityJob	Collects the user logon activity data incrementally on an ETL run-to-run basis.
<b>Imported Data Collection (Simbel)</b> TeamForge supports bulk data import through Simbel. This job collects Simbel-imported data. This job must be triggered manually post data import into TeamForge.	SimbelImportJob	Collects all Simbel-imported data such as the user logon activity, SCM commit and tracker artifacts data.

## fix\_data\_permission.sh

The `fix_data_permission.sh` script is used to fix the TeamForge filesystem permissions when you backup and restore the TeamForge data in a new box during TeamForge upgrade.

### Usage

Run this script as follows:

```
fix_data_permissons.sh [-r <repo_thread_count>>] [-d <dir_thread_count>] [-t <fix_permission item>]
```

### Example

Run the following command to fix the permission of GIT repositories.

```
./fix_data_permissons.sh -t git
```

Run the following command to fix the permission of GIT repositories and app directories.

```
./fix_data_permissons.sh -t git -t app
```

Run the following command to fix the permission of 15 Subversion repositories simultaneously.

```
./fix_data_permissons.sh -r 15 -t svn
```

### Options

The following options are available for the `fix_data_permission.sh` script:

<b>-r</b>	Number of repos to process at a time. Default=5
<b>-d</b>	Number of directories to process at a time. Default=<No of processors>
<b>-t</b>	Item to fix. Possible values are [app, http, svn, cvs, git]. Default=<all items>
<b>-h   --help</b>	Prints the help information.

## install.sh

The `install.sh` script handles all operations related to installing and removing the application.

### Options

**-a | --all**

Performs these operations in sequence - Install, create runtime and setup initscripts (equivalent to: -I -r -i)

**-b | --bootstrap**

Sets up initial site data.

**-C | --cleanup**

Stop/Kill the application processes, wipes out application packages and site directory.



**Caution:** THIS OPTION WILL WIPE OUT THE SITE DATA.

**-c | --conf-file**

The environment configuration file (usually of the form environment-<platform>.conf), used to identify the platform.



**Note:** This is not to be confused with the the `site-options.conf` file.

**-d | --directory type='str', argname='installation\_base\_dir',**

Path where the site would get installed.

**-E | --check\_environment**

Check if the system environment is suitable for this installation.

**-F | --force**

Force the operation that is performed wherever appropriate (e.g. install/uninstall).

**-f | --siteoptionsfile**

Path to the site configuration file (default: `./conf/site-options.conf`).

**-I | --initscript**

Start application services on reboot.

**-i | --install**

Install application packages.

**-n | --noninteractive**

Used to run the installer in non-interactive mode.

**-R | --internalruntime**

Create the internal runtime instance using the site configuration file.

**-r | --runtime**

Create the runtime instance using the site configuration file. (Also does configuration for Apache and PostgreSQL.).

**-S | --startnow**

Start application services after completing the other operations (if possible).

**-s | --startup**

Start application services on reboot and start it now: equivalent to -I -S.

**-u | --uninstall**

Uninstall application packages.

**-V | --verbose**

Show all output in noninteractive mode.

## pbl.py

The `pbl.py` utility enables you to upload files to the Project Build Library and perform various operations on them.

## Options

<code>--help -h</code>	Print out a help message and exit.
<code>--api-user -u username</code>	Your TeamForge Lab Management login name. Required for all upload operations.
<code>--api-key -k key</code>	Your TeamForge Lab Management API key. Required for all upload operations.
<code>--api-url -l url</code>	The URL to the TeamForge Lab Management API's. Will generally be <code>https://\$external_host/cubit_api/1</code> . Required for all upload operations.
<code>--comment -c "your comment"</code>	Print out a comment on this operation. The comment is always optional. The comment string will be logged in the audit log, but is not recorded in the PBL. For example, if you are deleting some files, you might want to use a comment to explain why you were deleting those files, for future auditing purposes.
<code>--verbose -v</code>	Print out more detail on what the <code>pbl.py</code> is doing.
<code>--xml-server-output --xml</code>	If this option is not specified, the <code>pbl.py</code> client reads in the XML returned from the server and presents the results to you in nicely formatted text. If you'd like to instead see the raw XML returned from the server, select this option.
<code>--no-auth-cache</code>	As a convenience, the <code>pbl.py</code> function caches the value of the <code>--api-user</code> and <code>--api-key</code> parameters in your home directory, in a subdirectory named <code>.cubit</code> , the first time a successful authentication is performed against the server. This is analogous to the Subversion client's use of the <code>.subversion</code> directory to store authentication credentials. Selecting the <code>--no-auth-cache</code> option turns off this caching.
<code>--project p projname</code>	The TeamForge Lab Management project in which the file you are operating on is located.
<code>--type t {pub priv}</code>	The visibility type of the file, either <code>pub</code> (the file is in the public area of the PBL) or <code>priv</code> (the file is in the private area of the PBL).
<code>--remotepath r path</code>	The remote path on the server, excluding the base directory, the project, and the visibility type. Examples are below.

## password\_util.sh

The `password_util.sh` script is used to get the encrypted or decrypted password value for the user `scmviewer`.

### Usage

To encrypt:

```
sudo /opt/collabnet/teamforge/runtime/scripts/password_util.sh -encrypt 'teamforge'
```

```
[root@xx scripts]# ./password_util.sh -encrypt 'teamforge'
```

```
Input String:teamforge
Encrypted password:VBxJJvzbXb5tNx2SxR26egA==
```

To decrypt:

```
sudo /opt/collabnet/teamforge/runtime/scripts/password_util.sh -decrypt
'VBxJJvzbXb5tNx2SxR26egA=='
```

```
[root@xx scripts]# ./password_util.sh -decrypt 'VBxJJvzbXb5tNx2SxR26egA=='
Input String:VBxJJvzbXb5tNx2SxR26egA==
Decrypted password:teamforge
```

## post-install.py

Run the `post-install.py` script to perform TeamForge post installation/upgrade actions. For example, run this script for setting up CLI reports, integrating Gerrit SCM with TeamForge, integrating SVN and CVS, and so on. This script is intelligent enough to identify the list of post installation actions required based on the TeamForge applications that are configured and running locally.

`post-install.py [-D|--debug] [-r|--reconfigure]`

Option	Description
<code>-D --debug</code>	Use this option to enable debug mode. Enabling the debug mode displays all the back-end debug messages on the console.
<code>-r --reconfigure</code>	To reconfigure the post installation with default credentials.
<code>-O  --option</code>	The <code>post-install.py</code> script prompts for input for items such as TeamForge admin user name, password, etc. Use this option to pass such parameters as command line arguments so that you are not prompted for the same later.  Valid values:  <pre>teamforge- tf.siteadmin.username=&lt;USERNAME&gt; teamforge- tf.siteadmin.passwd.secure=&lt;PASSWORD&gt; gerrit- tf.gerritdb.passwd.secure=&lt;PASSWORD&gt;</pre>
<code>-h --help</code>	Use this option to display help information about <code>post-install.py</code> .

### Usage examples

- Run the following command with no arguments to run the all post installation actions based on TeamForge applications that are configured in the current server.

```
./post-install.py
```

- Run the following command to reconfigure Gerrit integration.

```
./post-install.py -r gerrit
```

## Manually create SCM integration in case your integration server uses self-signed (or wildcard) certificates

If your integration server uses self-signed (or wildcard) certificates, the `post-install.py` script may fail to create SCM integration. In such case, you may have to manually create the SCM integration using the TeamForge **Admin > Integrations > SCM Integrations** UI. The following message is shown when you run the `post-install.py` script on integration servers that use self-signed (or wildcard) certificates.

```
Please create SVN integration through TeamForge Admin UI if your integration
server uses self signed certificate. Failed.
```

### Troubleshooting

See [Post install fails for Git. What should I do?](#) on page 954 for more information.

## psql-wrapper

The `psql-wrapper` script is used to connect to the TeamForge application.


### Usage

Run this script as below:

```
sudo [RUNTIME_DIR]/scripts/psql-wrapper
```

### Comments

- Run this script as a sudo user.
- Run this script with the postgres backend.
- You have full write access to the database for executing queries.

 **Note:** This script is not supported in the Oracle backend.

## psql-reporting-wrapper

The `psql-reporting-wrapper` script is used to connect to the datamart.


### Usage

Run this script as below:

```
sudo [RUNTIME_DIR]/scripts/psql-reporting-wrapper
```

### Comments

- Run this script as a sudo user.
- Run this script with the postgres backend.
- You have full write access to the datamart for executing queries.

 **Note:** This script is not supported in the Oracle backend.

## restore-data.py

The `restore-data.py` script restores the compressed data from the named source directory and deletes any existing data. By default, the TeamForge and the reporting database are backed up to the destination directory. If reporting is disabled, only the TeamForge database is backed up.

## Overview

`restore-data.py` finds and unpacks these data resources:

- Subversion repositories
- CVS repositories (if any)
- The data directory ( `/var` )
- The SourceForge database.

## Usage

Run this script as follows:

```
./restore-data.py --source=<directory name>
```

where `<directory-name>` is the directory to which you backed up the data with the `backup-data.py` script.

## Location

```
/opt/collabnet/teamforge/runtime/scripts/
```

## Options

### `--source`

Directory where the compressed copy of the site's data is available.

### `--help | -h`

Print this usage message and exit.

## restore-rb-data.py

The `restore-rb-data.py` script is used to restore the Review Board application data from the backup directory.

## Overview

This script removes the existing Review Board application data present in the system and restores data from the backup directory.

## Usage

Run this script as follows:

```
python ./restore-rb-data.py -backupdir={dir}
```

## Options

The following options are available for the `restore-rb-data.py` script:

- |                              |   |
|------------------------------|---|
| <code>-b   --both</code>     | Restore both the database and the filesystem. This is the default option. |
| <code>-d   --database</code> | Restore the database.   |
| <code>-f   --files</code>    | Restore the filesystem.   |
| <code>-h   --help</code>     | Provides a list of all available options for this script.                 |

## SearchReindex.py

The `SearchReindex.py` script allows you to reindex the entire TeamForge data.

## Overview

You can use this script to reindex the entire TeamForge data or you can choose to reindex the subset of data types.

## Usage

Run this script as follows:

```
./SearchReindex.py --<component name>
```

## Example

To perform a search reindex for the tracker, run this command:

```
./SearchReindex.py --trackers-only
```

To perform a search reindex for the wiki, run this command:

```
./SearchReindex.py --wiki-only
```

To perform a search reindex for documents run this command:

```
./SearchReindex.py --documents-only
```

## Options

<b>--single-item</b> <i>itemId</i> ,   <b>-i</b>	Schedules a re-index for just the given item. If the item id is for a project the scheduling results in the server re-indexing all of the project data.
<b>---force-index</b>   <b>-f</b>	Force indexing (doesn't check if item is searchable already).
<b>--artifacts only</b>   <b>-a</b>	Reindex all artifacts on the site that are currently not searchable or all artifacts if option <b>f</b> is selected.
<b>--documents only</b>   <b>-d</b>	Reindex all documents on the site that are currently not searchable or all artifacts if option <b>f</b> is selected.
<b>---posts only</b>	Reindex all posts on the site that are currently not searchable or all artifacts if option <b>f</b> is selected.
<b>---trackers only</b>	Reindex all trackers on the site.
<b>---document_folders-only</b>	Reindex all document folders on the site.
<b>---topics-only</b>	Reindex all topics.
<b>---forums-only</b>	Reindex all forums on the site.
<b>---news-only</b>	Reindex all news.
<b>---project_pages-only</b>	Reindex all project pages.
<b>---packages</b>	Reindex all packages.
<b>---commits-only</b>	Reindex all commits.

<code>---frs_files-only</code>	Reindex all frs files.
<code>---releases-only</code>	Reindex all releases.
<code>---wikis-only</code>	Reindex all wikis.
<code>--project-id projectID   -p</code>	Limit the re-indexing to data for single project when re-indexing only artifacts and or documents.
<code>--verify   -x</code>	Searches for each item that is scheduled for re-indexing. There is a one minute wait limit for each item to be re-indexed by the server.
<code>--dryrun   -r</code>	Executes all the steps for scheduling a re-index without actually sending any re-index requests to the server. This provides a list of items that need re-indexing.
<code>output-file filePath,   -o</code>	Prints the output for the given file.
<code>--verbose   -v</code>	Chronicles the process of scheduling the re-index a bit more.

## set\_auth\_key.py

The `set_auth_key.py` script sets the authorized key to the scmviewer user profile.

### Overview

The script accepts the authorized key file, reads the key from the file and sets it to the scmviewer user profile.

```
$ sudo /opt/collabnet/teamforge/runtime/scripts/codesearch/set_auth_key.py
--help
```

### Usage

Run this script as follows:

```
./set_auth_key.py --authkey-file=<path_of_authkey_file>
```

When you run the script, you will be prompted to enter the TeamForge site-admin credentials to update the key for the scmviewer user.

## sync-orchestrate.py

Run the `sync-orchestrate.py` script to synchronize a TeamForge project with Orchestrate by triggering the world-state message for the given project.

### Usage

Use the following command to run the this script.

```
[RUNTIME_DIR]/scripts/sync-orchestrate.py --projectId=<proj1234> [--
tfuser=<username>] [--tfpasswd=<password>]]
```

### Parameters

Paramater	Description
<code>--projectId</code>	The TeamForge project ID.



Paramater	Description
-u   --tfuser	The TeamForge site administrator user name.
-p   --tfpasswd	The password for TeamForge user.
-h   --help	To view help information for the script.

## Log files in TeamForge

System administrators can use logs to debug problems and ensure that the application is performing to expectations.

### JBoss logs

The JBoss application server writes several different logs under the <SOURCEFORGE\_INSTALL\_DIR>/log directory.

<b>boot.log</b>	Lgs the JBoss startup and shut down notifications. This log is overwritten each time JBoss is (re)started.
<b>localhost_access</b>	The Records access to the application from a remote host, similar to the Apache access_log. This log is rotated each day, and the files have a date stamp appended to their name, such as localhost_access2004-11-26.log.
<b>server.log</b>	Logs all the activities of the application server, including any exceptions. This log is the best place to begin debugging CollabNet TeamForge server error exception ids (exid).
<b>session-info.log</b>	Records when new sessions are created. This log is overwritten each time JBoss is (re)started.
<b>vamessages.log</b>	Records CollabNet TeamForge -specific actions, including some SQL queries that are sent to the backend database. This log is rotated each time it reaches 100MB in size. When rotated the older files have a number appended to the end, such as vamessages.log.1 and vamessages.log.2.

### Oracle logging

The most important Oracle log is the `alert` log, which is found in `$ORACLE_HOME/admin/$SID/bdump/alert_$SID.log`.

An Oracle database performs logging on a wide array of functionality. The majority of the logs that are generated are stored under `$ORACLE_HOME/admin/$SID/`. Many logs are stored under this directory hierarchy, but `alert` is the most important. This log records all database activity, including serious problems.

The `alert` log is not rotated or overwritten, and can become quite large over time, especially on an active database.

Additional logs are created under the same directory hierarchy, for specific incidents. If a problem is recorded in the `alert` log, the other logs should be inspected for additional details.

For more information, as well as support in the maintenance of an Oracle database, contact Oracle Support or Oracle's [Metalink](#) site.

### SCM (CVS and Subversion) logs

Software configuration management (SCM) servers generate several logs from the CollabNet TeamForge ; however, in the interest of completeness they are all documented here.

<b>catalina.out</b>	This log contains information on the startup and runtime activities of the Tomcat server. This log is not rotated, nor is it overwritten, and is appended continuously over the lifetime of the server.
<b>localhost_log</b>	This log contains a record of CVS or Subversion browsing URL construction. When a user attempts to browse a CVS or Subversion repository in his or her web browser, the URL construction process is documented in this log. This log is rotated for each date that there is activity.
<b>localhost_admin_log</b>	This log contains a record of the initial startup and deployment of the managed integration server. A new date stamped log is generated each time the integration server is started.
<b>vaexternalintegration.log</b>	This log contains information on the operations that are being executed by the managed integration server. This log is stored in <SOURCEFORGE_INSTALL_DIR>/log.

## Email logs

Both the CollabNet TeamForge email and search backends are managed from a parent daemon known as Phoenix. If the mail backend is not operating properly, the first troubleshooting step is to check the phoenix.log to see if it encountered difficulties starting up.

### Overview

The Phoenix daemon logs its activities to the phoenix.log file, which is stored under SOURCEFORGE\_INSTALL\_DIR/james/james-2.2.0/logs. This log is overwritten each time Phoenix is (re)started. Phoenix is run as part of the CollabNet TeamForge standalone server init script.

CollabNet TeamForge email is handled by the JAMES server. JAMES logs all of its activities under SOURCEFORGE\_INSTALL\_DIR/james/james-2.2.0/apps/james/logs. A new log is created for each date when there is activity, and additional logs are created if james is restarted on a date when there is activity. The date is embedded in the log name (such as james-2005-04-28-01-00.log).

### Active logs

Sixteen different logs are created by james for different components of its functionality. This topic describes only the ones that are used actively by CollabNet TeamForge .

<b>james-<math>\\$</math>date.log</b>	The James log records the overall mail handling behavior of the James server.
<b>mailet-<math>\\$</math>date.log</b>	The mailet log records how each piece of email is handled. If there is a mail delivery problem, this log is the best place to begin investigation.
<b>mailstore-<math>\\$</math>date.log</b>	The mailstore log records the behavior of mail spools, and the storage of mail. This log should normally not contain errors unless James is unable to write or read mail to or from the file system.
<b>smtpserver-<math>\\$</math>date.log</b>	The smtpserver log records all inbound mail handling results. If email to discussion forums is not posting, or is getting rejected, this log would be the best place to begin investigation.

**spoolmanager-`$date`.log**

The spoolmanager log records the processing of mail spools. This log could be of value in troubleshooting mail delivery or handling problems.

**Search logs**

Both the CollabNet TeamForge search and email backends are managed from a parent daemon known as Phoenix. If the search backend is not operating properly, the first troubleshooting step is to check the `phoenix.log` file to see if it encountered difficulties starting up.

The Phoenix daemon logs its activities to the `phoenix.log` file, which is stored under `SOURCEFORGE_INSTALL_DIR/james/james-2.2.0/logs`. This log is overwritten each time Phoenix is (re)started.

Phoenix is run as part of the CollabNet TeamForge standalone server init script.

Once started successfully, the search server waits for new content to be indexed or searches to be performed. The search server logs its activities under `SOURCEFORGE_INSTALL_DIR/james/james-2.2.0/apps/search/logs`. The logs that are created are all named `default` with the date stamp appended to them (such as `default-20041126.log`). A new log is created for each date that there is indexing activity.

If the search server is not running, or expected search results are not being provided, the default log is the best place to investigate further.


**Project Build Library audit log**

You can use this screen to view the complete list of actions performed in the Project Build Library.

**Contents**

Information about the following types of actions is displayed in this screen:

- Change a File Description
- Create a Directory
- Delete a File or Directory
- Download a File
- Move a File or Directory
- Upload a File

 **Note:** The value displayed in the **Event** field is the value passed in the `--comment` parameter from the Project Build Library client.

**Getting there**

On the project home page, click **Build Library** in the left navigation bar and select the **Audit Log** tab.

**Access**

This screen is accessible for all users who have at least the view permission for the project.

**Profile audit log**

Use this screen to view the complete list of actions performed on a profile.

**Getting there**

On the **Profile Library** screen, click the **Audit Log** tab.

## Access

This screen is accessible for all users who have at least the view permission for the project to which the profile is allowed.

x

## Example

When a user updates any of the profile fields on the **Profile Admin** screen, the following details are displayed in this screen:

- The old value for the field.
- The new value for the field.
- The name of user who updated the field.
- The time when the change occurred.

## User Audit Log

You can use this screen to view the list of actions performed by the user in the TeamForge Lab Management system.

For example, when a user logs into the web interface of the TeamForge Lab Management system, the event is displayed in this screen.

## Access

This screen is accessible to all users who have at least the Domain Administrator role.

## Getting there

On the **Administration** tab, click **User Audit Logs** in the left navigation bar.

## Host audit log

You can use this screen to view the complete list of actions performed on a host.

## Getting there

On the TeamForge Lab Management Host home page, click the **Audit log** tab.

## Access

This screen is accessible for all users who have at least the view permission for the project to which the host is assigned.

## Example

When the IP address for the host is changed, the following details are displayed in this screen:

- The old IP address.
- The new IP address.
- The name of user who changed the IP address.
- The time when the change occurred.

## Project audit log

The **Project audit log** screen shows the complete list of changes applied to a project.

## Getting there

On the TeamForge Lab Management Project home page, click **Audit Logs** in the left navigation bar.

## Access

This screen is accessible for all users who have at least the view permission for the project.


## Example

When a profile is added to the list of buildable profiles for this project, the following information appears on this screen:

- The action that was taken.
- The user who performed the change.
- The time when this change occurred.

## etl.log

This file contains information from extract-transform-load runs, including data transformation warnings and errors.

 **Note:** Transformation errors do not constitute a failed ETL run. For example, if a corrupt row of data in one of the source tables causes transformation errors, this is treated as a "skipped record" and gets logged.

## Reference information about integrated applications

---

Here is some stuff you may need to know to work with integrated applications.

### IntegratedAppSupport

CollabNet provides servlet helper libraries that take care of authenticating integrated application requests with TeamForge. They also provide context information for each subsequent form processing within the application or links clicked within the application.

The heart of this is the `IntegratedAppSupport` class, which must be called during every request of an integrated application. This class takes the `HttpServletRequest` and `HttpServletResponse` for each request and determines whether the user is already authenticated. It also provides project- and user-related information that can be used throughout the request.

 **Tip:** It is a good idea to store this as a `ThreadLocal` so that it can be used from anywhere in the application.

`IntegratedAppSupport`, on successful validation, provides several objects that can be used through the integrated application. .

### Parameters

The method is specified within parentheses at the end of each parameter.

#### SoapSessionId

This identifies a user. It can be used for making any TeamForge webservice calls over SOAP. TeamForge expects that every call be associated with a valid `SoapSessionId`. This method lets you fetch the `SoapSessionId` for the current session of the user. (`getSoapSessionId()`)

#### WebSessionId

This is the `jsessionid` for the current user. This is the servlet-container-specific id that might be used for some integrated applications. (`getWebSessionId()`)

#### CtfBaseUrl

Retrieves the base URL for Collabnet TeamForge (`getCtfBaseUrl()`).

#### ProjectPath

Retrieves the `TeamForgeProject` path for this project. This would typically be "projects.<projectname>" (`getProjectPath()`)


<b>CollabNetSoap handle</b>	Retrieves the handle for CollabNet SOAP (Refer to ctf_xxx_sdk.zip for SOAP calls that can be made using TeamForge) (getCollabNetSoap()). Documentation for the SOAP methods exposed are available in ctf_xxx_sdk.zip.
<b>PluggableAppSoap handle</b>	Retrieves the handle for pluggable application SOAP calls (getPluggableAppSoap()). Documentation for the SOAP methods exposed are available in ctf_xxx_sdk.zip.
<b>RbacAppSoap</b>	Retrieves the handle for RBAC SOAP Calls (getRbacAppSoap()). Documentation for the SOAP methods exposed are available in ctf_xxx_sdk.zip.
<b>IntegratedAppId</b>	Retrieves the linkid for this request. (getIntegratedAppId())
<b>IntegratedAppPrefix</b>	Retrieves the Prefix to use for go-urls, associations and linkifications (getIntegratedAppPrefix()).
<b>IntegratedAppName</b>	Retrieves the name of this integrated app (getIntegratedAppName()).
<b>ProjectId</b>	Retrieves the TeamForge project id (getProjectId()).
<b>userSoapDO</b>	Retrieves the SOAP data object information for the current user (getUserSoapDO()).

## SOAP calls for integrated applications

When you integrate an application with TeamForge, you will need to make SOAP calls from the integrated application into TeamForge to respond to events on either side.

To integrate an application with TeamForge using the Integrated Application Framework, you have to implement the following methods for TeamForge to communicate with the integrated application.

The interface should be exposed via SOAP, and the endpoint for this interface should be defined by the `<endpoint>` tag in the XML Application configuration file.

 **Tip:** There is no direct API for managing users/groups from TeamForge to an integrated application. However, if your integrated application has the notion of users and user groups, you can listen to “user events” or “group events” through an event handler in TeamForge and call this interface using that data. You can also make SOAP calls from the integrated application directly into TeamForge to get user information.

### Javadoc

For a complete description of all methods, see the [Javadoc for the Pebble integration service](#).

### SOAP calls for adding, editing or deleting an integrated application

You can use these SOAP calls to enable project administrators to add an integrated application to their projects, edit an integrated application, or delete it from their projects.

#### createProjectConfig

This method is called when an integrated application is added to a project from **Project Admin > Tools > Add Tool**.

#### Example

```
In com/collabnet/pebble/adapters/
PebbleIntegrationService.java, see the
createProjectConfig method, line 48.
```

#### editProjectConfig

This method is called when editing an integrated application in a project from **Project Admin > Tools**.

**Example**

In `com/collabnet/pebble/adapter/PebbleIntegrationService.java`, see the `editProjectConfig` method, line 202.

**deleteProjectConfig**

This method is called when deleting an integrated application in a project from **Project Admin > Tools**.

**Example**


In `com/collabnet/pebble/adapter/PebbleIntegrationService.java`, see the `deleteProjectConfig` method, line 213.

**SOAP calls for SCM in integrated applications**

You can use these SOAP calls to enable users to change data in an integrated application based on actions in a TeamForge source code repository.

**scmPreCommit**

This serves as the pre-Commit hook for an integrated application. TeamForge will call this method if `require-scm-integration` flag is set to true in the integrated app xml descriptor. This would mean that a commit made into TeamForge will only succeed if this method returns a “true” (String).

 **Note:** All TeamForge ids and ids from other integrated applications are removed from this list even though it could be part of the original commit message. Also these ids are without the prefix. The integrated application can use these ids to find if they are valid for the particular commit and respond with a “true” or a “false” followed by an optional error message that can be displayed to the user who is making the commit.

**scmPostCommit**

This method is called after a commit is made and post processing needs to happen for that particular commit in the integrated application. A typical use case will be to store the commit information as part of some object within the integrated application. Please note that any errors thrown as part of this will not stop the commit from happening as the commit has been already made.

**SOAP calls for search in integrated applications**

The `getSearchResults` method is called by TeamForge as part of doing a regular search. Use this method to include data from your integrated application in TeamForge search results.

**Example**

In `com/collabnet/pebble/adapter/PebbleIntegrationService.java`, see the `getSearchResults` method, line 409.

**SOAP calls for project templates in integrated applications**

You can use these SOAP calls to enable an integrated application to be included in project templates created by TeamForge.

**createTemplate**

This method is called when a TeamForge template is created on a project that contains this integrated application.

**Example**

In `com/collabnet/pebble/adapter/PebbleIntegrationService.java`, see the `createTemplate` method, line 334.

**getTemplateMetadata**

This method is called when displaying the template content in “Project Tools/Included in Template” section at the time of template creation.

**Example**

In `com/collabnet/pebble/adapter/PebbleIntegrationService.java`, see the `getTemplateMetadata` method, line 351.

**getTemplateContent**

This method is called when viewing the content of an existing template.

**Example**

In `com/collabnet/pebble/adapter/PebbleIntegrationService.java`, see the `getTemplateContent` method, line 367.

**validateParametersForTemplatizedProject**

This method is called to validate the configuration parameters provided when creating a project from a template.

**Example**

In `com/collabnet/pebble/adapter/PebbleIntegrationService.java`, see the `validateParametersForTemplatizedProject` method, line 381.

**createTemplatizedProjectConfig**

This method is called when a project is created from a template and this integrated application is part of the template. This is the equivalent of the “`createProjectConfig`” call except that it can be called when a project gets created from a template.

**Example**

In `com/collabnet/pebble/adapter/PebbleIntegrationService.java`, see the `createTemplatizedProjectConfig` method, line 395.

**SOAP calls for objects in integrated applications**

If your integrated application has the notion of objects, you can use these SOAP calls to enable users to act on those objects from within TeamForge.

**getObjectTitle**

Return the title for the provided object id. This will be used when associating an integrated app object to a TeamForge object. If the user enters an integrated application id, say, `PB_20100602`, the first part of it is the per-project prefix (PB) and the second part is the actual object id in the integrated application. This is the value passed into “`getObjectTitle`” so that it can return the Title for the object identifier. The title is then displayed as part of the association information in TeamForge.

**Example**

In `com/collabnet/pebble/adapter/PebbleIntegrationService.java`, see the `getObjectTitle` method, line 306.

**findProjectIdForObject**

Return the project to which the object belongs. This method is called only when the `require-per-project-prefix` parameter in the XML is false. When `require-per-project-prefix` is true, TeamForge knows the project to which the object belongs.

**Example**

In `com/collabnet/pebble/adapter/PebbleIntegrationService.java`, see the `findProjectIdForObject` method, line 294.



## Pebble

### Pebble integration service

Here's the WSDL for the Pebble blogging application's integration service.

```
<?xml version="1.0" encoding="UTF-8"?>
<wsdl:definitions targetNamespace="https://cul23.cubit.maa.collab.net:13000/pebble/services/PebbleIntegrationService" xmlns:apachesoap="http://xml.apache.org/xml-soap" xmlns:impl="https://cul23.cubit.maa.collab.net:13000/pebble/services/PebbleIntegrationService" xmlns:intf="https://cul23.cubit.maa.collab.net:13000/pebble/services/PebbleIntegrationService" xmlns:soapenc="http://schemas.xmlsoap.org/soap/encoding/" xmlns:tns1="http://util.integratedapps.soap60.ce.collabnet.com" xmlns:tns2="qConfigParam" xmlns:tns3="http://com.collab.net/pebble/XMLSchema" xmlns:tns4="qConfigParamDO" xmlns:tns5="qSearchInfo" xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/" xmlns:wsdlsoap="http://schemas.xmlsoap.org/wsdl/soap/" xmlns:xsd="http://www.w3.org/2001/XMLSchema">
<!--WSDL created by Apache Axis version: 1.4
Built on Apr 22, 2006 (06:55:48 PDT)-->
  <wsdl:types>
    <schema targetNamespace="qConfigParamDO" xmlns="http://www.w3.org/2001/XMLSchema">
      <import namespace="http://com.collab.net/pebble/XMLSchema"/>
      <import namespace="http://util.integratedapps.soap60.ce.collabnet.com"/>
      <import namespace="qSearchInfo"/>
      <import namespace="http://xml.apache.org/xml-soap"/>
      <import namespace="https://cul23.cubit.maa.collab.net:13000/pebble/services/PebbleIntegrationService"/>
      <import namespace="qConfigParam"/>
      <import namespace="http://schemas.xmlsoap.org/soap/encoding"/>
      <complexType name="qConfigParamDO">
        <sequence>
          <element name="defaultValue" nillable="true" type="soapenc:string"/>
          <element name="description" nillable="true" type="soapenc:string"/>
          <element name="displayType" nillable="true" type="soapenc:string"/>
          <element name="maxLength" type="xsd:int"/>
          <element name="name" nillable="true" type="soapenc:string"/>
          <element name="title" nillable="true" type="soapenc:string"/>
          <element name="valueType" nillable="true" type="soapenc:string"/>
        </sequence>
      </complexType>
    </schema>
    <schema targetNamespace="http://com.collab.net/pebble/XMLSchema" xmlns="http://www.w3.org/2001/XMLSchema">
      <import namespace="qConfigParamDO"/>
      <import namespace="http://util.integratedapps.soap60.ce.collabnet.com"/>
      <import namespace="qSearchInfo"/>
      <import namespace="http://xml.apache.org/xml-soap"/>
      <import namespace="https://cul23.cubit.maa.collab.net:13000/pebble/services/PebbleIntegrationService"/>
      <import namespace="qConfigParam"/>
      <import namespace="http://schemas.xmlsoap.org/soap/encoding"/>
      <complexType name="qArrayOfConfigParamDO">
        <complexContent>
          <restriction base="soapenc:Array">
            <attribute ref="soapenc:arrayType" wsdl:arrayType="tns4:qConfigParamDO[]" />
          </restriction>
        </complexContent>
      </complexType>
    </schema>
  </wsdl:types>

```

```

<schema targetNamespace="qConfigParam" xmlns="http://www.w3.org/2001/
XMLSchema">
  <import namespace="http://com.collab.net/pebble/XMLSchema"/>
  <import namespace="qConfigParamDO"/>
  <import namespace="http://util.integratedapps.soap60.ce.collabnet.com"/>
  <import namespace="qSearchInfo"/>
  <import namespace="http://xml.apache.org/xml-soap"/>
  <import namespace="https://cul23.cubit.maa.collab.net:13000/pebble/
services/PebbleIntegrationService"/>
  <import namespace="http://schemas.xmlsoap.org/soap/encoding"/>
  <complexType name="qConfigParam">
    <sequence>
      <element name="configParam" nillable="true"
type="tns3:qArrayOfConfigParamDO"/>
      <element name="deleted" type="xsd:boolean"/>
      <element name="locked" type="xsd:boolean"/>
      <element name="parentProjectId" nillable="true" type="soapenc:string"/>
      <element name="projectId" nillable="true" type="soapenc:string"/>
      <element name="projectName" nillable="true" type="soapenc:string"/>
      <element name="projectUrl" nillable="true" type="soapenc:string"/>
      <element name="template" type="xsd:boolean"/>
      <element name="userId" nillable="true" type="soapenc:string"/>
    </sequence>
  </complexType>
</schema>
<schema targetNamespace="http://xml.apache.org/xml-soap" xmlns="http://
www.w3.org/2001/XMLSchema">
  <import namespace="http://com.collab.net/pebble/XMLSchema"/>
  <import namespace="qConfigParamDO"/>
  <import namespace="http://util.integratedapps.soap60.ce.collabnet.com"/>
  <import namespace="qSearchInfo"/>
  <import namespace="https://cul23.cubit.maa.collab.net:13000/pebble/
services/PebbleIntegrationService"/>
  <import namespace="qConfigParam"/>
  <import namespace="http://schemas.xmlsoap.org/soap/encoding"/>
  <complexType name="mapItem">
    <sequence>
      <element name="key" nillable="true" type="xsd:anyType"/>
      <element name="value" nillable="true" type="xsd:anyType"/>
    </sequence>
  </complexType>
  <complexType name="Map">
    <sequence>
      <element maxOccurs="unbounded" minOccurs="0" name="item"
type="apachesoap:mapItem"/>
    </sequence>
  </complexType>
  <complexType name="Vector">
    <sequence>
      <element maxOccurs="unbounded" minOccurs="0" name="item"
type="xsd:anyType"/>
    </sequence>
  </complexType>
</schema>
<schema targetNamespace="https://cul23.cubit.maa.collab.net:13000/pebble/
services/PebbleIntegrationService" xmlns="http://www.w3.org/2001/XMLSchema">
  <import namespace="http://com.collab.net/pebble/XMLSchema"/>
  <import namespace="qConfigParamDO"/>
  <import namespace="http://util.integratedapps.soap60.ce.collabnet.com"/>
  <import namespace="qSearchInfo"/>
  <import namespace="http://xml.apache.org/xml-soap"/>
  <import namespace="qConfigParam"/>
  <import namespace="http://schemas.xmlsoap.org/soap/encoding"/>
  <complexType name="ArrayOf_xsd_anyType">

```

```

    <complexContent>
      <restriction base="soapenc:Array">
        <attribute ref="soapenc:arrayType" wsdl:arrayType="xsd:anyType[]" />
      </restriction>
    </complexContent>
  </complexType>
  <complexType name="ArrayOf_soapenc_string">
    <complexContent>
      <restriction base="soapenc:Array">
        <attribute ref="soapenc:arrayType" wsdl:arrayType="soapenc:string[]" />
      </restriction>
    </complexContent>
  </complexType>
</schema>
<schema targetNamespace="qSearchInfo" xmlns="http://www.w3.org/2001/
XMLSchema">
  <import namespace="http://com.collab.net/pebble/XMLSchema"/>
  <import namespace="qConfigParamDO"/>
  <import namespace="http://util.integratedapps.soap60.ce.collabnet.com"/>
  <import namespace="http://xml.apache.org/xml-soap"/>
  <import namespace="https://cul23.cubit.maa.collab.net:13000/pebble/
services/PebbleIntegrationService"/>
  <import namespace="qConfigParam"/>
  <import namespace="http://schemas.xmlsoap.org/soap/encoding"/>
  <complexType name="qSearchInfo">
    <sequence>
      <element name="attachment" nillable="true" type="soapenc:boolean"/>
      <element name="comment" nillable="true" type="soapenc:boolean"/>
      <element name="document" nillable="true" type="soapenc:boolean"/>
      <element name="keyWord" nillable="true" type="soapenc:string"/>
      <element name="plugAppName" nillable="true" type="soapenc:string"/>
      <element name="selectProject" nillable="true"
type="impl:ArrayOf_xsd_anyType"/>
    </sequence>
  </complexType>
</schema>
</wsdl:types>

  <wsdl:message name="validateParametersForTemplatizedProjectResponse">
</wsdl:message>

  <wsdl:message name="getPageComponentContentRequest">
    <wsdl:part name="oneTimeToken" type="soapenc:string"/>
    <wsdl:part name="projectId" type="soapenc:string"/>
    <wsdl:part name="projectIntegratedApplicationId"
type="soapenc:string"/>
    <wsdl:part name="pageComponentParam" type="soapenc:string"/>
    <wsdl:part name="rowsPerPage" type="soapenc:int"/>
    <wsdl:part name="locale" type="soapenc:string"/>
  </wsdl:message>

  <wsdl:message name="getSearchResultsRequest">
    <wsdl:part name="oneTimeToken" type="soapenc:string"/>
    <wsdl:part name="searchInfo" type="tns5:qSearchInfo"/>

```

```

    <wsdl:part name="locale" type="soapenc:string"/>
</wsdl:message>
<wsdl:message name="getPageComponentContentResponse">
    <wsdl:part name="getPageComponentContentReturn" type="soapenc:string"/
>
</wsdl:message>
<wsdl:message name="findProjectIdForObjectRequest">
    <wsdl:part name="oneTimeToken" type="soapenc:string"/>
    <wsdl:part name="objectId" type="soapenc:string"/>
</wsdl:message>
<wsdl:message name="createProjectConfigRequest">
    <wsdl:part name="oneTimeToken" type="soapenc:string"/>
    <wsdl:part name="iaConfig" type="tns2:qConfigParam"/>
</wsdl:message>
<wsdl:message name="createTemplateResponse">
</wsdl:message>
<wsdl:message name="getPageComponentParametersRequest">
    <wsdl:part name="oneTimeToken" type="soapenc:string"/>
    <wsdl:part name="projectId" type="soapenc:string"/>
    <wsdl:part name="locale" type="soapenc:string"/>
</wsdl:message>
<wsdl:message name="getSoapSessionIdRequest">
    <wsdl:part name="oneTimeToken" type="soapenc:string"/>
    <wsdl:part name="iAppSupport" type="xsd:anyType"/>
</wsdl:message>
<wsdl:message name="getObjectTitleResponse">
    <wsdl:part name="getObjectTitleReturn" type="soapenc:string"/>
</wsdl:message>
<wsdl:message name="editProjectConfigRequest">
    <wsdl:part name="oneTimeToken" type="soapenc:string"/>
    <wsdl:part name="iaConfig" type="tns2:qConfigParam"/>
</wsdl:message>

```

```

<wsdl:message name="findProjectIdForObjectResponse">
  <wsdl:part name="findProjectIdForObjectReturn" type="soapenc:string"/>
</wsdl:message>

<wsdl:message name="getConfigurationParametersRequest">
  <wsdl:part name="oneTimeToken" type="soapenc:string"/>
  <wsdl:part name="projectId" type="soapenc:string"/>
</wsdl:message>

<wsdl:message name="createProjectConfigResponse">
</wsdl:message>

<wsdl:message name="deleteProjectConfigRequest">
  <wsdl:part name="oneTimeToken" type="soapenc:string"/>
  <wsdl:part name="iaConfig" type="tns2:qConfigParam"/>
</wsdl:message>

<wsdl:message name="getConfigurationParametersResponse">
  <wsdl:part name="getConfigurationParametersReturn"
type="tns3:qArrayOfConfigParamDO"/>
</wsdl:message>

<wsdl:message name="getTemplateContentRequest">
  <wsdl:part name="oneTimeToken" type="soapenc:string"/>
  <wsdl:part name="templateName" type="soapenc:string"/>
  <wsdl:part name="locale" type="soapenc:string"/>
</wsdl:message>

<wsdl:message name="deleteProjectConfigResponse">
</wsdl:message>

<wsdl:message name="validateParametersForTemplatizedProjectRequest">
  <wsdl:part name="oneTimeToken" type="soapenc:string"/>
  <wsdl:part name="iaConfig" type="tns2:qConfigParam"/>
  <wsdl:part name="projectConfigurationParameters"
type="apachesoap:Map"/>
</wsdl:message>

<wsdl:message name="getObjectTitleRequest">
  <wsdl:part name="oneTimeToken" type="soapenc:string"/>
  <wsdl:part name="objectId" type="soapenc:string"/>

```

```

</wsdl:message>
<wsdl:message name="createTemplatizedProjectConfigResponse">
</wsdl:message>
<wsdl:message name="editProjectConfigResponse">
</wsdl:message>
<wsdl:message name="getTemplateMetadataResponse">
  <wsdl:part name="getTemplateMetadataReturn" type="apachesoap:Map"/>
</wsdl:message>
<wsdl:message name="getTemplateMetadataRequest">
  <wsdl:part name="oneTimeToken" type="soapenc:string"/>
  <wsdl:part name="locale" type="soapenc:string"/>
</wsdl:message>
<wsdl:message name="createTemplatizedProjectConfigRequest">
  <wsdl:part name="oneTimeToken" type="soapenc:string"/>
  <wsdl:part name="iaConfig" type="tns2:qConfigParam"/>
  <wsdl:part name="templateId" type="soapenc:string"/>
  <wsdl:part name="title" type="soapenc:string"/>
</wsdl:message>
<wsdl:message name="getSoapSessionIdResponse">
  <wsdl:part name="getSoapSessionIdReturn" type="soapenc:string"/>
</wsdl:message>
<wsdl:message name="getTemplateContentResponse">
  <wsdl:part name="getTemplateContentReturn" type="apachesoap:Map"/>
</wsdl:message>
<wsdl:message name="createTemplateRequest">
  <wsdl:part name="oneTimeToken" type="soapenc:string"/>
  <wsdl:part name="sourceProjectId" type="soapenc:string"/>
  <wsdl:part name="templateId" type="soapenc:string"/>
  <wsdl:part name="templateName" type="soapenc:string"/>
  <wsdl:part name="creatorId" type="soapenc:string"/>
</wsdl:message>
<wsdl:message name="getPageComponentParametersResponse">

```

```

    <wsdl:part name="getPageComponentParametersReturn"
type="impl:ArrayOf_soapenc_string"/>
  </wsdl:message>
  <wsdl:message name="getSearchResultsResponse">
    <wsdl:part name="getSearchResultsReturn" type="soapenc:string"/>
  </wsdl:message>
  <wsdl:portType name="PebbleIntegrationService">
    <wsdl:operation name="getSoapSessionId" parameterOrder="oneTimeToken
iAppSupport">
      <wsdl:input message="impl:getSoapSessionIdRequest"
name="getSoapSessionIdRequest"/>
      <wsdl:output message="impl:getSoapSessionIdResponse"
name="getSoapSessionIdResponse"/>
    </wsdl:operation>
    <wsdl:operation name="createProjectConfig"
parameterOrder="oneTimeToken iaConfig">
      <wsdl:input message="impl:createProjectConfigRequest"
name="createProjectConfigRequest"/>
      <wsdl:output message="impl:createProjectConfigResponse"
name="createProjectConfigResponse"/>
    </wsdl:operation>
    <wsdl:operation name="getConfigurationParameters"
parameterOrder="oneTimeToken projectId">
      <wsdl:input message="impl:getConfigurationParametersRequest"
name="getConfigurationParametersRequest"/>
      <wsdl:output message="impl:getConfigurationParametersResponse"
name="getConfigurationParametersResponse"/>
    </wsdl:operation>
    <wsdl:operation name="editProjectConfig" parameterOrder="oneTimeToken
iaConfig">
      <wsdl:input message="impl:editProjectConfigRequest"
name="editProjectConfigRequest"/>
      <wsdl:output message="impl:editProjectConfigResponse"
name="editProjectConfigResponse"/>
    </wsdl:operation>
    <wsdl:operation name="deleteProjectConfig"
parameterOrder="oneTimeToken iaConfig">
      <wsdl:input message="impl:deleteProjectConfigRequest"
name="deleteProjectConfigRequest"/>

```

```

    <wsdl:output message="impl:deleteProjectConfigResponse"
name="deleteProjectConfigResponse"/>

  </wsdl:operation>

  <wsdl:operation name="findProjectIdForObject"
parameterOrder="oneTimeToken objectId">

    <wsdl:input message="impl:findProjectIdForObjectRequest"
name="findProjectIdForObjectRequest"/>

    <wsdl:output message="impl:findProjectIdForObjectResponse"
name="findProjectIdForObjectResponse"/>

  </wsdl:operation>

  <wsdl:operation name="getObjectTitle" parameterOrder="oneTimeToken
objectId">

    <wsdl:input message="impl:getObjectTitleRequest"
name="getObjectTitleRequest"/>

    <wsdl:output message="impl:getObjectTitleResponse"
name="getObjectTitleResponse"/>

  </wsdl:operation>

  <wsdl:operation name="createTemplate" parameterOrder="oneTimeToken
sourceProjectId templateId templateName creatorId">

    <wsdl:input message="impl:createTemplateRequest"
name="createTemplateRequest"/>

    <wsdl:output message="impl:createTemplateResponse"
name="createTemplateResponse"/>

  </wsdl:operation>

  <wsdl:operation name="getTemplateMetadata"
parameterOrder="oneTimeToken locale">

    <wsdl:input message="impl:getTemplateMetadataRequest"
name="getTemplateMetadataRequest"/>

    <wsdl:output message="impl:getTemplateMetadataResponse"
name="getTemplateMetadataResponse"/>

  </wsdl:operation>

  <wsdl:operation name="getTemplateContent" parameterOrder="oneTimeToken
templateName locale">

    <wsdl:input message="impl:getTemplateContentRequest"
name="getTemplateContentRequest"/>

    <wsdl:output message="impl:getTemplateContentResponse"
name="getTemplateContentResponse"/>

  </wsdl:operation>

  <wsdl:operation name="validateParametersForTemplatizedProject"
parameterOrder="oneTimeToken iaConfig projectConfigurationParameters">

```



```

    <wsdl:input
message="impl:validateParametersForTemplatizedProjectRequest"
name="validateParametersForTemplatizedProjectRequest"/>

    <wsdl:output
message="impl:validateParametersForTemplatizedProjectResponse"
name="validateParametersForTemplatizedProjectResponse"/>

    </wsdl:operation>

    <wsdl:operation name="createTemplatizedProjectConfig"
parameterOrder="oneTimeToken iaConfig templateId title">

        <wsdl:input message="impl:createTemplatizedProjectConfigRequest"
name="createTemplatizedProjectConfigRequest"/>

        <wsdl:output message="impl:createTemplatizedProjectConfigResponse"
name="createTemplatizedProjectConfigResponse"/>

    </wsdl:operation>

    <wsdl:operation name="getSearchResults" parameterOrder="oneTimeToken
searchInfo locale">

        <wsdl:input message="impl:getSearchResultsRequest"
name="getSearchResultsRequest"/>

        <wsdl:output message="impl:getSearchResultsResponse"
name="getSearchResultsResponse"/>

    </wsdl:operation>

    <wsdl:operation name="getPageComponentParameters"
parameterOrder="oneTimeToken projectId locale">

        <wsdl:input message="impl:getPageComponentParametersRequest"
name="getPageComponentParametersRequest"/>

        <wsdl:output message="impl:getPageComponentParametersResponse"
name="getPageComponentParametersResponse"/>

    </wsdl:operation>

    <wsdl:operation name="getPageComponentContent"
parameterOrder="oneTimeToken projectId projectIntegratedApplicationId
pageComponentParam rowsPerPage locale">

        <wsdl:input message="impl:getPageComponentContentRequest"
name="getPageComponentContentRequest"/>

```

### pebble-dep.xml

The `pebble-dep.xml` file, also known as the Pebble deployment configuration file, contains the data that Pebble needs to interact with the TeamForge site.

#### A sample `pebble-dep.xml` file for the REST service type

```

<?xml version="1.0" encoding="utf-8"?>
<!DOCTYPE integrated-application
PUBLIC "-//CollabNet, Inc.//DTD Integrated Application Descriptor 1.0//EN"
"http://schema.open.collab.net/sfee50/dtd/sf-pluggable-deploy-
descriptor_1_0.dtd">

```

```
<integrated-application>
  <name>Pebble Blog</name>
  <baseurl>https://cu064.cloud.sp.collab.net:13001/pebble/index.jsp</
baseurl>
  <gourl>https://cu064.cloud.sp.collab.net:13001/pebble/gourl/%p/%o</gourl>
  <endpoint>https://cu064.cloud.sp.collab.net:13001/pebble/services/rest/
ctfapi</endpoint>
  <servicetype>REST</servicetype>
</integrated-application>
```

### A sample pebble-dep.xml file for the SOAP service type

```
<?xml version="1.0" encoding="utf-8"?>
<!DOCTYPE integrated-application
  PUBLIC "-//CollabNet, Inc.//DTD Integrated Application Descriptor 1.0//EN"
  "http://schema.open.collab.net/sfee50/dtd/sf-pluggable-deploy-
descriptor_1_0.dtd">
<integrated-application>
  <name>Pebble Blog</name>
  <baseurl>https://cul77.cloud.maa.collab.net:13001/pebble/index.jsp</
baseurl>
  <gourl>https://cul77.cloud.maa.collab.net:13001/pebble/gourl/%p/%o</gourl>
<endpoint>https://cul77.cloud.maa.collab.net:13001/pebble/services/
PebbleIntegrationService</endpoint>
  <servicetype>SOAP</servicetype>
</integrated-application>
```

### pebble-app.xml

The `pebble-app.xml` file, also known as the Pebble application configuration file, contains the text that the Pebble application displays in the TeamForge user interface.

This is an example of a default (unedited) `pebble-app.xml` file. To create your own integrated application config file, copy this one into a new file and replace the values with the values appropriate for the application you are integrating.

```
<?xml version="1.0" encoding="utf-8"?>
<!DOCTYPE integrated-application
  PUBLIC "-//CollabNet, Inc.//DTD Integrated Application Descriptor 1.0//EN"
  "http://schema.open.collab.net/sfee50/dtd/sf-pluggable-application-
descriptor_1_0.dtd">
<integrated-application>
  <name>Pebble Blog</name>
  <description>l10n.application.description</description>
  <permissions>
    <permission dapMappedTo="View">Blog Reader</permission>
    <permission>Blog Contributor</permission>
    <permission>Blog Publisher</permission>
    <permission>Blog Owner</permission>
  </permissions>
  <prefix>PB</prefix>
  <id-pattern></id-pattern>
  <require-per-project-prefix>true</require-per-project-prefix>
  <require-scm-integration>true</require-scm-integration>
  <is-search-supported>>false</is-search-supported>
  <!-- Page components for Integrated apps is not implemented for Alpha -->
  <page-component>
    <require-page-component>true</require-page-component>
    <page-component-details>
      <inputtype>text</inputtype>
      <resultformat>html</resultformat>
      <description>l10n.pce.description</description>
      <title>l10n.pce.title</title>
    </page-component-details>
```

```

</page-component>
<config-parameters>
  <!-- Pebble Configuration Parameters -->
  <param>
    <title>l10n.blogname.title</title>
    <name>blogName</name>
    <description>l10n.blogname.description</description>
    <defaultvalue>My Blog</defaultvalue>
    <displaytype valuetype="String" maxlength="25">TEXT</displaytype>
    <editable>>false</editable>
  </param>
  <param>
    <title>l10n.blogdescription.title</title>
    <name>blogDescription</name>
    <description>l10n.blogdescription.description</description>
    <defaultvalue>My Awesome Blog</defaultvalue>
    <displaytype valuetype="String" maxlength="40">TEXT</displaytype>
    <editable>>true</editable>
  </param>
  <param>
    <title>l10n.richtexteditor.title</title>
    <name>richTextEditorEnabled</name>
    <description>l10n.richtexteditor.description</description>
    <defaultvalue>checked</defaultvalue>
    <displaytype valuetype="String">CHECKBOX</displaytype>
    <editable>>true</editable>
  </param>
  <param>
    <title>l10n.noofrecentblogentries.title</title>
    <name>recentBlogEntries</name>
    <description>l10n.noofrecentblogentries.description</description>
    <defaultvalue>3</defaultvalue>
    <displaytype valuetype="String">SELECT</displaytype>
    <option name="3">l10n.three.value</option>
    <option name="5">l10n.five.value</option>
    <option name="7">l10n.seven.value</option>
    <option name="9">l10n.nine.value</option>
    <editable>>true</editable>
  </param>
</config-parameters>
<bundles>
  <bundle locale="en">
    <key name="l10n.application.description">Pebble Blog App</key>
    <key name="l10n.pce.description">Display Blog Title for Given
Date.</key>
    <key name="l10n.pce.title">Enter Blog Date (in yyyy-mm-dd)</key>
    <key name="l10n.blogname.title">Blog Name</key>
    <key name="l10n.blogname.description">Please provide a name for
the Blog. This appears on all blog pages</key>
    <key name="l10n.blogdescription.title">Blog Description</key>
    <key name="l10n.blogdescription.description">Please provide a
description for the Blog. This appears below blog name on all pages</key>
    <key name="l10n.richtexteditor.title">Rich Text Editor</key>
    <key name="l10n.richtexteditor.description">Enable Rich Text
Editor for comments and Blog entries</key>
    <key name="l10n.noofrecentblogentries.title">Recent Blog Entries</
key>
    <key name="l10n.noofrecentblogentries.description">How many recent
blog entries do you want to see in the home page?</key>
    <key name="l10n.three.value">3</key>
    <key name="l10n.five.value">5</key>
    <key name="l10n.seven.value">7</key>
    <key name="l10n.nine.value">9</key>
  </bundle>

```

```
</bundles>  
</integrated-application>
```

## Extend TeamForge

---

TeamForge provides you with the interrelated extension features to suit your organization's specific needs.

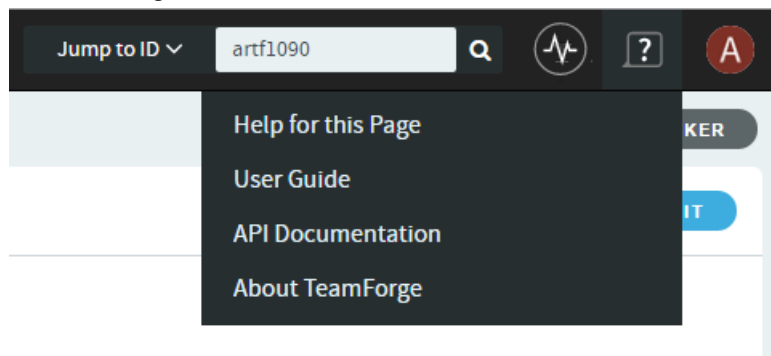
### Building on the TeamForge platform

---

TeamForge is an open platform that you can build on to provide functionality beyond what's in the basic application using both SOAP and REST APIs.

Use some or all of the interrelated extension features to tailor TeamForge to your organization's unique needs. TeamForge supports both SOAP and REST APIs.

See TeamForge's API Documentation online.



### Event handlers: Create custom work flows in TeamForge

When you add an event handler to your TeamForge site, you can automatically react to system events in ways that help your site's project members or administrators.

#### How TeamForge event handlers work

The TeamForge custom event handling framework allows third-party event handlers to register for TeamForge-specific application events and notifies them whenever such an event occurs.

The event handling framework implements an extended flavor of the observer pattern. The TeamForge application events are triggered whenever a property of a TeamForge object (e.g. tracker item, discussion item, wiki page) has been changed or is going to be changed, if no event handler objects (i.e. blocks the event).

For example, you can block deletions of projects for all users, add a comment to a tracker item whenever an association has been modified, or design your own tracker workflow engine.

Writing custom event handlers requires at least some basic knowledge in the Java programming language or (if you use the examples shipped with this post) a script language that is installed on the TeamForge server, such as shell, Python, or Perl. For these instructions, we'll assume that you are familiar with basic programming techniques.

The event handler framework differentiates between two types of events:

- **Asynchronous:** If a handler registers for asynchronous events, it is informed that a change has just happened. The handler can decide to trigger further changes by calling TeamForge web services, but it cannot block the change because it has already happened.

Asynchronous event handlers are good for triggering system events, such as changing an artifact status or sending an email. See [Using an asynchronous event handler: trigger follow-up events](#) on page 827.

- **Synchronous:** If a handler has registered for synchronous events, it gets informed whenever a change has been anticipated by a user. It can examine the properties that should be changed and decide whether to accept the change or block it. A synchronous event handler cannot trigger further changes on the currently processed object, since other handlers in the event handler chain must also have the chance to block the anticipated change.

A synchronous event handler is the appropriate way to show an alert directly in the TeamForge UI, for example. See [Using a synchronous event handler: Send event handler output to the TeamForge UI](#) on page 826.

Technically, all event handlers have to be part of a Java archive (JAR) file with a TeamForge specific deployment descriptor that describes which events should be intercepted. This JAR file then has to be uploaded to the TeamForge application server. No restart is necessary, but the event handling cache has to be refreshed.

In practice, you can customize a TeamForge site's behavior without any knowledge of Java if you can write scripts in a language that can deal with environment variables, write to standard out/error (to control what will be displayed in the TeamForge UI as result of the handler's execution) and control the return code (to decide whether to block the event or not).


We will show you how to come up with your own custom event handlers based on two examples.

### Event handler example: Comment on associations

This event handler adds a comment to a tracker item whenever an association is added to or deleted from this tracker item.

This example illustrates how to intercept a specific event, trigger a follow-up action by calling TeamForge web services, and add a comment based on the formatting template which is specified as part of a property file.

The code for example one can be found [here](#).

 **Tip:** You can treat a JAR file like a zip file: just extract it with your favorite zip program and have a look at the files that are part of it.

1. When you extract the JAR file you'll find a structure like this:


```
com/vasoftware/sf/plugin
```

This directory contains the class file of your event handler. There may be additional directories containing Java class files. If you like to include Java libraries, you have to unpack their JAR files and add their class files (including directory structure) in the event handler JAR.

```
META-INF/event.xml
```

This file contains the events and operations your handler class will intercept. It is also common to have additional files in the META-INF directory, such as property files to control the behavior of the event handler.

This example event handler is provided as-is (i.e. not supported as part of any TeamForge release). As with all event handlers, use it at your own risk. CollabNet cannot guarantee any SLAs on third-party code.

 **Tip:** If you want to change the behavior of custom event handlers at runtime (without redeploying the JAR file), you may want to look at the [TeamForge integration data API](#). For the purpose of this example, we will stick with property files.

2. In the `event.xml` file, you will find some lines like this:

```
<event api="6.1" mode="asynchronous">
  <type>relationship</type>
  <operation>*</operation>

  <handler>com.vasoftware.sf.plugin.AsynchronousRelationshipEventListener</
handler>
</event>
```

This tells the event handling framework that the class `AsynchronousRelationshipEventListener` is responsible for intercepting events of type “relationship” (aka associations) for every possible operation. The handler will be called after the event has happened (asynchronous mode) and the passed data structures will be compatible with the events format defined in TeamForge 6.0 (in other words, the event handler extends the `EventHandler60` class). If you only want to intercept certain operations, you can specify those instead of the wildcard character (\*). Supported operations are usually create, update, move and delete, but every event can define its own operations.


3. The `config.properties` file is used to control the formatting of the comment that gets added when an association has been modified.

The initialize-method of the handler class (`AsynchronousRelationshipEventListener`) shows how property files can be parsed within a custom event handler.

4. By default, the user triggering the event is also the user executing the event handler. If you want to run your event handler with a different user account, specify it in the user element, like this:

```
<event api="6.1" mode="synchronous">
  <type>*</type>
  <operation>*</operation>

  <handler>com.vasoftware.sf.plugin.AsynchronousRelationshipEventListener</
handler>
  <user>foo</user>
</event>
```

 **Note:** Be careful with this option, because running code on behalf of a different user opens the door for all kind of exploits if you do not check the user's input properly. Also, this option will not allow you to access the original user's session id any more, but you can always create a new session with a super user account (credentials saved in a property file) if you have to.


### Event handler example: Execute a hook script

When a TeamForge event arrives, this event handler looks to see whether there is a script in the TeamForge file system with the name/operation of the event, and then calls that script with all information from the event contained within environment variables.

This example illustrates how to intercept arbitrary TeamForge events, examine the event's properties, map them to system environment variables and call a script in the file system with a name corresponding to the intercepted event.

You can use this event handler to customize your TeamForge site's behavior without any knowledge of the Java programming language as long as you can write scripts in a language that can deal with environment variables, write to standard out/error (to influence what will be displayed in TeamForge's UI as result of the handler's execution) and influence the return code (to decide whether to block the event or not).

The code for this example can be found [here](#).

 **Tip:** You can treat a JAR file like a zip file: just extract it with your favorite zip program and have a look at the files that are part of it.

1. When you extract the JAR file you'll find a structure like this:


```
com/collabnet/ctf/events
```

This directory contains the class file of your event handler. There may be additional directories containing Java class files. If you like to include Java libraries, you have to unpack their JAR files and add their class files (including directory structure) in the event handler JAR.

```
META-INF/event.xml
```

This file contains the events and operations your handler class will intercept. It is also common to have additional files in the META-INF directory, such as property files to control the behavior of the event handler.

This example event handler is provided as-is (i.e. not supported as part of any TeamForge release). As with all event handlers, use it at your own risk. CollabNet cannot guarantee any SLAs on third-party code.

 **Tip:** If you want to change the behavior of custom event handlers at runtime (without redeploying the JAR file), you may want to look at the [TeamForge integration data API](#). For the purpose of this example, we will stick with property files.

2. In the `event.xml` file, you will find some lines like this:

```
<event api="6.1" mode="asynchronous">
```

```

<type>*</type>
<operation>*</operation>
<handler>com.collabnet.ctf.events.AsynchronousHookScriptEventListener</
handler>
</event>

<event api="6.1" mode="synchronous">
<type>*</type>
<operation>*</operation>
<handler>com.collabnet.ctf.events.SynchronousHookScriptEventListener</
handler>
</event>

```

These lines tell TeamForge to register two event handlers, one asynchronous (`AsynchronousHookScriptEventListener`) and one synchronous (`SynchronousHookScriptEventListener`) for arbitrary events (wildcard `*`). It also tells TeamForge that the handler classes will extend the `EventHandler60` interface.

3. By default, the user triggering the event is also the user executing the event handler. If you want to run your event handler with a different user account, specify it in the user element, like this:

```

<event api="6.1" mode="synchronous">
  <type>*</type>
  <operation>*</operation>
  <handler>com.collabnet.ctf.events.SynchronousHookScriptEventListener</
handler>
  <user>foo</user>
</event>

```



**Note:** Be careful with this option, because running code on behalf of a different user opens the door for all kind of exploits if you do not check the user's input properly. Also, this option will not allow you to access the original user's session id any more, but you can always create a new session with a super user account (credentials saved in a property file) if you have to.

It is possible to register multiple handlers for different events, but you can also use one handler to intercept both synchronous and asynchronous events.

### Event handler example: Hook scripts

These sample hook scripts should give you an idea how custom event handlers can be written. Feel free to adjust them to your own needs.



**Note:** These sample event handlers are not officially supported by CollabNet and must be used at your own risk.

- Hooks must be owned by `sf-admin` for security, and must have the executable bit set.
- To configure a site to prevent projects being deleted, we could create this file: `/opt/collabnet/teamforge/hooks/synchronous/project_delete`

```

#!/bin/sh
echo Sorry, projects cannot be deleted on this site 1>&2
exit 1

```

- To automatically create an initial directory structure in an SVN repository when the repository is created, you might create this file: `/opt/collabnet/teamforge/hooks/asynchronous/repository_create`

```

#!/bin/sh
/usr/bin/svn
mkdirhttp://localhost/svn/repos/${tf_original_RepositoryDirectory:9:999}/
trunk
http://localhost/svn/repos/${tf_original_RepositoryDirectory:9:999}/tags

```



```

http://localhost/svn/repos/${tf_original_RepositoryDirectory:9:999}/
branches
-m "Initial Structure"
--username admin --password mypassword --non-interactive --no-auth-cache
exit 0

```

### Extend `EventHandler60` for your TeamForge event handler classes

All event handlers (synchronous and asynchronous) have to extend

`com.vasoftware.sf.events.EventHandler60`.

This is how our example "Comment on associations" event handler does it:

```

package com.vasoftware.sf.plugin;
...
public class AsynchronousRelationshipEventListener extends EventHandler60 {
...
}

```

The following table lists the JAR file location for different versions of TeamForge that contains `EventHandler60`.

TeamForge Version	JAR File
6.1	/opt/collabnet/teamforge/dist/deployments/saturn.ear/sfevents.jar
6.1.1	/opt/collabnet/teamforge/dist/deployments/saturn.ear/lib/sfevents.jar
6.2	/opt/collabnet/teamforge/dist/deployments/saturn.ear/lib/sfevents.jar
6.2.0.1	/opt/collabnet/teamforge/dist/deployments/saturn.ear/lib/sfevents.jar
7.0	/opt/collabnet/teamforge/dist/deployments/saturn.ear/lib/sfevents.jar
7.1	/opt/collabnet/teamforge/dist/deployments/saturn.ear/lib/sfevents.jar
8.0	/opt/collabnet/teamforge/dist/deployments/saturn.ear/lib/events-api.jar
8.0.0.1	/opt/collabnet/teamforge/dist/deployments/saturn.ear/lib/events-api.jar
8.1	/opt/collabnet/teamforge/dist/deployments/saturn.ear/lib/events-api.jar
8.1.0.1	/opt/collabnet/teamforge/dist/deployments/saturn.ear/lib/events-api.jar
16.3	/opt/collabnet/teamforge/dist/deployments/lib/events-sdk.jar

TeamForge Version	JAR File
16.7 and later	/opt/collabnet/teamforge/dist/deployments/lib/event.jar

- EventHandler60 has one abstract method you have to override:

```
@Override
public void processEvent() throws Exception {
    ...
}
```

- This method does not take any parameters, so how do we access the relevant information?

There are some interesting methods of the base class (EventHandler60) you can call:

`getSessionKey`: Returns a session id of the user who is going to (synchronous handler) / has triggered (asynchronous handler) the event we just intercepted. If you used the user element in event.xml, it will contain a session id for the user you specified there.

`getEventContext`: Returns a data structure containing the event type, operation, project, comment and user name

`getOriginalData`: In case of a synchronous event handler, this will return a representation of the object the event is going to change. In case of an asynchronous event handler, this will return the representation of the object before it was changed by the event. The data structure used to represent the object is the same that would have been used in CollabNet's SOAP API.

`getUpdatedData`: In case of a synchronous event handler, this will return a representation of the object how it will look after the event has happened (you can still block it). In case of an asynchronous event handler, this will return the representation of the object after it was changed by the event.

Let's assume a user wants to change the priority of a tracker item from 3 to 4. If you have registered a synchronous event handler, this one is triggered before the change can actually be performed.

`getOriginalData` returns an `ArtifactSoapDO` object of the tracker item with the priority field set to 3. `getUpdatedData` contains an `ArtifactSoapDO` object of the tracker item with the priority field set to 4.

- If you block the event (by throwing an exception), the change does not happen and the user is presented with an error message. (See next section for how to influence this error message.)
- If you do not block the event (by just returning from the `processEvent` method), all registered asynchronous handlers are called. `getOriginalData` and `getUpdatedData` contain exactly the same objects as in the synchronous case. However, the semantic is different: They are no longer representing the current and anticipated state, but the previous and current state of the object in question.
- The following code snippet (taken from our "Hook script" event handler example) shows how to retrieve all information available to an event handler.

```
String type = getEventContext().getType();
String operation = getEventContext().getOperation();
String projectId = getEventContext().getProjectId();
String comment = getEventContext().getComment();
String userName = getEventContext().getUsername();
String originalDataClassName = getOriginalData().getClass().toString();
String updatedDataClassName = getUpdatedData().getClass().toString();
Object originalData = getOriginalData();
Object updatedData = getUpdatedData();
```

### Using a synchronous event handler: Send event handler output to the TeamForge UI

When we have extracted all data available to the event handler, how do we interact with the user interface?

- 👉 **Note:** Only synchronous event handlers can directly communicate with the UI, because if the event has already happened (as it has, in the case of asynchronous handlers), the user who triggered the event may already have been logged out.

- You use three independent actions to interact with the TeamForge UI:
  - a) Add a success message to the UI that gets displayed as the result of the action just triggered by the user.  
This can be done by calling the `addSuccessMessage` method of the `EventHandler` base class (see `SynchronousHookScriptEventListener.java` of example two for details).
  - b) Add an error message to the UI that gets displayed as the result of the action just triggered by the user.  
This can be done by calling the `addErrorMessage` method of the `EventHandler` base class.
  - c) Block the event you intercepted.  
This can be done by throwing an exception in the `processEvent` method. The payload of your exception will be displayed in the UI.

All three forms of UI feedback can be used in combination. For example, it is possible to display an error message even if you did not block the event, and it is possible to show many error and success messages together.

- What happens if the event in question was not triggered by a user logged into the TeamForge Web UI but by a client using the TeamForge web services?  
In this case, error and success messages do not reach the SOAP client. However, the payload of the exception object thrown when the event was blocked is delivered as part of the SOAP fault element.
- While synchronous event handlers enable you to block events and/or to provide additional feedback to the currently logged in user, they should not be used to trigger follow-up actions (like changing TeamForge artifacts or interacting with external systems).

Remember that these handlers are running in the main TeamForge event loop and nothing else will happen until you return from the `processEvent` method, so return as fast as you can.

### Using an asynchronous event handler: trigger follow-up events

Use an asynchronous event handler to communicate with TeamForge, external systems, processes or system resources.

- 👉 **Important:** To avoid accidentally locking the main TeamForge event queue down (and essentially rendering the system unusable), use only asynchronous event handlers (not synchronous event handlers) to trigger events.

Interacting with TeamForge is done as you would do it if you had to write a Java program to interact with TeamForge using its web services API. The only difference is that you will connect to `localhost` (since your handler is running locally) and that you already have a valid session ID.

- You do not have to include the SOAP SDK classes in your event JAR file, because they are already in the TeamForge class path.

This code snippet extracted from our association converter example ([Event handler example: Comment on associations](#) on page 822) shows how to do it:

```
ITrackerAppSoap trackerClient = (ITrackerAppSoap)
    ClientSoapStubFactory.getSoapStub(
    ITrackerAppSoap.class, "http://localhost:8080");
...
ArtifactSoapDO artifact = trackerClient.getArtifactData(getSessionKey(),
    originId);
trackerClient.setArtifactData(getSessionKey(), artifact, finalComment,
    null, null, null);
```

- 👉 **Tip:** You may use external libraries in your event handler by placing their `.class` files into your event JAR file. The only tricky part is if TeamForge is using a different version of this library (which will take precedence). In this case, you would have to recompile your library with a different package namespace.
- Using the session key provided by the event handler is actually only going to work if the SOAP call you are using is not throwing an exception.

The session ID passed into your handler is associated with an already running transaction that will be aborted if an exception is thrown as part of this session. Part of rolling back the transaction is rolling back the JVM's call stack which contains your event handling code, so you will not be able to catch the web services exception. If you like

to to handle web service exceptions, you have to create your own session id by logging into TeamForge again by calling `ICollabNetSoap.login` with some credentials stored as part of your handler. (You can store them in a property file in the `META-INF` directory.)

### Best practices for working with custom TeamForge event handlers

In general, watch out for deadlocks and favor asynchronous over synchronous event handlers.

#### Beware of deadlocks.

Having custom event handlers that modify other objects can be dangerous if there it is possible for that handler or another handler to chain in the opposite direction. An example of this is an event handler that updates a task when an associated artifact is updated and updates the artifact when the associated task is updated. It is possible for two users to modify each object at the same time causing the two event handlers to wait on each other. The task handler would have a lock on the task bean in the application server while the artifact handler would have a lock on the artifact bean. When the custom event handlers fired, they would wait for the locks to be released but since the two threads have the locks each other needs and are waiting on the opposite objects, a deadlock would occur.

#### Asynchronous is safer

Custom event handlers will be the least worrisome when they are responsible for data validation or secondary object creation (or association creation). Object modification is possible but adds greater complexity due to the risks involved with locking multiple objects across many threads. If you are unsure, use asynchronous handlers to modify objects instead since the lock on the original object will be gone by the time the asynchronous handler is executed.

Calling and waiting for synchronous hooks currently doesnt have a timeout. As long as your synchronous hook is running, the whole TeamForge site will be blocked for all users accessing the site. Some events trigger other events. For example creating a project actually calls the create project hook, wiki page hooks, and so on. Badly written or slow hooks can cripple a site.

#### Write to a file

Write your diagnostics messages in a file and not on `stdout/stderr`, since TeamForge does not read from `stdout/stderr` before the script completes. In the case of synchronous hooks, this could lead to a situation where the script blocks because the pipes buffer between the script process and the TF process is completely filled.

#### No cascading

Due to the nature of custom event handling, custom events cannot cascade. This means that if a custom event handler catches an event and creates an object that it or another custom event handler would normally process, the event bypasses the custom event handlers. This is to prevent looping and infinite object creation. While there are ways for event handlers to avoid this, it would be a fairly difficult task since all of your event handlers would have to use a circular event detection algorithm. Rather than adding that complexity, we just eliminated the possibility.

#### Event handler life cycle

For every single call to the `processEvent` method, a new object of your class will be instantiated. A best practice to avoid costly reinitialization every time (remember that the TeamForge event loop thread is blocked while you are doing this) is to delegate all synchronization work to a method you always call in your constructor which checks a static variable whether the initialization has already been done and if not, just returns without any further action (code snippet from example one):

```
private static boolean initialized = false;
public AsynchronousRelationshipEventListener() {
    initialize();
}

private synchronized void initialize() {
    if (initialized) {
        return;
    }
}
```

```

initialized = true;
// proceed with initialization
...
}

```

Logging in into TeamForge and initializing network connections file resources are costly operations that should be handled in such a method instead of doing it all over again.

### Event spooling

While it is true that asynchronous handlers may consume considerably more time than synchronous ones, there is only one thread for those handlers, so events may queue up if you do expensive operations. A best practice is to capture the event in your asynchronous event handler, write all necessary information to the local file system (comparable to a mail spooling directory) and return. At the same time, you can have a separate application reading from the spooling directory. This way, you never get into a situation where you miss TeamForge events, or things queue up just because you run into a blocking operation.

### Incremental changes

The event handler parser is really picky on the exact format of your JAR file. A best practice is to base your work on an already existing event handler and then adapt it to your own needs by doing incremental changes while checking whether it still works.

### Watch out for loops

Your follow-up actions may trigger your handler to be called again. You have to protect your handler from an infinite update loop if that happens. A best practice is to add a check to your event handler to see whether the user initiating the event is the same user you are using to perform follow-up actions.

### Roll back sparingly

Throwing an exception in a synchronous event handler blocks the intercepted event and rolls back the transaction associated with the change. Rolling back transactions also means that the data the user entered is not saved. If this happens accidentally due to a wrongly programmed event handler, it can be frustrating to your users, so make sure that you only throw exceptions in your handler code when you really want to enforce the rollback.

### Catch errors generically

It is quite easy to miss an exception you did not expect (like a null pointer exception, parsing exception, time out exception, any other malfunction in your own code). A best practice is to introduce a generic catch block in your handler and only rethrow the exception if it was an intended exception (see `SynchronousHookScriptEventListener`):

```

} catch (Exception e) {
    if (!intendedException) {
        log.error("Exception occurred: " + e.getMessage(), e);
    } else {
        ...
        throw e;
    }
}

```

### Reference information about event handlers

Here is some stuff you may need to know to work with event handlers.

#### event.xml

Here is an example of an `event.xml` file. Use it to create a new `event.xml` file for your custom event handler in the `META-INF` directory of your `jar` file.

In this example, the `CreateTestCaseWorkflow` class will be invoked only on artifact creation and will run after the artifact creation is successfully committed. The `EventTestListener` will run for every single event and will run synchronously thereby having the opportunity to cancel the event.

```

<?xml version="1.0"?>
<!DOCTYPE event-handler
PUBLIC "-//VA Software, Inc.//DTD Data Object 1.1//EN"
"http://schema.vasoftware.com/sf/dtd/sf-event-handler_1_0.dtd">

<event-handler>
  <event api="6.1" mode="asynchronous">
    <type>artifact</type>
    <operation>create</operation>
    <user>admin</user>
    <handler>com.collabnet.ctf.events.CreateTestCaseWorkflow</
handler>
  </event>

  <event api="6.1" mode="synchronous">
    <type>*</type>
    <operation>*</operation>
    <handler>com.collabnet.ctf.events.EventTestListener</handler>
  </event>
</event-handler>

```

### DTD for custom event handler

An event handler works by subscribing to system events and then responding when event occurs. You subscribe to events by placing an XML file in the custom event handler JAR that describes what events are being monitored and who is doing the monitoring.

Here is the current DTD.

```

<!ELEMENT event-handler (description?, event+)>

<!-- Optional description for the Event Handler -->
<!ELEMENT description (#PCDATA)>

<!ELEMENT event (type, operation, user?, handler)>

<!-- This is the API level that the event handler expects the
ObjectSoapDO objects to be marshaled for -->
<!ATTLIST event api (6.1|6.2) #REQUIRED>

<!--
Whether the event handler should run synchronously (allows it
to cancel the event) or asynchronously after the event
has completed and is committed to the database.
-->
<!ATTLIST event mode (synchronous|asynchronous) #REQUIRED>

<!-- The object type being handled (e.g., artifact, task,
document). '*' is all object types. -->
<!ELEMENT type (#PCDATA)>

<!-- This is the operation that is being listened for (e.g.,
create. delete. move, update). '*' is all operations. -->
<!ELEMENT operation (#PCDATA)>

<!--
If the event handler needs to run as a different user (ie., a
site admin account) the username should be

```

```

    passed in this attribute. If the username is not valid the
    event handler will throw an abort exception. That
    means that if the handler is processing synchronously the
    event will be canceled and the user will get an
    exception message.
-->
<!ELEMENT user (#PCDATA)>

<!--
    Name of a fully qualified class to handle the event.
-->
<!ELEMENT handler (#PCDATA)>

```

### ENABLE\_UI\_FOR\_CUSTOM\_EVENT\_HANDLERS

To support branding and customization changes, set the *ENABLE\_UI\_FOR\_CUSTOM\_EVENT\_HANDLERS* variable to true.

#### Values

true or false

#### Default

true

## API guide: Use TeamForge features via SOAP

When you set up another application to communicate with the CollabNet TeamForge API, the remote app's users can do almost anything they can do in the CollabNet product's Web interface.


The CollabNet TeamForge web services offer external access to core application functionality and data, such as work tracking, task management, and hardware provisioning.

See [CollabXchange](#) for examples of successful integrations using CollabNet APIs.

### Get started with the CollabNet TeamForge SOAP API

A CollabNet TeamForge plugin can provide any of the features a user can access through the Web interface.

1. Browse the [CollabNet user help](#) to get a full picture of the functionality you can deliver with the SOAP API.
2. Get an idea of the sorts of applications that have been developed.
  - Look at the applications available for download on [CollabXchange](#).
  - Review the reference applications on the [CollabNet community site](#). These samples illustrate some of the more common methods for using TeamForge via the SOAP API. They interact with most of the TeamForge data objects, including users, projects, trackers and artifacts, code commits and tasks.

 **Note:** If you find a defect in a sample program or have a suggestion for improvement, please post a message in the [developer discussion forum](#).

3. Set up the tools you will need:
  - A Java JDK, 1.8.0\_112 or later, from <http://www.oracle.com/technetwork/java/index.html>.
  - Apache Ant, version 1.7.0 or later, from <http://ant.apache.org>.
4. Choose a place to host your plugin project online.

We recommend <http://www.collab.net/community>, because you get access to all the collaboration support tools that CollabNet TeamForge provides.
5. Get the SDK from <http://www.collab.net/community/teamforge>.

The SDK includes everything needed to develop and deploy your application:

  - Source files and compiled output.

- Full JavaDoc reference material. Inside the package, look for `doc/com/collabnet/ce/soap60/webservices`.
- Annotated code examples.


6. Get help from other CollabNet TeamForge users and CollabNet staff in the TeamForge [API discussion forum](#).

### Example: Send commit data to TeamForge via SOAP

Use the Commit Tool provided with TeamForge to transmit your commit data to TeamForge.

The SCM Adapter allows you to develop integrations with almost any SCM tool, then exchange commit data with TeamForge. After you have created an SCM integration using the TeamForge API, you will use the Commit Tool provided with TeamForge to transmit your commit data to TeamForge.

For any SCM tool that supports triggers, you can use the tool's triggering mechanism to script the following actions. Otherwise, you can perform them manually for each commit.

 **Note:** The SCM Adapter option is always available, but will work only if you have developed an integration with an SCM tool using the TeamForge API.

1. Get the Commit Tool from `$SOURCEFORGE_HOME/integration/CommitTool.py`
2. Add the SCM server to your TeamForge installation.  
For the end user instructions, see [Integrate a source code server](#) on page 347 .
3. Create a repository on the SCM server.  
For the end user instructions, see [Create a source code repository](#) on page 93.
4. Use the Commit Tool to start a commit.

```
CommitTool.py create <systemId> <path> <username> <host> <port>
```

where:

- `systemId` is the external system identifier that was given to the integration server. You can find this on the **Repository Details** page.
- `path` is the path on the external system that was given to the repository. See the **Repository Details** page.
- `username` is the TeamForge user name that will be performing the commit.
- `host` is the TeamForge application server machine hostname.
- `port` is the TeamForge application server machine SOAP port.

5. After the commit has been created, use the Commit Tool to add files with versions and actions.

```
CommitTool.py add <filename <version [<status [<fromfile <fromversion]]
```


where:

- `filename` is the name of the file that is being placed into the commit.
- `version` is the version of the file that is being placed into the commit.
- `status` is the status of the file in the repository. Valid status values are 'added', 'deleted', 'modified', 'moved', and 'copied'. Status is optional and will default to 'added' if no value is provided.
- `fromfile` is the name of the original filename if this file was copied or moved. It is required on a copy or move, but is not allowed on an add, delete, or modify.
- `fromversion` is the version of the original filename is this file was copied or moved. It is required on a copy or move, but is not allowed on an add, delete, or modify.

6. When all actions and files have been added to the commit, use the Commit Tool to transmit the commit to CollabNet TeamForge.


```
CommitTool.py commit <commitmessage>
```

where `<commitmessage>` is a description about the commit that will be displayed along with the files, versions, and operations.

 **Note:** To create an association, use the standard associate command [`<item id>`]. No special syntax is required when using the Commit Tool.



For all of the above commands, when run correctly, there will be no output, and the return value of the program will be “0”, indicating success. Failure will show a message and return with a non “0” return value.

 **Tip:** To see the status of the current commit, you can use the Commit Tool to print the output of the current commit:

```
CommitTool.py print
```

It will show output similar to this:

```
Repository: <username>@<systemid>:<path> Modified filename 1.2
Added document 1.1 Copied newfile 1.1 (From: origfile 1.3) <status>
<filename> <version> (From: <fromfile> <fromversion>)
```

### Update an application to the TeamForge 17.4 API

To use any of the new API calls introduced in TeamForge 17.4, you must update your existing applications to use the 17.4 API.

### What's new in TeamForge 17.4

Click the links in the following table and refer to the Change History section to know the new and changed methods in TeamForge 17.4.

TeamForge service	Changed in 17.4?	Link to change history
CollabNet service	No	Javadoc for the <a href="#">CollabNet</a> service
DiscussionApp service	No	Javadoc for the <a href="#">DiscussionApp</a> service
DocumentApp service	<b>Yes</b>	Javadoc for the <a href="#">DocumentApp</a> service
FileStorageApp and SimpleFileStorageApp services	No	Javadoc for the <a href="#">FileStorageApp</a> and <a href="#">SimpleFileStorageApp</a> services
FrsApp service	No	Javadoc for the <a href="#">FrsApp</a> service
IntegrationDataApp service	No	Javadoc for the <a href="#">IntegrationDataApp</a> service
MonitoringApp service	No	Javadoc for the <a href="#">MonitoringApp</a> service
NewsApp service	No	Javadoc for the <a href="#">NewsApp</a> service
PageApp service	No	Javadoc for the <a href="#">PageApp</a> service
PlanningApp service	No	Javadoc for the <a href="#">PlanningApp</a> service
PluggableApp service	No	Javadoc for the <a href="#">PluggableApp</a> service
RbacApp service	No	Javadoc for the <a href="#">RbacApp</a> service
ScmApp service	No	Javadoc for the <a href="#">ScmApp</a> service
TaskApp service	No	Javadoc for the <a href="#">TaskApp</a> service
TeamApp service	No	Javadoc for the <a href="#">TeamApp</a> service
TrackerApp service	No	Javadoc for the <a href="#">TrackerApp</a> service
WikiApp service	No	Javadoc for the <a href="#">WikiApp</a> service

TeamForge service	Changed in 17.4?	Link to change history
CategorizationApp service	No	Javadoc for the <a href="#">CategorizationApp</a> service

1. For calls that use the TeamForge-provided Java SDK classes, update all references to the package containing the classes.

**For this TeamForge version...**

**6.x/7.x/8.x/16.x**

**the package is located at...**

`com.collabnet.ce.soap60`

2. For calls that access the WSDL directly, point your application to the new WSDL location.

**For this TeamForge version...**

**6.x/7.x/8.x/16.x**

**the WSDL is located at...**

`http://<teamforge#server>/ce#soap60/services/<service#name>?wsdl`


3. Update any calls that have been changed in the 17.4 API.

### TeamForge SOAP API reference

TeamForge provides a SOAP service for each application.

#### TeamForge services available via SOAP

TeamForge provides a SOAP service for each tool in the application.

 **Note:** The WSDL for each TeamForge service is available on your TeamForge site at `https://<mysite.com>/ce-soap60/services/<myService>?wsdl`. For an example, to review the WSDL for the Tracker service, see <http://ctf.open.collab.net/ce-soap60/services/TrackerApp?wsdl>. For information on the SOAP API, including the pluggable component package and classes, see the [SOAP package summary](#).

#### The CollabNet service

The CollabNet service is the primary service that handles logging into and out of CollabNet TeamForge, finding and creating users, and retrieving lists of projects, project members, and project data.

The CollabNet service must always be accessed first, to authenticate the user and return a sessionId. The sessionId is a necessary parameter that must be passed to the methods in the other services.

The CollabNet service also handles retrieving and managing lists of users and projects, which are often the first steps that an application must perform.

Some examples of activities that are managed by the CollabNet service are:

- Returning a list of all CollabNet TeamForge users or projects.
- Creating a new user.
- Providing a user with a list of all projects of which he or she is a member.
- Creating an association between two items.
- Managing user groups and their membership.

User groups, added in CollabNet TeamForge Enterprise Edition 4.4 Service Pack 1, simplify permission management for projects. User groups are site-wide sets of users, managed by site administrators. Project administrators can add user groups to roles, just like they can add individual project members.

After authenticating the user through the CollabNet service, you can then access all of the other services that manage the items and activities associated with the other CollabNet TeamForge applications, such as trackers, tasks, documents, or file releases. A new “Anonymous Login” capability introduced in CollabNet TeamForge Enterprise Edition 4.4 adds the ability to log in with the privileges and access rights of a non-authenticated user.

### Detailed information

For a complete description of each method, including its parameter definitions and SOAP faults, see the [JavaDoc for the cemain service](#).

## WSDL

The WSDL for this service is available on your TeamForge site at <https://<mysite.com>/ce-soap60/services/CollabNet?wsdl>. For example, here's the copy on the CollabNet web site: <http://ctf.open.collab.net/ce-soap60/services/CollabNet?wsdl>.

### New and changed methods

Refer to the Change History section on the [CollabNet service Javadoc](#) for a list of new and changed methods.

#### *The discussion service*

The DiscussionApp service handles the activities and items associated with forums.

Some examples of activities that are managed by the DiscussionApp service are:

- Creating a forum, forum topic, or forum post.
- Returning a list of all posts in a forum.
- Returning a list of forum posts created by a specified user.
- Returning a list of forum posts matching a search string.
- Deleting a forum, forum topic, or forum post.

After authenticating the user through the CollabNet service, you can then access the DiscussionApp service to work with forums, forum topics, and forum posts.

### Detailed information

For a complete description of each method, including its parameter definitions and SOAP faults, see the [JavaDoc for the DiscussionApp service](#).

## WSDL

The WSDL for this service is available on your TeamForge site at <https://<mysite.com>/ce-soap60/services/DiscussionApp?wsdl>. For example, here's the copy on the CollabNet web site: <http://ctf.open.collab.net/ce-soap60/services/DiscussionApp?wsdl>.

### New and changed methods

Refer to the Change History section on the [DiscussionApp service Javadoc](#) for a list of new and changed methods.

#### *The document service*

The DocumentApp service handles the activities and items associated with documents and document folders.

Some examples of activities that are managed by the DocumentApp service are:

- Creating a document or document folder.
- Editing a document's details.
- Returning the name of the user who last edited or locked a document.
- Returning a list of documents matching a search string.
- Returning a list of all document folders in a project.
- Returning a list of all open document reviews in a project.
- Moving a document or document folder.

After authenticating the user through the CollabNet service, you can access the DocumentApp service to work with documents and document folders.

### Detailed information

For a complete description of each method, including its parameter definitions and SOAP faults, see the [JavaDoc for the DocumentApp service](#).

## WSDL

The WSDL for this service is available on your TeamForge site at <https://<mysite.com>/ce-soap60/services/DocumentApp?wsdl>. For example, here's the copy on the CollabNet web site: <http://ctf.open.collab.net/ce-soap60/services/DocumentApp?wsdl>.

## New and changed methods

Refer to the Change History section on the [DocumentApp service Javadoc](#) for a list of new and changed methods.

### *The file release service*

The FrsApp service handles the activities and items associated with file releases.

Some examples of activities that are managed by the FrsApp service are:

- Creating a package, release, or file.
- Editing a package, release, or file.
- Returning a list of all packages in a project.
- Returning a list of all releases in a package.
- Returning the status, maturity value, or other attribute of a release.

After authenticating the user through the CollabNet service, you can access the FrsApp service to work with file releases.

## Detailed information

For a complete description of each method, including its parameter definitions and SOAP faults, see the [JavaDoc for the file release service](#).

## WSDL

The WSDL for this service is available on your TeamForge site at <https://<mysite.com>/ce-soap60/services/FrsApp?wsdl>. For example, here's the copy on the CollabNet web site: <http://ctf.open.collab.net/ce-soap60/services/FrsApp?wsdl>.

## New and changed methods

Refer to the Change History section on the [FrsApp service Javadoc](#) for a list of new and changed methods.

### *The file storage service*

The FileStorageApp and SimpleFileStorageApp services handle uploading of files and simple data bytes.

Some examples of activities that are managed by the FileStorageApp and SimpleFileStorageApp services are:

- Uploading a file to CollabNet TeamForge .
- Downloading a file from CollabNet TeamForge .
- Checking the size of a file.

After authenticating the user through the CollabNet service, you can then access the FileStorageApp and SimpleFileStorageApp services to work with file uploads.

## Detailed information

For a complete description of each method, including its parameter definitions and SOAP faults, see the [JavaDoc for the file storage service](#).

## WSDL


The WSDL for this service is available on your TeamForge site at <https://<mysite.com>/ce-soap60/services/FileStorageApp?wsdl>. For example, here's the copy on the CollabNet web site: <http://ctf.open.collab.net/ce-soap60/services/FileStorageApp?wsdl>.

## New and changed methods

Refer to the Change History section on the [FileStorageApp service Javadoc](#) and [SimpleFileStorageApp service Javadoc](#) for a list of new and changed methods.

### *The integration data service*

The IntegrationDataApp service enables applications to associate their own data or metadata with CollabNet TeamForge objects.

 **Tip:** For information about installing integrated apps on a TeamForge site, see the [Integrated applications](#).

An application can register its own namespace, to prevent collisions with other applications, and then store key#value pairs of data that are associated with any CollabNet TeamForge object.

- Developers who are creating integrations between CollabNet TeamForge and external applications can use this service to maintain associations between CollabNet TeamForge objects and counterparts in other systems.
- Developers creating customizations or enhancements can use this service to store additional data not directly supported by other CollabNet TeamForge APIs. For example, you can create extra, hidden fields on any object (such as an artifact, a user or a project) and use a Velocity customization to display that data where you need it.

## Detailed information

For a complete description of each method, including its parameter definitions and SOAP faults, see the [Javadoc for the integration data service](#).

## WSDL

The WSDL for this service is available on your TeamForge site at <https://<mysite.com>/ce-soap60/services/IntegrationDataApp?wsdl>. For example, here's the copy on the CollabNet web site: <http://ctf.open.collab.net/ce-soap60/services/IntegrationDataApp?wsdl>.

## New and changed methods

Refer to the Change History section on the [IntegrationDataApp service Javadoc](#) for a list of new and changed methods.

### *The news service*

The NewsApp service handles the activities and items associated with news posts.

News posts for a project are visible on the project's home page, and the application's main page lists news posts for all projects visible to the current user.

Starting with CollabNet TeamForge , project news is presented on the project's home page via a "Project News" page component, which can optionally be deleted or added to other pages. Each instance of the Project News component is a view of the same set of project news posts.

Some examples of activities that are managed by the NewsApp service are:

- Returning a list of all news posts in a project.
- Returning a list of all news posts matching a search string.
- Returning a list of all news posts in all projects of which a specified user is a member.
- Creating a news post.
- Deleting a news post.

After authenticating the user through the CollabNet service, you can access the NewsApp service to work with news posts.

## Detailed information

For a complete description of each method, including its parameter definitions and SOAP faults, see the [JavaDoc for the news app](#).

## WSDL

The WSDL for this service is available on your TeamForge site at <https://<mysite.com>/ce-soap60/services/NewsApp?wsdl>. For example, here's the copy on the CollabNet web site: <http://ctf.open.collab.net/ce-soap60/services/NewsApp?wsdl>.

### New and changed methods

Refer to the Change History section on the [NewsApp service Javadoc](#) for a list of new and changed methods.

#### *The planning folder service*

The Planning Folder service enables you to create and maintain planning folders from external applications.

Planning folders are visible in the tracker tool of each project, as an organizing level sitting on top of trackers. Planning folders are designed to support agile management practices.

Some examples of activities managed by the PlanningApp service are:

- Creating and deleting planning folders.
- Reordering planning folders.
- Setting the statuses of planning folders.

### Detailed information

For a complete description of each method, including its parameter definitions and SOAP faults, see the [JavaDoc for the planning folder app](#).

## WSDL

The WSDL for this service is available on your TeamForge site at <https://<mysite.com>/ce-soap60/services/PlanningApp?wsdl>. For example, here's the copy on the CollabNet web site: <http://ctf.open.collab.net/ce-soap60/services/PlanningApp?wsdl>.

### New and changed methods

Refer to the Change History section on the [PlanningApp service Javadoc](#) for a list of new and changed methods.

#### *The project pages service*

The PageApp service handles the activities and items associated with project pages.

The project pages feature enables project owners to customize their project home by creating their own custom pages. Multiple pages can be added, and these pages can be assembled into a hierarchy, browseable via a "tree" navigation on the left side of the project's home. The content of each page can be customized by adding components of various types, such as Text, News, and Document Folder.

Some examples of activities that are managed by the PageApp service are:

- Creating pages and components.
- Managing the page hierarchy.
- Managing the order of components on a page.
- Updating text components with new HTML content.

### Detailed information

For a complete description of each method, including its parameter definitions and SOAP faults, see the [JavaDoc for the project pages service](#).

## WSDL

The WSDL for this service is available on your TeamForge site at <https://<mysite.com>/ce-soap60/services/PageApp?wsdl>. For example, here's the copy on the CollabNet web site: <http://ctf.open.collab.net/ce-soap60/services/PageApp?wsdl>.

## New and changed methods

Refer to the Change History section on the [PageApp service Javadoc](#) for a list of new and changed methods.

### *The project categorization service*

The CategorizationApp service handles the activities and items associated with creation and management of Project Categorization.

Some examples of activities that are managed by the CategorizationApp service are:

- Creating categories
- Adding categories to projects
- Retrieving category and project data.

## Detailed information

For a complete description of each method, including its parameter definitions and SOAP faults, see the [JavaDoc for the CategorizationApp service](#).

## WSDL

The WSDL for this service is available on your TeamForge site at <https://<mysite.com>/ce-soap60/services/CategorizationApp?wsdl>. For example, here's the copy on the CollabNet web site: <http://ctf.open.collab.net/ce-soap60/services/CategorizationApp?wsdl>.

## New and changed methods

Refer to the Change History section on the [CategorizationApp service Javadoc](#) for a list of new and changed methods.

### *The role-based access control service*

The RbacApp service handles the activities and items associated with role creation and management.

Some examples of activities that are managed by the RbacApp service are:

- Creating a role.
- Adding a user to a role.
- Adding permissions to a role.
- Adding a user group to a role.

Permissions are managed in terms of “clusters.” Permission clusters correspond to the sets of permissions that are managed in the role administration section of the application’s user interface, such as “view” or “create/edit/view.” Clusters can be associated with top#level folders in the application (e.g. with trackers in the Tracker application, or packages in the File Releases application). The RbacApp service allows this same permission management to be performed via SOAP.

## Detailed information

For a complete description of each method, including its parameter definitions and SOAP faults, see the [JavaDoc for the role-based access control service](#).

## WSDL

The WSDL for this service is available on your TeamForge site at <https://<mysite.com>/ce-soap60/services/RbacApp?wsdl>. For example, here's the copy on the CollabNet web site: <http://ctf.open.collab.net/ce-soap60/services/RbacApp?wsdl>.

## New and changed methods

Refer to the Change History section on the [RbacApp service Javadoc](#) for a list of new and changed methods.

### *The software configuration management (SCM) service*

The ScmApp service handles the activities and items associated with integrated software configuration management (SCM) applications.

Some examples of activities that are managed by the ScmApp service are:

- Returning a list of files associated with a code commit.
- Returning a list of all repositories in a project.
- Returning a list of all code commits in a repository.
- Editing commit information.

After authenticating the user through the CollabNet service, you can access the ScmApp service to work with integrated SCM applications.

## Detailed information


For a complete description of each method, including its parameter definitions and SOAP faults, see the [JavaDoc for the software configuration management service](#).

## WSDL

The WSDL for this service is available on your TeamForge site at <https://<mysite.com>/ce-soap60/services/ScmApp?wsdl>. For example, here's the copy on the CollabNet web site: <http://ctf.open.collab.net/ce-soap60/services/ScmApp?wsdl>.

## SCM Adapter

The SCM Adapter allows you to develop integrations with almost any SCM tool, then exchange commit data with CollabNet TeamForge . An SCM Adapter option is added to the Type menu on the CollabNet TeamForge Create Integration page. Use this option to add your SCM integration server to CollabNet TeamForge .

 **Note:** The SCM Adapter option is always available, but will work only if you have developed an integration with an SCM tool using the CollabNet TeamForge API.

After you have created an SCM integration using the CollabNet TeamForge API, you will use the Commit Tool provided with CollabNet TeamForge to transmit your commit data to CollabNet TeamForge .

## New and changed methods

Refer to the Change History section on the [ScmApp service Javadoc](#) for a list of new and changed methods.

### *The tag service*

The TagApp service provides SOAP services for tag objects.

Some examples of activities that are managed by the TagApp service are:

- Create a tag in a specific project.
- Deleting a tag from a project.
- Get tag data for a given tag ID.
- Gets the list of tags in a project.

## Detailed information

For a complete description of each method, including its parameter definitions and SOAP faults, see the [JavaDoc for the tag service](#).



## WSDL

The WSDL for this service is available on your TeamForge site at <https://<mysite.com>/ce-soap60/services/TagApp?wsdl>. For example, here's the copy on the CollabNet web site: [https://www.open.collab.net/community/cif\\_sfee/samples/TagApp.xml](https://www.open.collab.net/community/cif_sfee/samples/TagApp.xml).

## New and changed methods

Refer to the Change History section on the [TagApp service Javadoc](#) for a list of new and changed methods.


### *The tasks service*

The TaskApp service handles the activities and items associated with tasks and task folders.

Some examples of activities that are managed by the Task App service are:

- Returning a list of all tasks assigned to a specified user.
- Returning a list of all task folders in a project.
- Returning a list of all tasks matching a search string.
- Listing task dependencies.
- Creating a task or task folder.
- Editing a task.
- Moving tasks and task folders.

After authenticating the user through the CollabNet service, you can access the TaskApp service to work with tasks and task folders.

 **Note:** Task folders are referred to as task groups in the method descriptions.

## Detailed information

For a complete description of each method, including its parameter definitions and SOAP faults, see the [JavaDoc for the tasks service](#).

## WSDL

The WSDL for this service is available on your TeamForge site at <https://<mysite.com>/ce-soap60/services/TaskApp?wsdl>. For example, here's the copy on the CollabNet web site: <http://ctf.open.collab.net/ce-soap60/services/TaskApp?wsdl>.

## New and changed methods

Refer to the Change History section on the [TaskApp service Javadoc](#) for a list of new and changed methods.

### *The tracker service*

The TrackerApp service handles the activities and items associated with trackers and tracker artifacts.

Some examples of activities that are managed by the TrackerApp service are:

- Returning a list of all trackers in a project.
- Returning a list of all tracker artifacts matching a search string.
- Returning a list of all tracker artifacts assigned to a specified user.
- Creating a tracker artifact.
- Editing a tracker artifact.
- Moving a tracker artifact.

After authenticating the user through the CollabNet service, you can access the TrackerApp service to work with trackers and tracker artifacts.

## Detailed information

For a complete description of each method in the TrackerApp service, including the SOAP data objects (DO), data lists, and data rows, see the [JavaDoc for the tracker service](#).

## WSDL

The WSDL for this service is available on your TeamForge site at <https://<mysite.com>/ce-soap60/services/TrackerApp?wsdl>. For example, here's the copy on the CollabNet web site: <http://ctf.open.collab.net/ce-soap60/services/TrackerApp?wsdl>.

## New and changed methods

Refer to the Change History section on the [TrackerApp service Javadoc](#) for a list of new and changed methods.

### *The wiki service*

The WikiApp service handles activities associated with project wiki pages.

Some examples of activities that are managed by the WikiApp service are:

- Creating a Wiki page.
- Adding Wiki content.
- Retrieving Wiki content in HTML format.

## Detailed information

For a complete description of each method, including its parameter definitions and SOAP faults, see the [JavaDoc for the wiki service](#).

## WSDL

The WSDL for this service is available on your TeamForge site at <https://<mysite.com>/ce-soap60/services/WikiApp?wsdl>. For example, here's the copy on the CollabNet web site: <http://ctf.open.collab.net/ce-soap60/services/WikiApp?wsdl>.

## New and changed methods

Refer to the Change History section on the [WikiApp service Javadoc](#) for a list of new and changed methods.

### **Context-specific objects for manipulating Velocity pages**

The rendering context is a set of objects made available to the Velocity template. You can access the members of these objects in a template using the standard Velocity syntax.

#### *ArrayTool*

This is necessary because Velocity does not provide accessors into arrays and some of the API calls will return SoapNamedValues objects that have two matching arrays of name and value.

#### **length(Object[] array)**

Get the length of the array.

##### **array**

The array being examined.

Returns: The length of the array (null array will return 0).

#### **get(Object[] array, int index)**

Get the object at the specified index of the array.

##### **array**

The array that contains the object of interest.

**index**

The index we are pulling the object from.

Returns: The object stored at that array or null if the array is empty or position is invalid.

**FORM [FormTool]**

Manages opening and closing of forms.

**startForm(String action, String formId)**

Get the opening tag for a form.

**action**

The path to the action that is handling the form (e.g. /cemain/do/ login).

**formId**

The id/name of the form (they will be the same).

Returns: The open tag for the form and any standard hidden elements.

**startForm(String action, String formId, String formName)**

Get the opening tag for a form.

**action**

The path to the action that is handling the form (e.g. /cemain/do/ login).

**formId**

The id of the form.

**formName**

The name of the form.

Returns: The open tag for the form and any standard hidden elements.

**endForm()**

Close out a form.

Returns: The form closing tag.

**LINK [LinkTool]**

Creates links; see default templates for examples.

**getPageUrl(String urlBase, String action)**

Get the url for a page being linked to.

**urlBase**

The base url.

**action**

The action being linked to.

Returns: The value to place in the href attribute of a link.

**getUserUrl(String username, String fullName)**

Generates URL for user details.

**username**

User name.

**fullName**

User's full name.

Returns: URL to user details page.

**MESSAGE [MessageTool]**

Enables internationalization of templates, providing functions for localizing strings.

**get(String bundle, String key[, String arg0[, String arg1[, String arg2]])**

Returns the message.

**bundle**

Resource bundle.

**key**

Message key.

**arg0**

Message argument.

**arg1**

Message argument.

**arg2**

Message argument.

Returns: Message.

**getFieldLabel(String bundle, String key, boolean required)**

Returns the message with the appropriate label modifier ('!' and possibly an asterisk if field is required).

**bundle**

The bundle where the message lives.

**key**

The key where the message is stored.

**required**

If the field is required.

Returns: The message with appropriate label modifiers.

**STRING [StringTool]**

Returns a formatted file size string, such as 45KB, 100MB, or 1GB

**formatFileSize(long fileSize)****fileSize**

File size in bytes.

Returns: File size string.

**wordWrap(String text, int width)**

Wraps text string at word boundaries.

**s**

String to word wrap.

**width**

Number of characters at which to wrap.

Returns: Word wrapped string.

### escapeXml(String s)

Escape out the xml from the text

s

The text that we need to escape xml from

Returns: The text with xml escaped

*TEXTPARSER [TextParserTool]*

Converts text into “linkified” html, with IDs and wiki words converted into appropriate links.

### parseText(String text)

Return linkified text (add href tags to object ids and urls).

text

Text to process.


Returns: Replacement text with link replacement.

### Example integration script

This integration script demonstrates how to use the SCM Adapter to create an integration with Subversion.

### Overview

This is a post#commit script that gets executed either from or as the hooks or post#commit trigger.

 **Note:** This example will work for Subversion only. It will not work in its present form for any other SCM tool. It is intended as an example that you can reference when creating scripts for other SCM tools. The example also contains error checking as an example only; post#commit triggers fail to abort a commit on Subversion.

```
#!/bin/sh
# Location of SOURCEFORGE_HOME
SOURCEFORGE_HOME="/opt/collabnet/teamforge/sourceforge_home"
# The Hostname of the application server
SOURCEFORGE_SOAP_HOST="ctf.open.collab.net"
# The port used for soap communication
SOURCEFORGE_SOAP_PORT="8080"
# The "Repository Path" of the repository to commit to.
SOURCEFORGE_REPOSITORY="myrepository"
# The external system id of the integration server.
SOURCEFORGE_SYSTEMID="exsy1001"
# Path to the svnlook command
SVNLOOK=/usr/bin/svnlook
# Path to the CommitTool.py command
COMMITTOOL="$SOURCEFORGE_HOME/integration/CommitTool.py"
# If the author metadata of the commit is empty, use this user. This
will normally
# happen if a commit occurs that bypasses authentication
DEFAULTUSER="admin"
# Repository storage path. This is passed in by Subversion
REPOS="$1"
# Transaction id of this commit. This is passed in by Subversion
TXN="$2"
# Execute an svnlook log to get the Commit Message of this commit
LOGMSG=`$SVNLOOK log -r "$TXN" "$REPOS"`
# Execute the svnlook command to get the author of the commit if
any.
```

```

AUTHOR=`"$SVNLOOK" author -r "$TXN" "$REPOS" `
if [ "X$AUTHOR" == "X" ]; then
AUTHOR="$DEFAULTUSER"
fi
$SVNLOOK changed -r "$TXN" "$REPOS" | (
# Create a CommitTool transaction.
"$COMMITTOOL" create "$SOURCEFORGE_SYSTEMID"
"$SOURCEFORGE_REPOSITORY" "$AUTHOR" \
"$SOURCEFORGE_SOAP_HOST" "$SOURCEFORGE_SOAP_PORT"
if [ $? -ne 0 ]; then
exit 1;
fi
read TYPE FILE
# Iterate through each line of the svnlook changes command and add a
file with
# the appropriate status to the CommitTool.py transaction.
while [ "X$TYPE" != "X" ]; do

# default to Added
COMMAND="Added"
case "$TYPE" in
U)
COMMAND="Modified"
;;
D)
COMMAND="Deleted"
;;
A)
COMMAND="Added"
;;
esac
"$COMMITTOOL" add "$FILE" "$TXN" "$COMMAND" >> /dev/null
if [ $? -ne 0 ]; then
exit 1;
fi
read TYPE FILE
done
# Submit the current transaction to sourceforge with the commit
message.
"$COMMITTOOL" commit "$LOGMSG" >> /dev/null

if [ $? -ne 0 ]; then
exit 1;
fi
)
if [ $? -ne 0 ]; then
exit 1;
fi
exit 0

```

### Context-specific object members

You can access the members of these objects in a template using the standard Velocity syntax.

For example:

```

${PAGE_INFO.actionName()}
${PAGE_INFO.isSuperUser()}
${GLOBAL.datePattern()}
${API44.discussionApp()}

```

etc.

**PAGE\_INFO [PageInformationTool]**

Allows access to information specific to the page that is passed in.

Information	Description
actionName	The name of the action that was requested (e.g. viewArtifact, listTrackers)
currentUser	User object that contains information about the current user: <ul style="list-style-type: none"> <li>• \$PAGE_INFO.currentUser.username (Login ID)</li> <li>• \$PAGE_INFO.currentUser.fullName (Full name)</li> <li>• \$PAGE_INFO.currentUser.email (E#mail address)</li> <li>• \$PAGE_INFO.currentUser.locale (Locale, such as en_US)</li> <li>• \$PAGE_INFO.currentUser.lastLogin (Date and time of last login)</li> </ul>
path	The full path string of the request's target object (e.g. projects.test/trackers.foo/bar). If no path is present on the request, this returns an empty string.
projectPath	The project path (e.g. projects.test) on the current folderPath or an empty string if none exists (no path context or just a project context).
itemName	The name of the item (e.g. artf1234, Home) on the current request or an empty string if there is no path context or just a project or folder context.
objectId	The id of the requested object (e.g. artf1234, proj1234).
objectType	The type string of the current object (i.e. An artifact would have Tracker.Artifact)
projectId	The ID of the project on the request or the project that contains the current folder or item.
requestUrl	The URL that was requested (e.g. /ce/tracker/do/viewArtifact/projects.test/tracker.foo/artf1234).
isSuperUser	True if the current user is a site admin.
isLoggedIn	True if the current user is logged into the system.

**GLOBAL [GlobalTool]**


Allows access to various kinds of static information about the site.

Static Information	Description
datePattern	The standard date pattern used for date conversions.
imageRoot	The root URL path of CollabNet TeamForge images.
isApprovalRequiredForNewUsers	A flag that indicates whether new users must be approved before they can use the site.
isRequireAssociationOnDocumentCreate	A flag that indicates whether creation of a document requires an association with an object.
isSelfCreationEnabled	A flag that indicates whether visitors are allowed to create new accounts in the application by themselves.
isUsingExternalAuthentication	A flag that indicates whether logins are authenticated against an external server. (Currently this means LDAP.)

**ApiTool**

API60 [ApiTool] exposes the CollabNet TeamForge SOAP service interfaces for use by Velocity templates.

While the interfaces corresponding to the SOAP API are provided, the implementation is efficient; the SOAP network protocol is not actually used.

 **Note:** API50 is present for backward compatibility and may be removed in future. API44 and API43 are removed completely.

**Interfaces**

Interface	Description
sessionKey	Gets the soap session id that should be passed in as the first parameter to soap service methods.
discussionApp	Gives access to the methods in the Discussion application.
documentApp	Gives access to the methods in the Documents application.
frsApp	Gives access to the methods in the File Releases application.
integrationDataApp	Gives access to the methods for Data Integration.
newsApp	Gives access to the methods in the Project News application.
rbacApp	Gives access to the methods for role#based access control.
scmApp	Gives access to the methods for version control (SCM) integration.
SourceForge	Gives access to the methods in the main CollabNet TeamForge application.
taskApp	Gives access to the methods in the Tasks application.
trackerApp	Gives access to the methods in the Tracker application.



Interface	Description
wikiApp	Gives access to the methods in the Wiki application.
categorizationApp	Gives access to the methods in project categorization.
emptyFilter	Gets an empty filter that can be passed into methods that require soap filters.
PageApp	Handles the activities and items associated with project pages.
PlanningApp	Gives access to methods in the Planning application.

### Example

This is an example of how you might use some popular API calls in a Velocity template:

```
## Display Server version, current project title and description
#set( $sessionKey = $API60.sessionKey )
#set( $projectId = $PAGE_INFO.projectId )
#set( $projectData = $API60.sourceForge.getProjectData($sessionKey,
$projectId) )


<p>Server Version: $API60.sourceForge.getVersion($sessionKey)</p>
<p>Project Title: ${projectData.title}</p>
<p>Project Desc: ${projectData.description}</p>
```

### Lab Management API methods

Almost all Lab Management functionality is available via an API method.

Lab Management API methods come in two forms: signed and unsigned.

- Signed methods require authentication, in accordance with [the User Authentication documentation](#). Most of the API methods in Lab Management, and all of the most useful ones, are *signed* methods which require authentication to use. The `cubit_api_client.py` client makes using these signed methods easy by handling the authentication negotiation for you.
- Unsigned methods do not require authentication.

 **Note:** This API is under active development. We have started with what we feel are the most common actions which would need to be performed on an automated basis, and will be expanding this list over time as we receive feedback about what our users want and need. Please let us know if there are missing features or bugs!

#### *CheckHostConsistency*

Check the setup of all hosts in the domain to make sure their representation is internally complete within Lab Management.

### Overview

This method is only available to users to Lab Management Domain Admins and can be considered a maintenance method that should rarely, if ever, need to be run. Three areas of setup consistency are checked for each host:

1. Verification and re-application of the host's internal database representation within Lab Management. This may be needed after an upgrade of the Lab Management application is installed.
2. Verification that the IP address assigned to the host matches the hostname. If you suspect DNS setup issues in your domain, this test should catch them (unless the errors are transient). Lab Management will not let you create hosts with invalid DNS, but if DNS for the host gets changed improperly after the host is created, your host will experience many problems in the Lab Management environment until it is fixed.
3. Verification of the IP mask/gateway configured for the host with the actual one it resolves to.

4. Re-generation of statistics system performance monitoring configuration files for the host. If monitoring is not properly working for a host, regenerating the configuration files will usually fix the problem.

Depending on the number of hosts in your domain, this method can take some time to complete. If you have questions on the output of this method, please see your local Lab Management Administrator or contact CollabNet directly for assistance.

## URL

/cubit\_api/1/check\_host\_consistency

## Authentication

This method requires [authentication](#) using an API key.

## Parameters

<b>sig ( Required, once )</b>	API authentication hash signature. Type: String
<b>userid ( Required, once )</b>	The login name of the user initiating the request. Type: String

## Example Response

Successful completion of host consistency check:

```
<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
  <output>All Hosts OK</output>
</cubit>
```

If the user is not a Lab Management Domain Admin:

```
<?xml version='1.0'?>
<cubit version='1'>
  <error>User 'grue' is not authorized to run this method.</error>
</cubit>
```

If there is an error re-generating the host's internal database representation for at least one host:

```
<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
  <output>Hosts which were not successfully updated:
  cu011.dev.cubitdemo.net</output>
</cubit>
```

If the system performance statistics configuration files are not properly generated for at least one host:

```
<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
  <output>Hosts which are not properly setup to collect
  performance data: cu011.dev.cubitdemo.net</output>
</cubit>
```

If there is a DNS mismatch for at least one host:

```
<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
```

```

        <output>Hosts which require re-ip'ing:
        cu011.dev.cubitdemo.net</output>
    </cubit>

    If there is a IP mask mis-match:
    <?xml version='1.0'?>
    <cubit version='1'>
        <status>OK</status>
        <output>Hosts whose IP mask/gateway does not match with the
        IP mask/gateway it resolves to:
        cu011.dev.cubitdemo.net</output>
    </cubit>

```

## Response Codes

- 200 - ok
- 400 - Login failed / Insufficient permissions

## CloudCreateHosts

Create one or more hosts from a cloud given the host-type and size.

The cloud will match the host-type and size to one (or more) source(s) from the cloud in order to fill the entire order. Host types and sizes cannot be mixed in a single API call. All the hosts must be the same type/size. The method will fail if the cloud's sources cannot accomodate that many hosts of the given type/size.

Any user who is a member of a project that has been given permission to allocate from this cloud can create instances.

## URL

/cubit\_api/1/allocate\_hosts\_from\_cloud

## Authentication

This method requires [authentication](#) using an API key.

## Parameters

<b>alloc_hours ( zero or once )</b>	Amount of time for which to allocate the host. The allocation limit is subject to the limit set by the project or cloud administrator. Type: Float
<b>alloc_minutes ( zero or once )</b>	Similar to alloc_hours parameter, but in minutes. It is mutually exclusive with the alloc_hours option, neither of these options should be specified, or only one of these options should be specified. If alloc_hours or alloc_minutes is set to 0, or if alloc_hours and alloc_minutes is unset, the allocation time defaults to the longest possible time available in the project. Type: Integer
<b>cloud ( Required, once )</b>	Name of Lab Management cloud to allocate instances from. Type: String
<b>count ( zero or once )</b>	The number of identical instances to bring up. The instances will all be running the same profile and

	version. Must be an integer greater than 0. If unset, defaults to 1. Type: Integer
<b>descr ( zero or once )</b>	Set the specified string as a description for the allocated hosts. This is useful if you wish to uniquely identify a group of hosts which have all been allocated to the same user at the same time for the same purpose. Type: String
<b>dont_delete ( zero or once )</b>	By default, EC2 instances will automatically be deleted when they are deallocated. If <code>dont_delete</code> is set to True, the instances created will not be deleted when they are deallocated, and will instead revert back to the project pool. Type: String
<b>host_type ( Required, once )</b>	The name of the <code>host_type</code> of machine to create. Must be one of the host-types supported by the cloud. Type: String
<b>profile ( Required, once )</b>	Name of the profile to assign to the host. The profile must, of course, be a profile that is eligible to be built by the host-type/size you selected. Type: String
<b>project ( Required, once )</b>	Name of Lab Management project to place hosts into. You must have permission to add hosts to this project: that is, you must either be a Project Admin or Delegated Host Management must be turned on in your project. Type: String
<b>revision ( zero or once )</b>	Revision number of the profile you wish to assign to the host. Mutually exclusive with the <code>version</code> option, but at least one of the <code>version</code> or <code>revision</code> options must be specified. Type: Integer
<b>sig ( Required, once )</b>	API authentication hash signature. Type: String
<b>size ( Required, once )</b>	The name of the machine size to create. Must be one of the sizes supported by the given <code>host_type</code> . Type: String
<b>userid ( Required, once )</b>	The login name of the user initiating the request. Type: String
<b>version ( zero or once )</b>	Version number or tag name of the profile you wish to assign to the host. The special version tag <code>HEAD</code> always denotes the latest version of the profile at the moment of execution. Note that profile tags can move between versions: this is a useful feature, but you should be aware of it. Mutually exclusive with the <code>revision</code> option, but at least one of the <code>version</code> or <code>revision</code> options must be specified. Type: String

## Example Response

Successful allocation of 10 instances:

```
<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
  <output>ec2-pending-1204832590-85</output>
  <output>ec2-pending-1204832595-09</output>
  <output>ec2-pending-1204832599-97</output>
  <output>ec2-pending-1204832611-49</output>
  <output>ec2-pending-1204832615-60</output>
  <output>ec2-pending-1204832620-29</output>
  <output>ec2-pending-1204832624-25</output>
  <output>ec2-pending-1204832628-46</output>
  <output>ec2-pending-1204832640-90</output>
  <output>ec2-pending-1204832645-60</output>
</cubit>
```

```
<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
</cubit>
```

If the user lacks permission to create hosts in the project:

```
<?xml version='1.0'?>
<cubit version='1'>
  <error>You are not authorized to create hosts in this project.</error>
</cubit>
```

If the project does not have the permissions to allocate from the requested cloud:

```
<?xml version='1.0'?>
<cubit version='1'>
  <error>This project is not authorized to create hosts in this cloud.</
error>
</cubit>
```

If the cloud does not have enough space on the sources to allocate that many hosts:

```
<?xml version='1.0'?>
<cubit version='1'>
  <error>This cloud does not have enough resources to create [n]
hosts.</error>
</cubit>
```

## Response Codes

- 400 - Login failed / Insufficient permissions

### *GetAuthToken*

GetAuthToken(userid) -> token

## URL

/cubit\_api/1/get\_auth\_token

## Authentication

This method does not require authentication.

## Parameters

This method does not take any parameters.

## Response Codes

This method does not have any documented response codes.

### *HostAllocate*

Allocate a host to a user.

Hosts must be in the `Free` state to be allocated, although users with Project Admin access can re-allocate hosts which are in the `Allocated` state to other users in their project. Domain Admin users can re-allocate machines between projects and users, although if a machine is reallocated to a user in another project, that user must also be a member of the destination project. The user must be authorized to allocate the host. If the host is in any of the following states: `Immutable`, `Powercycle`, `Rebuild`, or `Rebuilding`, this method will fail.

## URL

`/cubit_api/1/alloc_host`

## Authentication

This method requires [authentication](#) using an API key.

## Parameters

**alloc\_hours ( zero or once )**

Amount of time for which to allocate the host. The allocation limit is subject to the limit set by the project or cloud administrator.

Type: Float

**alloc\_minutes ( zero or once )**

Similar to `alloc_hours` parameter, but in minutes. It is mutually exclusive with the `alloc_hours` option, neither of these options should be specified, or only one of these options should be specified. If `alloc_hours` or `alloc_minutes` is set to 0, or if `alloc_hours` and `alloc_minutes` is unset, the allocation time defaults to the longest possible time available in the project.

Type: Integer

**alloc\_proj ( zero or once )**

The project to allocate the machine to. The `alloc_user` must be a valid Lab Management user in the project, or else this method will fail. Only Domain Admins can change the project that a host is assigned to. If not specified, the project is not changed. At least one of `alloc_user` or `alloc_proj` must be specified.

Type: String

**alloc\_user ( zero or once )**

Login name of user to allocate host to. Only Project Admins or better can specify a user other than themselves. Other users can only allocate machines to themselves. If not specified, defaults to `userid`. At least one of `alloc_user` or `alloc_proj` must be specified.

Type: String

<b>force ( zero or once )</b>	<p>If the machine is currently in the <code>Allocated</code> state, the <code>force</code> option must be given to reassign the host to another user. The <code>force</code> option is only available to Project Admins and above. If the host is a virtual host, and any virtual guests of this host are in the <code>Allocated</code> state, this option must also be used, or the entire allocation will fail. The only valid value for this parameter is <code>True</code>.</p> <p>Type: String</p>
<b>guests ( zero or once )</b>	<p>If the machine is a virtual host and has active virtual guests, setting this parameter to <code>True</code> will move the virtual guests along with the virtual host. Because moving a virtual host will move several hosts at once, we have a separate parameter to confirm this action. The only valid value for this parameter is <code>True</code>. If the host has no virtual guests, this option has no effect.</p> <p>Type: String</p>
<b>host ( Required, once )</b>	<p>Fully qualified hostname to allocate.</p> <p>Type: String</p>
<b>sig ( Required, once )</b>	<p>API authentication hash signature.</p> <p>Type: String</p>
<b>userid ( Required, once )</b>	<p>The login name of the user initiating the request.</p> <p>Type: String</p>

### Example Response

Successful host allocation:

```
<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
</cubit>
```

If the user is unauthorized, or the host does not exist:

```
<?xml version='1.0'?>
<cubit version='1'>
  <error>You are not authorized to allocate this host.</error>
</cubit>
```

If the host is allocated to a deleted project:

```
<?xml version='1.0'?>
<cubit version='1'>
  <error>Cannot allocate host 'cu001.dev.cubitdemo.net' to a
  deleted project 'look'</error>
</cubit>
```

### Response Codes

- 200 - ok
- 400 - Login failed / Insufficient permissions

#### *HostAssignProfileAndRebuild*

Assign a profile to the host and rebuild it with that profile/version.

The user must be authorized to rebuild the host. If the host is in any of the states which disallow rebuilding, such as `Immutable`, `Rebuild`, or `Rebuilding`, this method will fail. The host must also not be of a type which prohibits rebuilding (e.g., an EC2 host). If you do not want to change the profile a host is running, and just want to rebuild a host, you can save a few keystrokes and use the `Rebuild` method instead.

## URL

```
/cubit_api/1/assign_rebuild
```

## Authentication

This method requires [authentication](#) using an API key.

## Parameters

<b>host ( Required, once )</b>	Fully qualified hostname to assign profile to and rebuild. Type: String
<b>profile ( Required, once )</b>	Name of the profile to assign to the host. Type: String
<b>revision ( zero or once )</b>	Revision number of the profile you wish to assign to the host. Mutually exclusive with the <code>version</code> option, but at least one of the <code>version</code> or <code>revision</code> options must be specified. Type: String
<b>sig ( Required, once )</b>	API authentication hash signature. Type: String
<b>userid ( Required, once )</b>	The login name of the user initiating the request. Type: String
<b>version ( zero or once )</b>	Version number or tag name of the profile you wish to assign to the host. The special version tag <code>HEAD</code> always denotes the latest version of the profile at the moment of execution. Note that profile tags can move between versions: this is a useful feature, but you should be aware of it. Mutually exclusive with the <code>revision</code> option, but at least one of the <code>version</code> or <code>revision</code> options must be specified. Type: String

## Example Response

Successful initiation of rebuild:

```
<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
</cubit>
```

If the user is unauthorized, or the host does not exist:

```
<?xml version='1.0'?>
<cubit version='1'>
  <error>You are not authorized to allocate a profile to this host and
rebuild it.</error>
</cubit>
```



## Response Codes

- 200 - ok
- 400 - Login failed / Insufficient permissions

### *HostDelete*

Delete one or more hosts.

The host must not be in the `Immutable` state for this method to work. Hosts can be deleted by users under the following conditions:

- Lab Management Domain Admins can delete any host, regardless if it is physical, virtual, or from a remote cloud
- Users who are *not* Lab Management Domain Admins can delete *virtual guests* if the parent host is allocated to them and Delegated Host Management is turned on in their project.
- Users who are *not* Lab Management Domain Admins can also delete *remote cloud* hosts allocated to them if Delegated Host Management is turned on in their project.

## URL

```
/cubit_api/1/delete_host
```

## Authentication

This method requires [authentication](#) using an API key.

## Parameters

### **force ( zero or once )**

If the machine being deleted does not belong to the user initiating the request, the `force` option is required. The only valid value for this parameter is `True`.

Type: String

### **guests ( zero or once )**

If the machine is a virtual host and has active virtual guests, setting this parameter to `True` will delete the virtual guests along with the virtual host. Because deleting a virtual host will delete several hosts at once, we have a separate parameter to confirm this action. The only valid value for this parameter is `True`. If the host has no virtual guests, this option has no effect. Use of this option can be really convenient, or it can really ruin someone's day!

Type: String

### **hosts ( Required, once )**

List of comma-separated, fully qualified hostnames to delete. Hosts must not be in the `Immutable` state.

Type: String

### **sig ( Required, once )**

API authentication hash signature.

Type: String

### **userid ( Required, once )**

The login name of the user initiating the request.

Type: String

## Example Response

```

Successful host delete of the host cu011.cubit.example.com:
<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
  <output>cu011.cubit.example.com</output>
</cubit>

If the user is unauthorized:
<?xml version='1.0'?>
<cubit version='1'>
  <error>0 of 1 hosts deleted successfully. Failed hosts below.</error>
  <output>cu013.cubit.example.com: You are not authorized to delete this
host.</output>
</cubit>

If a deletion of hosts partially succeeds:
<?xml version='1.0'?>
<cubit version='1'>
  <error>6 of 8 hosts deleted successfully. Failed hosts below.</error>
  <output>cu013.cubit.example.com: You are not authorized to delete this
host</output>
  <output>cu016.cubit.example.com: Host deletion failed</output>
</cubit>

```

## Response Codes

- 200 - ok
- 400 - Login failed / Insufficient permissions

## HostFree

Free a host currently assigned to a user.

Hosts must be in the *Allocated* state to be freed. A user can only free hosts already assigned to them, although Project Admins and Domain Admins can free hosts in their project which are currently allocated. If it is a virtual guest and "Delete on Deallocation" flag is ON, the virtual guest would be deleted after it is set Free.

## URL

/cubit\_api/1/free\_host

## Authentication

This method requires [authentication](#) using an API key.

## Parameters

<b>force (zero or once)</b>	If the machine being freed does not belong to the user initiating the request, the <i>force</i> option is required. Only Project Admins and Domain Admins can free hosts which do not belong to them. The only valid value for this parameter is <code>True</code> . Type: String
<b>host (Required, once)</b>	Fully qualified hostname to free. Hosts must be in the <i>Allocated</i> state to be freed. Type: String
<b>sig (Required, once)</b>	API authentication hash signature. Type: String
<b>userid (Required, once)</b>	The login name of the user initiating the request.

Type: String

### Example Response

```

Successful freeing of host:
<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
</cubit>

If the user is unauthorized, or the host does not exist:
<?xml version='1.0'?>
<cubit version='1'>
  <error>You are not authorized to free this host.</error>
</cubit>

```

### Response Codes

- 200 - ok
- 400 - Login failed / Insufficient permissions

#### *HostGetAssignedProfile*

Get the assigned profile for the host, and returns the profile name along with the revision and version number.

Output is available in both text and XML formats. You can also get this information from the [QueryAlloc](#) on page 873 method, but the output of this method is easier to parse if you have scripts checking a profile. The user requesting information about a host must have at least login privileges to that host.

### URL

```
/cubit_api/1/assigned_profile
```

### Authentication

This method requires [authentication](#) using an API key.

### Parameters

<b>host ( Required, once )</b>	Fully qualified hostname to query. Type: String
<b>output ( Required, once )</b>	Output type, either 'txt' (for text output) or 'xml' (for XML output). Type: String
<b>sig ( Required, once )</b>	API authentication hash signature. Type: String
<b>userid ( Required, once )</b>	The login name of the user initiating the request. Type: String

### Example Response

```

Sample XML output:
<?xml version='1.0'?>

```

```

<cubit version='1'>
  <profile revision="1838" version="5">solaris10_x86</profile>
</cubit>

Sample text output:
solaris10_x86 1838 5

If user is unauthorized:
<?xml version='1.0'?>
<cubit version='1'>
  <error>You do not have permissions to view this host</error>
</cubit>

```

## Response Codes

- 200 - ok
- 400 - Login failed / Insufficient permissions

### *HostGetMyAssignedProfile*

Get the profile and version assigned to the host invoking this method.

## URL

/cubit\_api/1/my\_profile

## Authentication

This method does not require authentication.

## Parameters

<b>output ( Required, once )</b>	Output type, either 'txt' (for text output) or 'xml' (for XML output). Type: String
----------------------------------	--

## Example Response

```

Sample XML output:
<?xml version='1.0'?>
<cubit version='1'>
  <profile revision="26" version="2">rhel3_base</profile>
</cubit>

Sample text output:
rhel3_base 26 2

```

## Response Codes

- 200 - ok

### *HostPerfColl*

Get performance collection data from remote hosts managed by Lab Management.

This method takes one parameter, *data*. The hostname is auto-derived using reverse DNS. There is no other way to verify the sending host.

The currently supported statgroups are:

- `net`: Returns a dictionary of network interface transfer metrics. The key is the name of the interface. Value is a tuple containing: (*bytes\_in*, *bytes\_out*, *errors\_in*, *errors\_out*, *pkts\_in*, *pkts\_out*)
- `loadavg`: Returns a tuple containing the (1, 5, 15) minute load average on the system.
- `storage`: Returns a dictionary of disk storage metrics. The key is the name of the mounted partition and the value is a tuple of: (*block\_size*, *blocks\_avail*, *blocks\_used*, *block\_errors*)
- `systatd`: Returns a string containing the `systatd` output file for the system.

## URL

`/cubit_api/1/host_perfcoll`

## Authentication

This method does not require authentication.

## Parameters

This method does not take any parameters.

## Response Codes

This method does not have any documented response codes.

### *HostPowercycle*

Powercycle a host.

This method immediately executes a powercycle of the host. No graceful shutdown is run on the host at the operating system level. It is recommended that this method only be run when the normal operating shutdown/reboot sequence has failed or is unavailable. You can only powercycle machines in the `Allocated`, `Free`, and `Immutable` states.

## URL

`/cubit_api/1/powercycle`

## Authentication

This method requires [authentication](#) using an API key.

## Parameters

<b>force ( zero or once )</b>	If the machine being powercycled is a virtual host, the <code>force</code> option is required. The only valid value for this parameter is <code>True</code> . Type: String
<b>host ( Required, once )</b>	Fully qualified hostname to powercycle. Type: String
<b>sig ( Required, once )</b>	API authentication hash signature. Type: String
<b>userid ( Required, once )</b>	The login name of the user initiating the request. Type: String

## Example Response

```
Successful completion of powercycle:
```

```

<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
</cubit>

If the user is unauthorized, or the host does not exist:
<?xml version='1.0'?>
<cubit version='1'>
  <error>You are not authorized to powercycle this host.</error>
</cubit>

```

## Response Codes

- 200 - ok
- 400 - Login failed / Insufficient permissions

### *HostRebuildCancel*

Cancel a rebuild for a host.

The host must be in `Rebuild` state for this method to work. In order to cancel a rebuild for a host, the user must have permission to rebuild the host, that is, they must either:

- Have the host allocated to them
- Be a Lab Management Project Admin in the project, or
- Be a Lab Management Domain Admin

## URL

/cubit\_api/1/rebuild\_cancel

## Authentication

This method requires [authentication](#) using an API key.

## Parameters

<b>host ( Required, once )</b>	Fully qualified hostname to cancel rebuild of. Hosts must be in the <code>Rebuild</code> state. Type: String
<b>sig ( Required, once )</b>	API authentication hash signature. Type: String
<b>userid ( Required, once )</b>	The login name of the user initiating the request. Type: String

## Example Response

```

Successful host rebuild cancel:
<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
</cubit>

If the user is unauthorized:
<?xml version='1.0'?>
<cubit version='1'>

```

```
<error>You are not authorized to cancel a rebuild for this host.</error>
</cubit>
```

## Response Codes

- 200 - ok
- 400 - Login failed / Insufficient permissions

### *HostRebuild*

Rebuild a host with the currently selected profile and profile version.

The user must be authorized to rebuild the host. The host must also not be in any of the states which prohibit rebuilding, such as *Immutable*, *Powercycle*, *Rebuild*, or *Rebuilding*, or this method will fail.

The host must also indicate that it is rebuildable (e.g., not an EC2 host), or this method will fail.

## URL

/cubit\_api/1/rebuild

## Authentication

This method requires [authentication](#) using an API key.

## Parameters

<b>host ( Required, once )</b>	Fully qualified hostname to rebuild. Type: String
<b>sig ( Required, once )</b>	API authentication hash signature. Type: String
<b>userid ( Required, once )</b>	The login name of the user initiating the request. Type: String

## Example Response

Successful initiation of rebuild:

```
<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
</cubit>
```

If the user is unauthorized, or the host does not exist:

```
<?xml version='1.0'?>
<cubit version='1'>
  <error>You are not authorized to rebuild this host.</error>
</cubit>
```

## Response Codes

- 200 - ok
- 400 - Login failed / Insufficient permissions

### *PblChangeDesc*

Project Build Library (PBL) file and directory description changing interface. This function is used to change the description for files and directories in the Project Build Library. This function may be useful if you are building your own PBL client.

### URL

/cubit\_api/1/pbl\_changedesc

### Authentication

This method requires *authentication* using an API key.

### Parameters

<b>comment ( zero or once )</b>	An optional comment to leave about the operation being performed. The comment will not appear in the PBL, but it will be in the audit log entry for this event. Type: String
<b>desc ( Required, once )</b>	The new text description of the file. This description will completely replace any description currently set for the file. Type: String
<b>path ( Required, once )</b>	The path to the file being operated on. For example, if the complete file URL is /pbl/zork/pub/foo/bar/test.txt, the path would be /foo/bar/test.txt Type: String
<b>proj ( Required, once )</b>	The name of the project which contains the file we are operating on. Type: String
<b>sig ( Required, once )</b>	API authentication hash signature. Type: String
<b>type ( Required, once )</b>	The type of file that we are operating on. Valid values are 'pub' and 'priv'. Type: String
<b>userid ( Required, once )</b>	The login name of the user initiating the request. Type: String

### Example Response

```
<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
</cubit>
```

### Response Codes

- 200 - ok
- 400 - Login failed / Insufficient permissions



### *PblDelete*

Project Build Library (PBL) file and directory delete interface. This function is used to permanently remove files and directories in the Project Build Library. This function may be useful if you are building your own PBL client.

### URL

/cubit\_api/1/pbl\_delete

### Authentication

This method requires [authentication](#) using an API key.

### Parameters

<b>comment ( zero or once )</b>	An optional comment to leave about the operation being performed. This will not appear in the PBL, but it will be in the audit log entry for this event. Type: String
<b>dryrun ( Required, once )</b>	If this parameter is set to <code>True</code> , the specified path will not actually be deleted. The only valid value for this parameter is <code>True</code> . Type: String
<b>force ( zero or once )</b>	If the <code>force</code> option is present and set to <code>True</code> , and the specified <code>path</code> argument is a directory, a recursive delete of the directory will be performed. The only valid value for this parameter is <code>True</code> . Type: String
<b>path ( Required, once )</b>	The path to the file being operated on. For example, if the complete file URL is <code>/pbl/zork/pub/foo/bar/test.txt</code> , the path would be <code>/foo/bar/test.txt</code> Type: String
<b>proj ( Required, once )</b>	The name of the project which contains the file we are operating on. Type: String
<b>sig ( Required, once )</b>	API authentication hash signature. Type: String
<b>type ( Required, once )</b>	The type of file that we are operating on. Valid values are 'pub' and 'priv'. Type: String
<b>userid ( Required, once )</b>	The login name of the user initiating the request. Type: String

### Example Response

```
<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
</cubit>
```

## Response Codes

- 200 - ok
- 400 - Login failed / Insufficient permissions

### *PblMove*

Project Build Library (PBL) file and directory move interface. This function may be useful if you are building your own PBL client.

## URL

/cubit\_api/1/pbl\_move

## Authentication

This method requires [authentication](#) using an API key.

## Parameters

### **comment ( zero or once )**

An optional comment to leave about the operation being performed. This will not appear in the PBL, but it will be in the audit log entry for this event.

Type: String

### **destpath ( Required, once )**

The path to move the destination file to. For example, if the complete file URL is /pbl/zork/pub/foo/bar/test.txt, the srcpath would be /foo/bar/test.txt. Two important things to note about this parameter:

1. If you specify a path which does not exist, that path will be automatically created for you as part of the move.
2. If the destpath parameter ends with a slash ("/"), the destination will be assumed to be a directory. If it does not end with a slash, the destination will be assumed to be a file.

Type: String

### **destprj ( zero or once )**

The destination project for the file. You must have permissions to upload files into both the srcprj and the destprj, or the move will fail. If not specified, defaults to srcprj.

Type: String

### **desttype ( zero or once )**

The destination type of the file. Valid values are 'pub' and 'priv'. If not specified, defaults to srcprj.

Type: String

### **force ( zero or once )**

If the force option is present, the file will be moved even if a file with that name currently exists, and the previous file will be deleted and replaced with this file. The only valid value for this parameter is True.

Type: String

### **sig ( Required, once )**

API authentication hash signature.

<b>srcpath ( Required, once )</b>	Type: String The path to the source file being operated on. For example, if the complete file URL is /pbl/zork/pub/foo/bar/test.txt, the srcpath would be /foo/bar/test.txt
<b>srcprj ( Required, once )</b>	Type: String The project in which the source file is located.
<b>srctype ( Required, once )</b>	Type: String The source type of the file. Valid values are 'pub' and 'priv'.
<b>userid ( Required, once )</b>	Type: String The login name of the user initiating the request.

### Example Response

```
<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
</cubit>
```

### Response Codes

- 200 - ok
- 400 - Login failed / Insufficient permissions

### PblUpload

Project Build Library (PBL) file and directory upload interface. This function is used to create files and directories in the Project Build Library. Unless you are building your own PBL client, this function will not be very useful to you.

### URL

/cubit\_api/1/pbl\_upload

### Authentication

This method requires [authentication](#) using an API key.

### Parameters

<b>comment ( zero or once )</b>	An optional comment to leave about the operation being performed. This will not appear in the PBL, but it will be in the audit log entry for this event. Type: String
<b>desc ( zero or once )</b>	The text description of the file. Type: String
<b>file ( Required, once )</b>	URL-encoded name and contents of the file, encoded as per <a href="#">RFC 1867</a> .

<b>force ( zero or once )</b>	Type: String If the <code>force</code> option is present, the file will be uploaded even if a file with that name currently exists. A file cannot be uploaded over a directory of the same name, even if the <code>force</code> parameter is present. The only valid value for this parameter is <code>True</code> .
<b>md5sum ( Required, once )</b>	Type: String The md5 checksum of the file being uploaded.
<b>path ( Required, once )</b>	Type: String The path to the file being operated on. For example, if the complete file URL is <code>/pbl/zork/pub/foo/bar/test.txt</code> , the <code>path</code> would be <code>/foo/bar</code> . Note: the <code>path</code> parameter is used slightly differently in this method than in other methods, since the filename is not appended to the end of the parameter.
<b>proj ( Required, once )</b>	Type: String The name of the project which contains the file we are operating on.
<b>sig ( Required, once )</b>	Type: String API authentication hash signature.
<b>type ( Required, once )</b>	Type: String The type of file that we are operating on. Valid values are 'pub' and 'priv'.
<b>userid ( Required, once )</b>	Type: String The login name of the user initiating the request.

### Example Response

```
<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
</cubit>
```

### Response Codes

- 200 - ok
- 400 - Login failed / Insufficient permissions

#### *ProfileAdd*

Add a profile to the system.

The following conditions must hold for a profile to be added by name:

- The `profile_name` must be specified.
- The `profile_file` must not be specified.
- The profile must already exist in SVN.
- The user calling this method must be a Domain Admin.

The following conditions must hold for a profile to be uploaded:

- The `profile_file` must be specified.
- The `profile_name` must not be specified
- The user calling this method must have permission to add profiles to the specified project.

## URL

`/cubit_api/1/add_profile`

## Authentication

This method requires [authentication](#) using an API key.

## Parameters

<b><code>can_users_modify</code> ( Required, once )</b>	Set this to <code>True</code> if user's with Root Access or better should be able to make modifications to the project. If set to <code>False</code> , only the owning user and users with Project Admin access (or better) will be able to modify the profile. Type: String
<b><code>is_prebuilt</code> (Required, once )</b>	Set this to <code>True</code> if this profile is a prebuilt image, <code>False</code> if this profile is installed via a network install method. Type: String
<b><code>is_public</code> ( Required, once )</b>	If the profile is public (i.e can be built by any project). Valid values are <code>True</code> or <code>False</code> . Type: String
<b><code>owner</code> ( zero or once )</b>	Lab Management user that this profile will be associated with. The user must be a member of the owning project in order to be able to modify the profile. If unset, this profile will not belong to any particular user. Type: String
<b><code>profile_file</code> ( zero or once )</b>	The filename of a profile to be uploaded. The specified file must: <ol style="list-style-type: none"> <li>1. Exist on your local system.</li> <li>2. Contain a valid profile name. Profile names can only contain letters, numbers, underscores ("_"), and dashes ("-"). Profile names are case sensitive.</li> <li>3. Not currently exist as a profile name on the system.</li> <li>4. Be a valid XML file, and contain all the needed attributes for a Lab Management profile.</li> </ol>
<b><code>profile_name</code> ( zero or once )</b>	The name of a profile that already exists in Lab Management's Subversion Repository. Type: String
<b><code>project</code> ( zero or once )</b>	Project that will own this profile. If unset, this profile will not belong to any particular project. Type: String
<b><code>sig</code> ( Required, once )</b>	API authentication hash signature. Type: String

<b>summary ( Required, once )</b>	Summary of the purpose of this profile. Type: String
<b>userid ( Required, once )</b>	The login name of the user initiating the request. Type: String

### Example Response

```

Successful addition of the profile rhel3_base:
<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
  <output>rhel3_base</output>
</cubit>

If the user is unauthorized:
<?xml version='1.0'?>
<cubit version='1'>
  <error>Failed to add profile rhel3_base.</error>
  <output>User 'userid' does not have Domain Admin rights.</output>
</cubit>

```

### Response Codes

- 200 - ok
- 400 - Login failed / Insufficient permissions

#### *ProfileDelete*

Delete a profile from the system.

Profiles can be deleted under the following conditions:

- No host has ever been built or is currently building with this profile.
- The user calling this method has permission to modify this profile.

### URL

```
/cubit_api/1/delete_profile
```

### Authentication

This method requires [authentication](#) using an API key.

### Parameters

<b>profiles ( Required, once )</b>	List of comma-separated, profile names to delete. Profiles must never have been used to build any hosts. The user invoking this API must have modify rights on the profiles. Type: String
<b>sig ( Required, once )</b>	API authentication hash signature. Type: String
<b>userid ( Required, once )</b>	The login name of the user initiating the request. Type: String

## Example Response

```

Successful delete of the profile rhel3_base:
<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
  <output>rhel3_base</output>
</cubit>

If the user is unauthorized:
<?xml version='1.0'?>
<cubit version='1'>
  <error>0 of 1 profiles deleted successfully. Failed profiles below.</
error>
  <output>rhel3_base: You do not have modify rights on this profile.</
output>
</cubit>

If a deletion of profiles partially succeeds:
<?xml version='1.0'?>
<cubit version='1'>
  <error>7 of 8 profiles deleted successfully. Failed profiles below.</
error>
  <output>rhel3_base: This profile has been used to build hosts.</
output>
</cubit>

```

## Response Codes

- 200 - ok
- 400 - Login failed / Insufficient permissions

### *ProfileGetMyPkglist*

List the packages associated with the profile assigned to the calling host.

## URL

/cubit\_api/1/getmypkglist

## Authentication

This method does not require authentication.

## Parameters

<b>output ( Required, once )</b>	Output type, either 'txt' (for text output) or 'xml' (for XML output). Type: String
----------------------------------	--

## Example Response

```

Sample XML output:
<?xml version='1.0'?>
<cubit version='1'>
  <package version="1.4.3-6">libtool-libs</package>
  <package version="2.1.15-10">cyrus-sasl-md5</package>
  <package version="20020927-11.30.4">iputils</package>

```

```

<package version="3.03-28">perl-HTML-Tagset</package>
<package version="1.4.1-2">python-optik</package>
...
<package version="1.7-23">time</package>
<package version="3.2.3-54">cpp</package>
<package version="8.3.5-92.4">tcl</package>
<package version="0.2.3-7.1">audiofile</package>
</cubit>

```

```

Sample text output:
libtool-libs 1.4.3-6
cyrus-sasl-md5 2.1.15-10
iputils 20020927-11.30.4
perl-HTML-Tagset 3.03-28
python-optik 1.4.1-2
...
time 1.7-23
cpp 3.2.3-54
tcl 8.3.5-92.4
audiofile 0.2.3-7.1

```

## Response Codes

- 200 - ok

### *ProfileGetPkglst*

List packages associated with a particular version or revision of a profile.

If neither version nor revision is set, the latest (HEAD) revision of the profile is selected. You must have permissions to view the profile in order to retrieve a package listing. The profile must be public, or you must be a valid user in the project that owns the profile.

## URL

/cubit\_api/1/getpkglst

## Authentication

This method requires [authentication](#) using an API key.

## Parameters

<b>output ( zero or once )</b>	Output mode. 'txt' or 'xml' Type: String
<b>profile ( Required, once )</b>	Profile to get package list for. Type: String
<b>revision ( zero or once )</b>	Revision of profile to get package list for. Type: Integer
<b>sig ( Required, once )</b>	API authentication hash signature. Type: String
<b>userid ( Required, once )</b>	The login name of the user initiating the request. Type: String
<b>version ( zero or once )</b>	Version of profile to get package list for. Type: Integer



## Example Response

```
(output='xml') :
<?xml version='1.0'?>
<cubit version='1'>
  <package version="1.2.3">pkgname1</package>
  <package version="4.5.6">pkgname2</package>
  <package version="7.8.9">pkgname3</package>
</cubit>

(output='txt') :
pkg1  1.2.3
pkg2  4.5.6
pkg3  7.8.9
```

## Response Codes

- 200 - ok
- 400 - Login failed / Insufficient permissions

### QueryAlloc

Determine, at either a *domain*, *project*, *user*, or *host* level, how hosts are allocated and who they are allocated to.

Users are only allowed to get information about hosts which they have access to see. For example, if you are not a member of a project, you cannot query that project's host allocations. If you query the domain, you will only see allocation information for projects to which you belong. This method is used by Lab Management internally, and it is made available for users as well, since they may find this data useful.

## URL

/cubit\_api/1/query\_alloc

## Authentication

This method requires [authentication](#) using an API key.

## Parameters

### **full** ( zero or once )

If full=true then the complete information about the host is printed.

Type: String

### **maxhosts** ( zero or once )

The maximum number of hosts to return for the query. If type=domain, this will be the maximum number of hosts to return per project (i.e. if this is set to 5 and there are two projects in the domain, up to 10 hosts may be returned). If type=host, this parameter will not be used. If maxhosts is set, the returned hosts will be sorted by the host's name.

Type: Integer

### **name** ( zero or once )

The name of the host, project, or user to be queried. If type=domain, this argument is not required. If type={host,project} then this argument is required. If type=user, this argument is optional, defaulting to the name of the user who you authenticate as. Only Lab

	Management Domain Admins are allowed to query user allocations for users other than themselves.
	Type: String
<b>sig ( Required, once )</b>	API authentication hash signature.
	Type: String
<b>skiphosts ( zero or once )</b>	The number of initial hosts that will be skipped over when returning the results. This only has an effect when maxhosts is also used. If type=domain, the hosts skipped are per per project. If type=hosts, this parameter will not be used.
	Type: Integer
<b>type ( Required, once )</b>	The type of service. Valid values are host, project, user and domain.
	Type: String
<b>userid ( Required, once )</b>	The login name of the user initiating the request.
	Type: String

### Example Response

Here is a sample output where we query type=host and name=cu011.dev.cubitdemo.net:

```
<?xml version='1.0'?>
<cubit version='1'>
  <project name="zork">
    <host version="2">
      <name>cu011.dev.cubitdemo.net</name>
      <alloc_user>yz</alloc_user>
      <alloc_project>zork</alloc_project>
      <labels>
        <label name="cubit_user">yz</label>
        <label name="cubit_project">zork</label>
      </labels>
    </host>
  </project>
</cubit>
```

Here is a sample output where we query type=host, name=cu011.dev.cubitdemo.net and full=True:

```
<?xml version='1.0'?>
<cubit version='1'>
  <project name=zork>
    <host version="2">
      <name>cu011.dev.cubitdemo.net</name>
      <alloc_user>yz</alloc_user>
      <alloc_project>zork</alloc_project>
      <hardware>
        <vendor>Collabnet</vendor>
        <modelname>HP DL380</modelname>
        <ncpu>2</ncpu>
        <cpuarch>i386</cpuarch>
        <cpuname>Xeon</cpuname>
        <cpumhz>1400</cpumhz>
        <memsize unit="MB">2048</memsize>
```

```

    <disksize unit="GB">144</disksize>
</hardware>
<labels>
  <label name='cfprofile' rev='322' version='4'>rhel3_base</label>
    <label name='install_mode'>kickstart-net</label>
  <label name='cubit_state'>Allocated</label>
  <label name='cubit_alloc_time'>1170799967</label>
  <label name='cubit_state_mtime'>1170799967</label>
  <label name='cubit_build_time'>1169682748</label>
  <label name='cubit_project'>zork</label>
  <label name='cubit_user'>yz</label>
  <label name='vmware-server' />
  <label name='cubit_lom_addr'>cu3-1.sjc.collab.net</label>
</labels>
<groups>
  <aigroup>DEFAULT</aigroup>
</groups>
<network>
  <routing>
    <default>
      <ipv4_addr>192.168.202.1</ipv4_addr>
    </default>
  </routing>
  <interface>
    <name>eth0</name>
    <sol_name>e1000g0</sol_name>
    <mac_addr>00:16:35:c3:9b:ed</mac_addr>
    <bootproto>static</bootproto>
    <addr>
      <ipv4_addr>192.168.202.21</ipv4_addr>
      <ipv4_mask>255.255.255.0</ipv4_mask>
      <dnsname>cu011.dev.cubitdemo.net</dnsname>
    </addr>
  </interface>
</network>
</host>

  </project>
</cubit>

```

Here is a sample output where we query type=domain:

```

<?xml version='1.0'?>
<cubit version='1'>
  <domain>
    <project name="zork">
      <host>
        <name>cu011.dev.cubitdemo.net</name>
        <alloc_user>yz</alloc_user>
        <alloc_project>zork</alloc_project>
        <labels>
          <label name="cubit_user">yz</label>
          <label name="cubit_project">zork</label>
        </labels>
      </host>
      <host>
        <name>cu012.dev.cubitdemo.net</name>
        <alloc_user>yz</alloc_user>
        <alloc_project>zork</alloc_project>
        <labels>
          <label name="cubit_user">yz</label>
          <label name="cubit_project">zork</label>
        </labels>
      </host>
    </project>
  </domain>
</cubit>

```

```

    </project>
  </domain>
</cubit>

```

Here is a sample output where we query type=domain and full=True:

```

<?xml version='1.0'?>
<cubit version='1'>
  <domain>
    <project name="zork">
      <host version="2">
        <name>cu011.dev.cubitdemo.net</name>
        <alloc_user>yz</alloc_user>
        <alloc_project>zork</alloc_project>
        <hardware>
          <vendor>Collabnet</vendor>
          <modelName>HP_DL380</modelName>
          <ncpu>2</ncpu>
          <cpuarch>i386</cpuarch>
          <cpuname>Xeon</cpuname>
          <cpumhz>1400</cpumhz>
          <memsize unit="MB">2048</memsize>
          <disksize unit="GB">144</disksize>
        </hardware>
        <labels>
          <label name='cfprofile' rev='322' version='4'>rhel3_base</label>
          <label name='install_mode'>kickstart-net</label>
          <label name='cubit_state'>Allocated</label>
          <label name='cubit_alloc_time'>1170799967</label>
          <label name='cubit_state_mtime'>1170799967</label>
          <label name='cubit_build_time'>1169682748</label>
          <label name='cubit_project'>zork</label>
          <label name='cubit_user'>yz</label>
          <label name='vmware-server' />
          <label name='cubit_lom_addr'>cu3-1.sjc.collab.net</label>
        </labels>
        <groups>
          <aigroup>DEFAULT</aigroup>
        </groups>
        <network>
          <routing>
            <default>
              <ipv4_addr>192.168.202.1</ipv4_addr>
            </default>
          </routing>
          <interface>
            <name>eth0</name>
            <sol_name>e1000g0</sol_name>
            <mac_addr>00:16:35:c3:9b:ed</mac_addr>
            <bootproto>static</bootproto>
            <addr>
              <ipv4_addr>192.168.202.21</ipv4_addr>
              <ipv4_mask>255.255.255.0</ipv4_mask>
              <dnsname>cu011.dev.cubitdemo.net</dnsname>
            </addr>
          </interface>
        </network>
      </host>

      <host version="2">
        <name>cu012.dev.cubitdemo.net</name>
        <alloc_user>yz</alloc_user>
        <alloc_project>zork</alloc_project>
        <hardware>

```

```

    <vendor>Collabnet</vendor>
    <modelname>HP_DL380</modelname>
    <ncpu>1</ncpu>
    <cpuarch>i386</cpuarch>
    <cpuname>Xeon</cpuname>
    <cpumhz>1400</cpumhz>
    <memsize unit="MB">2048</memsize>
    <disksize unit="GB">144</disksize>
</hardware>
<labels>
  <label name='cfprofile' rev='322' version='4'>rhel3_all</label>
  <label name='install_mode'>kickstart-net</label>
  <label name='cubit_state'>Allocated</label>
  <label name='cubit_alloc_time'>1170800012</label>
  <label name='cubit_state_mtime'>1170800012</label>
  <label name='cubit_build_time'>1169687903</label>
  <label name='cubit_project'>zork</label>
  <label name='cubit_user'>yz</label>
  <label name='vmware-server' />
  <label name='cubit_lom_addr'>cu4-1.sjc.collab.net</label>
</labels>
<groups>
  <aigroup>DEFAULT</aigroup>
</groups>
<network>
  <routing>
    <default>
      <ipv4_addr>192.168.202.1</ipv4_addr>
    </default>
  </routing>
  <interface>
    <name>eth0</name>
    <sol_name>e1000g0</sol_name>
    <mac_addr>00:16:35:5a:fe:39</mac_addr>
    <bootproto>static</bootproto>
    <addr>
      <ipv4_addr>192.168.202.22</ipv4_addr>
      <ipv4_mask>255.255.255.0</ipv4_mask>
      <dnsname>cu012.dev.cubitdemo.net</dnsname>
    </addr>
  </interface>
</network>
</host>
</project>
</domain>
</cubit>

```

Here is a sample output where we query type=project and name=zork:

```

<?xml version='1.0'?>
<cubit version='1'>
  <project name="zork">
    <host version="2">
      <name>cu011.dev.cubitdemo.net</name>
      <alloc_user>yz</alloc_user>
      <alloc_project>zork</alloc_project>
      <labels>
        <label name="cubit_user">yz</label>
        <label name="cubit_project">zork</label>
      </labels>
    </host>
    <host version="2">
      <name>cu012.dev.cubitdemo.net</name>
      <alloc_user>yz</alloc_user>

```

```

    <alloc_project>zork</alloc_project>
    <labels>
      <label name="cubit_user">yz</label>
      <label name="cubit_project">zork</label>
    </labels>
  </host>
</host version="2">
  <name>cu013.dev.cubitdemo.net</name>
  <alloc_user>grue</alloc_user>
  <alloc_project>zork</alloc_project>
  <labels>
    <label name="cubit_user">grue</label>
    <label name="cubit_project">zork</label>
  </labels>
</host>
</host version="2">
  <name>cu014.dev.cubitdemo.net</name>
  <alloc_user>root</alloc_user>
  <alloc_project>zork</alloc_project>
  <labels>
    <label name="cubit_user">root</label>
    <label name="cubit_project">zork</label>
  </labels>
</host>
</host version="2">
  <name>cu019.dev.cubitdemo.net</name>
  <alloc_user>root</alloc_user>
  <alloc_project>zork</alloc_project>
  <labels>
    <label name="cubit_user">root</label>
    <label name="cubit_project">zork</label>
  </labels>
</host>
</host version="2">
  <name>cu022.dev.cubitdemo.net</name>
  <alloc_user>grue</alloc_user>
  <alloc_project>zork</alloc_project>
  <labels>
    <label name="cubit_user">grue</label>
    <label name="cubit_project">zork</label>
  </labels>
</host>
</host version="2">
  <name>cu091.dev.cubitdemo.net</name>
  <alloc_user>root</alloc_user>
  <alloc_project>zork</alloc_project>
  <labels>
    <label name="cubit_user">root</label>
    <label name="cubit_project">zork</label>
  </labels>
</host>
</project>
</cubit>

```

Here is a sample output where we query type=project, name=zork and full=True:

```

<?xml version='1.0'?>
<cubit version='1'>
  <project name="zork">
    <host version="2">
<name>cu011.dev.cubitdemo.net</name>
<alloc_user>yz</alloc_user>
<alloc_project>zork</alloc_project>

```

```

<hardware>
  <vendor>Collabnet</vendor>
  <modelname>HP DL380</modelname>
  <ncpu>2</ncpu>
  <cpuarch>i386</cpuarch>
  <cpuname>Xeon</cpuname>
  <cpumhz>1400</cpumhz>
  <memsize unit="MB">2048</memsize>
  <disksize unit="GB">144</disksize>
</hardware>
<labels>
  <label name='cfprofile' rev='322' version='4'>rhel3_base</label>
  <label name='install_mode'>kickstart-net</label>
  <label name='cubit_state'>Allocated</label>
  <label name='cubit_alloc_time'>1170799967</label>
  <label name='cubit_state_mtime'>1170799967</label>
  <label name='cubit_build_time'>1169682748</label>
  <label name='cubit_project'>zork</label>
  <label name='cubit_user'>yz</label>
  <label name='vmware-server' />
  <label name='cubit_lom_addr'>cu3-1.sjc.collab.net</label>
</labels>
<groups>
  <aigroup>DEFAULT</aigroup>
</groups>
<network>
  <routing>
    <default>
      <ipv4_addr>192.168.202.1</ipv4_addr>
    </default>
  </routing>
  <interface>
    <name>eth0</name>
    <sol_name>e1000g0</sol_name>
    <mac_addr>00:16:35:c3:9b:ed</mac_addr>
    <bootproto>static</bootproto>
    <addr>
      <ipv4_addr>192.168.202.21</ipv4_addr>
      <ipv4_mask>255.255.255.0</ipv4_mask>
      <dnsname>cu011.dev.cubitdemo.net</dnsname>
    </addr>
  </interface>
</network>
</host>

  <host version="2">
<name>cu012.dev.cubitdemo.net</name>
<alloc_user>yz</alloc_user>
<alloc_project>zork</alloc_project>
<hardware>
  <vendor>Collabnet</vendor>
  <modelname>HP DL380</modelname>
  <ncpu>1</ncpu>
  <cpuarch>i386</cpuarch>
  <cpuname>Xeon</cpuname>
  <cpumhz>1400</cpumhz>
  <memsize unit="MB">2048</memsize>
  <disksize unit="GB">144</disksize>
</hardware>
<labels>
  <label name='cfprofile' rev='322' version='4'>rhel3_all</label>
  <label name='install_mode'>kickstart-net</label>
  <label name='cubit_state'>Allocated</label>
  <label name='cubit_alloc_time'>1170800012</label>

```

```

<label name='cubit_state_mtime'>1170800012</label>
<label name='cubit_build_time'>1169687903</label>
<label name='cubit_project'>zork</label>
<label name='cubit_user'>yz</label>
<label name='vmware-server' />
<label name='cubit_lom_addr'>cu4-1.sjc.collab.net</label>
</labels>
<groups>
  <aigroup>DEFAULT</aigroup>
</groups>
<network>
  <routing>
    <default>
      <ipv4_addr>192.168.202.1</ipv4_addr>
    </default>
  </routing>
  <interface>
    <name>eth0</name>
    <sol_name>e1000g0</sol_name>
    <mac_addr>00:16:35:5a:fe:39</mac_addr>
    <bootproto>static</bootproto>
    <addr>
      <ipv4_addr>192.168.202.22</ipv4_addr>
      <ipv4_mask>255.255.255.0</ipv4_mask>
      <dnsname>cu012.dev.cubitdemo.net</dnsname>
    </addr>
  </interface>
</network>
</host>

  <host version="2">
<name>cu013.dev.cubitdemo.net</name>
<alloc_user>grue</alloc_user>
<alloc_project>zork</alloc_project>
<hardware>
  <vendor>Collabnet</vendor>
  <modelname>HP DL380</modelname>
  <ncpu>1</ncpu>
  <cpuarch>i386</cpuarch>
  <cpuname>Xeon</cpuname>
  <cpumhz>1400</cpumhz>
  <memsize unit="MB">1024</memsize>
  <disksize unit="GB">25</disksize>
</hardware>
<uuid>51 51 bb a0 00 15 15 51-81 9d fb 47 77 00 51 61</uuid>
<labels>
  <label name='cfprofile' rev='322' version='6'>rhel4_all</label>
  <label name='install_mode'>kickstart-net</label>
  <label name='cubit_state'>Allocated</label>
  <label name='cubit_alloc_time'>1170960674</label>
  <label name='cubit_state_mtime'>1170963906</label>
  <label name='cubit_build_time'>1170963906</label>
  <label name='cubit_project'>zork</label>
  <label name='cubit_user'>grue</label>
  <label name='vmware-guest'>cu011.dev.cubitdemo.net</label>
  <label name='cubit_lom_addr'>cu011.dev.cubitdemo.net</label>
</labels>
<groups>
  <aigroup>DEFAULT</aigroup>
</groups>
<network>
  <routing>
    <default>
      <ipv4_addr>192.168.202.1</ipv4_addr>

```



```

    </default>
  </routing>
  <interface>
    <name>eth0</name>
    <sol_name>e1000g0</sol_name>
    <mac_addr>00:50:56:00:00:01</mac_addr>
    <bootproto>static</bootproto>
    <addr>
      <ipv4_addr>192.168.202.23</ipv4_addr>
      <ipv4_mask>255.255.255.0</ipv4_mask>
      <dnsname>cu013.dev.cubitdemo.net</dnsname>
    </addr>
  </interface>
</network>
</host>
  </project>
</cubit>

```

## Response Codes

- 200 - ok
- 400 - Login failed / Insufficient permissions

### QueryMembership

Determine the projects and the hosts that the user has permission to access.

If name parameter is not passed, the permission details of the initiating user will be displayed. Only Lab Management Domain Admins are permitted to query users other than themselves.

## URL

/cubit\_api/1/query\_membership

## Authentication

This method requires [authentication](#) using an API key.

## Parameters

<b>name ( zero or once )</b>	The name of the user to get the permission details. This argument is optional, defaulting to the name of the user who you authenticate as. Only Lab Management Domain Admins are permitted to query users other than themselves. Type: String
<b>sig ( Required, once )</b>	API authentication hash signature. Type: String
<b>userid ( Required, once )</b>	The login name of the user initiating the request. Type: String

## Example Response

Here is a sample output where we query name=grue:

```

<?xml version='1.0'?>

```

```

<cubit version='1'>
  <user name="grue">
    <project name="testing">
      <summary>Cubit Testing</summary>
      <roles>CUBIT - Domain Admin</roles>
    </project>
    <project name="zim">
      <summary>invador</summary>
      <roles>CUBIT - Domain Admin</roles>
    </project>
    <project name="zork">
      <summary>For dorks</summary>
      <roles>CUBIT - Domain Admin</roles>
    </project>
  </user>
</cubit>

```

## Response Codes

- 200 - ok
- 400 - Login failed / Insufficient permissions

### Status

Return the status of the Lab Management web service. This provides some indication of whether or not the Lab Management web service is working properly. The only status returned is OK.

## URL

/cubit\_api/1/status

## Authentication

This method does not require authentication.

## Parameters

This method does not take any parameters.

## Example Response

```

<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
</cubit>

```

## Response Codes

- 200 - ok

### StatusSigned

Return the status of the Lab Management web service, and the name of the authenticated user when authentication is used via a signed method.

This provides proof that the API, and authentication to the API are working properly. The only status returned is OK. If the authentication is not successful, you will get a `Permission Denied` error. If the API key is expired, you will get a `Permission Denied: API key expired` error.

## URL

/cubit\_api/1/status\_signed

## Authentication

This method requires [authentication](#) using an API key.

## Parameters

<b>sig ( Required, once )</b>	API authentication hash signature. Type: String
<b>userid ( Required, once )</b>	The login name of the user initiating the request. Type: String

## Example Response

```
<?xml version='1.0'?>
<cubit version='1'>
  <status>OK</status>
  <userid>grue</userid>
</cubit>
```

In addition to returning the OK message, the username of the authenticated user is provided.

## Response Codes

- 200 - ok
- 400 - Login failed / Insufficient permissions

### *UserCanLogin*

Given a userid and a hostname, return whether or not the user is authorized to login provided that the appropriate credentials are presented.

The HTTP response code will always be 200 for a successful query whether or not access is permitted. The document body will be text only, and contain one of the following:

- 0: no access
- 1: user level access
- 2: root level access
- 3: host owner access
- 4: project admin access

## URL

/cubit\_api/1/can\_login

## Authentication

This method does not require authentication.

## Parameters

<b>host ( Required, once )</b>	The fully qualified hostname of the host to test for access. Type: String
--------------------------------	--

**userid ( Required, once )**

The Lab Management userid of the user to test for access.

Type: String

**Example Response**

3

**Response Codes**

- 200 - ok

**TeamForge EventQ API documentation**

TeamForge EventQ APIs can be classified into three groups: Extensibility APIs, Reporting APIs and HTTP APIs.

The API documentation here discusses the following APIs:

- Extensibility APIs: To add new integrations to EventQ.
- Reporting API: To retrieve data from the EventQ database.
- HTTP APIs: To administrate EventQ programmatically.

**Extending TeamForge EventQ**

TeamForge EventQ provides a message queue (MQ) based interface for adding custom adapters and the Extensible Data Source (XDS) system for adding new product domains and tools.

**Overview**


EventQ can be extended via the Message Queue (MQ) API. The MQ API is provided for purposes of creating custom adapters for stock product classes/domains; build, commit, code review, and work item sources can post activity information to the corresponding message queues using the formats described in the API Documentation (see references below).

New product classes may be established using the Extensible Data Source API, or XDS API. For instance, TeamForge EventQ does not have stock product class for "deployment", so the XDS API may be used to establish a new product domain for "deployment" which can then be used like other sources.

**MQ API — Writing your own adapters**

Each TeamForge EventQ instance has an associated "message queue" service, which is used to collect and queue data from integrated tools. Extensibility is achieved by writing "adapters" that communicate specially formatted messages to the TeamForge EventQ message queue service.

The message communication protocol is called "AMQP" and the [RabbitMQ web site](#) links to open source AMQP protocol libraries that can be used to create adapters for numerous languages and platforms. The API documentation referenced below houses detailed message format specifications, example messages, and code samples for interacting with AMQP libraries.



TeamForge EventQ is extensible via MQ for these types of tools: SCM and version control, continuous integration/build, code review, and work item systems.

### XDS API — Authoring new product/tool domains

EventQ can be configured to accept data from a wide range of sources through the XDS API. In addition to the stock activity APIs defined for commits, builds, reviews, and work items, the XDS feature enables you to create your own activity APIs by defining and registering an "XDS Schema" that describes the expected activity message format for a particular tool or product domain in question. Once an XDS Schema is registered with EventQ, adapters can send EventQ activity messages in compliance with the registered schema and the resulting XDS activities may show up in activity streams and as associations. In this way, XDS Schemas give users a way to publish new activity APIs for arbitrary, user-defined data sources.

XDS Schemas can be added to EventQ either through the MQ or HTTP API. Please read the [Extensible Data Sources \(XDS\) walk-through](#) on page 907 and refer to the API documentation below.

Once a Schema is established, "XDS Sources" can be added through the "Custom" step type. See [Extensible Data Source \(XDS\) Overview](#) on page 580.

## Authentication

For the MQ interface, authentication is provisioned by the RabbitMQ server. When a new source is created, a unique username/password combination is created and exposed on the EventQ source edit screen. These credentials must be supplied by the posting client. These user credentials are independent of the TeamForge user store and belong only to RabbitMQ.

## TeamForge EventQ Commits API

### Submission parameters

#### Submission queue

- **queue\_name**: eventq.commits
- **auto\_delete**: false
- **durable**: true

#### Top-level

Parameter	Required/Optional	Description
<b>api_version</b>	Required	A string that matches the API version this document is for.
<b>source_association_key</b>	Required	A key generated by TeamForge EventQ that links incoming commit messages with the appropriate source server.  Example: "source_association_key": "31e8dd90-164f-0130-c4d0-406c8f04c05d"
<b>commit_data</b>	Required	The commit.

#### Objects

*commit\_data* - The following fields should be in this commit post.

Field	Required/Optional	Description
<b>repo_id</b>	Optional	A unique identifier for this repository. For example, the repo UUID.
<b>revision_id</b>	Required	A unique identifier for the commit action (eg., revision number or commit id).
<b>branch_name</b>	Optional	Identifier for the applicable repository branch.

Field	Required/ Optional	Description
<b>type</b>	Required	A string representing the activity type.
		<b>Accepted values</b> <i>post-commit</i> - A commit has been completed.
<b>message</b>	Optional	The message associated with this activity.
<b>changes</b>	Optional	An array of changed items. Required information about affected paths (an array of hashes).
		<b>Accepted values</b> <ul style="list-style-type: none"> <li>• <b>path</b> (Required) - A string with the relative path to the file that created, modified, etc.</li> <li>• <b>action</b> (Required) - A string representing the action. <ul style="list-style-type: none"> <li>• added - Added</li> <li>• deleted - Deleted</li> <li>• modified - Modified</li> <li>• props_modified - File properties (not contents) modified</li> <li>• type_changed - Type changed (e.g., from 'file' to 'symlink', 'symlink' to 'directory', etc.)</li> <li>• copied - Path is copied from another source</li> <li>• renamed - Path is renamed from another source</li> </ul> </li> <li>• <b>from</b> (Optional) - If the activity is a copy or rename, the path from which it was copied or renamed (if it can be determined or inferred with some accuracy; this is not guaranteed to be absolutely correct)</li> </ul>
<b>event_time</b>	Required	A string containing a timestamp in UTC timezone

Field	Required/ Optional	Description
		and RFC 3339 format.
<b>created_by</b>	Required	The user name of the person who took the action.

## Examples

The following code examples are Copyright 2015 CollabNet, Inc., licensed under the Apache License, Version 2.0 (the "License"); you may not use this code except in compliance with the License. You may obtain a copy of the License at [<http://www.apache.org/licenses/LICENSE-2.0>]

### Sample Submission JSON

```
{
  "api_version": "1",
  "source_association_key" :
  "150e8951-19e1-4066-8b74-601026dd6cde",
  "commit_data" : {
    "repo_id":
    "390e8951-19e1-4066-8b74-601026dd6cde",
    "revision_id": "2315",
    "branch_name": "master",
    "type": "post_commit",
    "message": "here's my commit
[orc-277]",
    "changes": [
      { "path": "/source/
file/1", "action": "modified" },
      { "path": "/source/
file/2", "action": "added" },
      { "path": "/source/
file/3", "action": "deleted" },
      { "path": "/source/
file/4", "action": "modified" },
      { "path": "/"
source/file/4", "action":
"props_modified" }
    ],
    "event_time":
    "2012-10-02T17:15:32.320Z",
    "created_by": "username"
  }
}
```

### Ruby Submission Example

You will need to change the following example before you can run it. Set the Rabbit MQ hostname and the association key, and paste a JSON block following the commit format specified on this page.

```
require 'amqp'

QUEUE = 'eventq.commits'
COMMIT = '{
  "api_version": "1",
  "source_association_key" :
  "150e8951-19e1-4066-8b74-601026dd6cde",
  "commit_data" : {
```



```

      "repo_id":
"390e8951-19e1-4066-8b74-601026dd6cde",
      "revision_id": "2315",
      "branch_name": "master",
      "type": "post_commit",
      "message": "here\'s my
commit [orc-277]",
      "changes": [
        { "path": "/source/
file/1", "action": "modified" },
        { "path": "/source/
file/2", "action": "added" },
        { "path": "/source/
file/3", "action": "deleted" },
        { "path": "/source/
file/4", "action": "modified" },
        { "path": "/
source/file/4", "action":
"props_modified" }
      ],
      "event_time":
"2012-10-02T17:15:32.320Z",
      "created_by": "username"
    }
  }'

# Event loop
EventMachine.run do
  connection =
  AMQP.connect('amqp://
guest:guest@example-mq')

  # Set up our RabbitMQ
  information
  channel =
  AMQP::Channel.new(connection)
  queue =
  channel.queue(Queue, :auto_delete
=> false, durable: true)
  exchange =
  channel.direct('')

  # Publish the activity and
  exit the loop
  exchange.publish
  COMMIT, :routing_key => queue.name
  do
    connection.disconnect
  {EventMachine.stop}
  end
end
end

```

## TeamForge EventQ Builds API

### Submission parameters

#### Submission queue

- **queue\_name:** eventq.builds
- **auto\_delete:** false
- **durable:** true

**Top-level**

Parameter	Required/ Optional	Description
<b>api_version</b>	Required	A string that matches the API version this document is for.
<b>source_association_key</b>	Required	A key generated by TeamForge EventQ that links incoming builds with the appropriate source server.  Example: "source_association_key": "6t5qM5AuLLWLkmqFJeKp"
<b><i>build_data</i></b>	Required	The raw build data, as provided by the build system.

**Objects**

*build\_data* - The following fields should be in this build object.

Field	Required/ Optional	Description
<b>remote_id</b>	Required	A string to uniquely identify the build. It will appear in the UI.
<b>duration</b>	Required	A string containing time (in seconds) it took for a build agent to build.
<b>event_time</b>	Required	A string containing a timestamp in UTC timezone and RFC 3339 format.
<b>build_url</b>	Required	A string with the URL to the build summary. Please note that the fully qualified URL with the URL scheme is required.
<b>created_by</b>	Optional	A string containing the username of

Field	Required/ Optional	Description
		the person who triggered the build.
<b>status</b>	Required	An object that declares the status of the build object.

#### Accepted values

- **type** (Required) - A string representing the type of the status. Please note that any type other than the listed below will be used as UNKNOWN inside TeamForge EventQ.
  - SUCCESS
  - FAILURE
  - UNSTABLE
  - ABORTED
- **name** (Required) - A user-friendly display name for the status.

<b>test_results</b>	Optional	An object containing the test results summary.
---------------------	----------	--

#### Accepted values

- **passed\_count** - A string containing the number of passed tests.
- **failed\_count** - A string containing the number of failed tests.
- **ignored\_count** - A string containing the number of ignored tests.
- **url** - A string containing the URL to the test results. Please note that the fully qualified URL with the URL scheme is required.

<b>revisions</b>	Optional	An array of objects containing the SCM revisions related to this build.
------------------	----------	---

#### Accepted values

Field	Required/ Optional	Description
		<ul style="list-style-type: none"> <li>• <b>revision</b> (Required) - A string containing the commit revision id.</li> <li>• <b>repository_url</b> (Required) - A string containing the repository url.</li> </ul>
<b>event_time</b>	Required	A string containing a timestamp in UTC timezone and RFC 3339 format.
<b>created_by</b>	Required	The user name of the person who took the action.

## Examples

The following code examples are Copyright 2015 CollabNet, Inc., licensed under the Apache License, Version 2.0 (the "License"); you may not use this code except in compliance with the License. You may obtain a copy of the License at [<http://www.apache.org/licenses/LICENSE-2.0>]

### Sample Submission JSON

```
{
  "api_version": "1",
  "source_association_key":
  "29bf0c90-3b40-0130-ae2d-dddd5893",
  "build_data": {
    "remote_id": "700",
    "event_time":
    "2012-10-02T17:15:32.320Z",
    "duration": "200",
    "build_url": "http://
example.com/builds/600",
    "created_by": "admin",
    "status": {
      "type": "SUCCESS",
      "name": "SUCCESSFUL"
    },
    "test_results": {
      "passed_count": "10",
      "failed_count": "0",
      "ignored_count": "0",
      "url": "http://
example.com/builds/600/tests"
    },
    "revisions":
    [{ "revision":
    "fb849a2440dda438f4c6ab25f8c3266ed82d8797",
      "repository_url":
    "ssh://gitserver/project"
    }]
  }
}
```

## Ruby Submission Example

You will need to change the following example before you can run it. Set the Rabbit MQ hostname and the association key, and paste a JSON block following the build format specified on this page.

```

}

require 'amqp'

  QUEUE = 'eventq.builds'
  BUILD = '{
    "api_version": "1",
    "source association key":
"29bf0c90-3b40-0130-ae2d-dddd5893",
    "build_data": {
      "remote_id": "700",
      "event_time":
"2012-10-02T17:15:32.320Z",
      "duration": "200",
      "build_url": "http://
example.com/builds/600",
      "created_by": "admin",
      "status": {
        "type": "SUCCESS",
        "name": "SUCCESSFUL"
      },
      "test_results": {
        "passed_count":
"10",
        "failed_count": "0",
        "ignored_count":
"0",
        "url": "http://
example.com/builds/600/tests"
      },
      "revisions":
      [{ "revision":
"fb849a2440dda438f4c6ab25f8c3266ed82d8797",
        "repository_url":
"ssh://gitserver/project"
      }]
    }
  }'
  # Event loop
  EventMachine.run do
    connection =
  AMQP.connect('amqp://
guest:guest@example-mq')

    # Set up our RabbitMQ
  information
    channel =
  AMQP::Channel.new(connection)
    queue =
  channel.queue(QUEUE, :auto_delete
=> false, durable: true)
    exchange =
  channel.direct('')

    # Publish the build and exit
  the loop

```

```

        exchange.publish
BUILD, :routing_key => queue.name
do
    connection.disconnect
{EventMachine.stop}
end
end

```

## TeamForge EventQ Review Requests API

### Submission parameters

#### Submission queue

- **queue\_name**: eventq.reviews
- **auto\_delete**: false
- **durable**: true

#### Top-level

Parameter	Required/ Optional	Description
<b>api_version</b>	Required	A string that matches the API version this document is for.
<b>source_association_key</b>	Required	A key generated by TeamForge EventQ that links incoming review requests with the appropriate source server.  Example: "source_association_key": "31e8dd90-164f-0130-c4d0-406c8f04c05d"
<i>request_data</i>	Required	The raw review request data, as provided by the review server.

#### Objects

*request\_data* - The following fields should be in this review request object.

Field	Required/ Optional	Description
<b>remote_id</b>	Required	A string to uniquely identify the review request.
<b>link</b>	Optional	A URL that points back to this individual review.

Field	Required/ Optional	Description
<b>summary</b>	Optional	A string to summarize the review request.
<b>description</b>	Optional	A longer string that can be used to further describe the review request. It will appear in the UI.
<b>created_by</b>	Required	A string containing the username of the person who submitted the request (e.g., initiated this request event, NOT necessarily the creator of the review; in some instances, this may consistently be an administrative user).
<b>status</b>	Required	An object that declares the status of the review request.

#### Accepted values

- **type** (Required) - A string representing the type of the status.
  - open - The review request is currently open for review.
  - submitted - The review request changes have been submitted.
  - rejected - The code in the request was rejected.
  - discarded - The review request has been discarded.
  - other - A status type that may not fit into the above list.
- **name** (Required) - A user-friendly display name for the status.

Field	Required/ Optional	Description
<b>event_time</b>	Required	A string containing a timestamp in UTC timezone and RFC 3339 format.
<b>updated_by</b>	Optional	A string containing the username of the person performing the update.
<b>files</b>	Optional	An array of objects with optional SCM revision information. If included, only the file name is required.

#### Accepted values

- **file** (Required) - A string with the full path to the file that was being reviewed.
- **revision** (Optional) - A string representing the revision number of the commit that includes this file.
- **repository\_url** (Optional) - A string with the URL to the SCM server. For associations to commits, this must match one defined as a source in the same pipeline.

<b>reviewers</b>	Optional	A array of user hashes. Optional reviewer information. If included, only the user name is required. For some systems (e.g., Crucible), reviewers are only reported when they have marked the review 'completed' or 'reviewed'; other
------------------	----------	--



Field	Required/ Optional	Description
		systems (e.g., Review Board) report reviewers when they add comments or make other changes to a review.

#### Accepted values

- **username** (Required) - A string of the user's name.
- **link** (Optional) - A string representing the URL to the users profile.

## Examples

The following code examples are Copyright 2015 CollabNet, Inc., licensed under the Apache License, Version 2.0 (the "License"); you may not use this code except in compliance with the License. You may obtain a copy of the License at [<http://www.apache.org/licenses/LICENSE-2.0>]

### Sample Submission JSON

```
{
  "api_version": "1",
  "source_association_key" :
  "31e8dd90-164f-0130-
c4d0-406c8f04c05d",
  "request_data" : {
    "remote_id": "42",
    "link": "http://reviewserver/
review/42/",
    "summary": "The summary of the
request",
    "description": "An extended
description of what the request
covers",
    "status": {
      "type": "open",
      "name": "Review"
    },
    "event_time":
    "2012-10-02T17:15:32.320Z",
    "created_by": "username",
    "updated_by": "username",
    "files": [
      {
        "revision":
        "ef9f8fe661992f48c3539c05fc5c69c61918fa17",
        "file":
        "commit1_file1.txt",
        "repository_url": "ssh://
gitserver/project"
      },
      {
        "revision":
        "ef9f8fe661992f48c3539c05fc5c69c61918fa17",
```

```

        "file":
        "commit1_file2.txt",
        "repository_url": "ssh://
gitserver/project"
    },
    {
        "revision":
        "fb849a2440dda438f4c6ab25f8c3266ed82d8797",
        "file":
        "commit2_file1.txt",
        "repository_url": "ssh://
othergitserver/project"
    }
],
"reviewers": [
    { "username": "adent",
      "link": "http://reviewserver/user/
adent" },
    { "username": "fpreferct",
      "link": "http://reviewserver/user/
fpreferct" }
]
}
}

```

### Ruby Submission Example

You will need to change the following example before you can run it. Set the Rabbit MQ hostname and the association key, and paste a JSON block following the review format specified on this page.

```

require 'amqp'

QUEUE = 'eventq.reviews'
REVIEW = '{
  "api_version": "1",
  "source_association_key" :
  "31e8dd90-164f-0130-
c4d0-406c8f04c05d",
  "request_data" : {
    "remote_id": "42",
    "link": "http://
reviewserver/review/42/",
    "summary": "The summary of
the request",
    "description": "An extended
description of what the request
covers",
    "status": {
      "type": "open",
      "name": "Review"
    },
    "event_time":
    "2012-10-02T17:15:32.320Z",
    "created_by": "username",
    "files": [
      {
        "revision":
        "ef9f8fe661992f48c3539c05fc5c69c61918fa17",
        "file":
        "commit1_file1.txt",
        "repository_url":
        "ssh://gitserver/project"
      }
    ]
  }
}'

```

```

    },
    {
      "revision":
"ef9f8fe661992f48c3539c05fc5c69c61918fa17",
      "file":
"commit1_file2.txt",
      "repository_url":
"ssh://gitserver/project"
    },
    {
      "revision":
"fb849a2440dda438f4c6ab25f8c3266ed82d8797",
      "file":
"commit2_file1.txt",
      "repository_url":
"ssh://othergitserver/project"
    }
  ],
  "reviewers": [
    { "username": "adent",
      "link": "http://reviewserver/user/
adent" },
    { "username": "fprefect",
      "link": "http://reviewserver/user/
fprefect" }
  ]
}
}'

# Event loop
EventMachine.run do
  connection =
  AMQP.connect('amqp://
guest:guest@example-mq')

  # Set up our RabbitMQ
  information
  channel =
  AMQP::Channel.new(connection)
  queue =
  channel.queue(Queue, :auto_delete
=> false, durable: true)
  exchange = channel.direct('')

  # Publish the activity and
  exit the loop
  exchange.publish
  REVIEW, :routing_key => queue.name
do
  connection.disconnect
  {EventMachine.stop}
end
end
end

```

## TeamForge EventQ Work Item API

### Standards

Integrating TeamForge EventQ with the Work Item systems requires for the adapters to be able to send the configuration details of the underlying systems (trackers configuration etc. (see "TODO put a link on configuring work item sources")) as well as the actual data (work items) generated by these systems. These are two distinct types of messages that use different formats.

## TeamForge EventQ Work Item Configuration API

## Submission parameters

## Submission queue

- **queue\_name**: eventq.work\_items
- **auto\_delete**: false
- **durable**: true

## Top-level

Parameter	Required/ Optional	Description
<b>api_version</b>	Required	A string that matches the API version this document is for.
<b>source_association_key</b>	Required	A key generated by TeamForge EventQ that links incoming work items with the appropriate source server.  Example: "source_association_key": "6t5qM5AuLLWLkmqFJeKp"
<b>workitem_settings</b>	Required	Tracker configuration details.

## Objects

*workitem\_settings* - The following fields should be in this *workitem\_settings* object.

Field	Required/ Optional	Description
<b>event_time</b>	Required	A string containing a timestamp in UTC timezone and RFC 3339 format.
<b>url</b>	Required	A string with the browse URL for the system. EventQ will append an individual work item's id to that url to link back to the work item
<b>tracker</b>	Required	Tracker configuration information. A JSON object describing the

Field	Required/ Optional	Description
-------	-----------------------	-------------

structure of the tracker.

#### Accepted values

- **name** (Required) - A name for the tracker.
- **id** (Required) - A string containing a unique (on the reporting system) identifier for the tracker.
- **key** (Optional) - A string containing the tracker key (usually used to prefix the work-item ids).
- **regex** (Required) - A string containing a regular expression that EventQ will use to associate work-items with commits using the commit message.
- **icon** (Optional) - A string containing a URL for this tracker's icon.
- **statuses** (Required) - An array of status objects associated with the tracker.
  - **id** (Required) - A string containing a unique (on the underlying system) identifier for the status.
  - **name** (Required) - A string containing the name of the status.

## Examples

The following code examples are Copyright 2015 CollabNet, Inc., licensed under the Apache License, Version 2.0 (the "License"); you may not use this code except in compliance with the License. You may obtain a copy of the License at [<http://www.apache.org/licenses/LICENSE-2.0>]

### Sample Submission JSON

```
{
  "api_version": "1",
  "source_association_key" :
  "150e8951-19e1-4066-8b74-601026dd6cde",
  "workitem_settings" : {
    "url": "http://
www.example.com/jira",
    "event_time":
    "2012-10-02T17:15:32.320Z",
    "tracker": {
      "name": "EventQ/
Stories",
```

```

      "id":
"01234567abcdef01234567",
      "key": "ORC",
      "regex": "ORC-\d
+",
      "icon": "http://
www.example.com/jira/browse/
bug.jpg",
      "statuses": [
        {"name":
"Open", "id": "abcdef"},
        {"name": "In
progress", "id": "123abc"},
        {"name":
"Closed", "id": "456123"}
      ]
    }
  }
}

```

### Ruby Submission Example

You will need to change the following example before you can run it. Set the Rabbit MQ hostname and the association key, and paste a JSON block following the build format specified on this page.

```

require 'amqp'

QUEUE = 'eventq.work_items'
WORK_ITEM_CONFIGURATION = '
{
  "api_version": "1",
  "source_association_key" :
"150e8951-19e1-4066-8b74-601026dd6cde",
  "workitem_settings" : {
    "url": "http://
www.example.com/jira",
    "event_time":
"2012-10-02T17:15:32.320Z",
    "tracker": {
      "name":
"EventQ/Stories",
      "id":
"01234567abcdef01234567",
      "key": "ORC",
      "regex":
"ORC-\d+",
      "icon":
"http://www.example.com/jira/
browse/bug.jpg",
      "statuses": [
        {"name": "Open", "id": "abcdef"},
        {"name": "In progress", "id":
"123abc"},
        {"name": "Closed", "id": "456123"}
      ]
    }
  }
}

```

```

    }'
    # Event loop
    EventMachine.run do
      connection =
        AMQP.connect('amqp://
        guest:guest@example-mq')

      # Set up our RabbitMQ
      information
      channel =
        AMQP::Channel.new(connection)
      queue =
        channel.queue(Queue, :auto_delete
        => false, durable: true)
      exchange =
        channel.direct('')

      # Publish the work item
      configuration and exit the loop
      exchange.publish
      WORK_ITEM_CONFIGURATION, :routing_key
      => queue.name do
        connection.disconnect
      {EventMachine.stop}
      end
    end
  end
end

```

### TeamForge EventQ Work Item Message API

#### Submission parameters

##### Submission queue

- **queue\_name:** eventq.work\_items
- **auto\_delete:** false
- **durable:** true

##### Top-level

Parameter	Required/ Optional	Description
<b>api_version</b>	Required	A string that matches the API version this document is for.
<b>source_association_key</b>	Required	A key generated by TeamForge EventQ that links incoming work items with the appropriate source server.  Example: "source_association_key": "6t5qM5AuLLWLkmqFJeKp"
<b>workitem_settings</b>	Required	The work item data.

##### Objects

*work\_item* - The following fields should be in this work item object.

Field	Required/ Optional	Description
<b>created_by</b>	Optional	A string containing the username of the person who submitted the work item.
<b>updated_by</b>	Optional	A string containing the username of the person performing the update.
<b>id</b>	Required	A string to uniquely identify the work item.
<b>event_time</b>	Required	A string containing a timestamp in UTC timezone and RFC 3339 format.
<b>summary</b>	Required	A string to summarize the work item.
<b>description</b>	Optional	A longer string that can be used to further describe the work item.
<b>status_id</b>	Required	A string to uniquely identify the status. This value must match the identifier within the work item configuration message.
<b>tracker_id</b>	Required	A string to uniquely identify the tracker. This value must match the identifier within the work item configuration message.
<b>closed</b>	Optional	A boolean for the open or closed



Field	Required/ Optional	Description
		state of the work item. True if the work item is closed.
<b>deleted</b>	Optional	A boolean for the deleted state of the work item. True if the work item has been deleted.
<b>priority</b>	Optional	A string containing the priority of the work item.
<b>assigned_to</b>	Optional	A string containing the username of the person who is assigned to the work item.
<b>creation_time</b>	Optional	A string containing a timestamp in UTC timezone and RFC 3339 format. The time the item was first created in the source system.
<b>tags</b>	Optional	An array of strings containing the tags, or labels, for this work item.

## Examples

The following code examples are Copyright 2015 CollabNet, Inc., licensed under the Apache License, Version 2.0 (the "License"); you may not use this code except in compliance with the License. You may obtain a copy of the License at [<http://www.apache.org/licenses/LICENSE-2.0>]

### Sample Submission JSON

```
{
  "api_version": "1",
  "source_association_key" :
  "150e8951-19e1-4066-8b74-601026dd6cde",
  "work_item": {
    "created_by": "admin",
    "updated_by": "username",
    "summary": "title",
    "id": "KDA-59",
```

```

        "description":
"description",
        "status_id": "10001",
        "event_time":
"2014-03-12T20:47:41.295Z",
        "closed": false,
        "deleted": false,
        "tracker_id": "10200-2",
        "priority": "Major",
        "assigned_to": "admin",
        "creation_time":
"2014-03-29T00:26:20.091Z",
        "tags": [
            "Tag1",
            "Tag2",
            "Tag3"
        ]
    }
}

```

### Ruby Submission Example

You will need to change the following example before you can run it. Set the Rabbit MQ hostname and the association key, and paste a JSON block following the build format specified on this page.

```

require 'amqp'

QUEUE = 'eventq.work_items'
WORK_ITEM = '
{
  "api_version": "1",
  "source_association_key" :
"150e8951-19e1-4066-8b74-601026dd6cde",
  "work_item": {
    "created_by": "admin",
    "updated_by": "admin",
    "summary": "title",
    "id": "KDA-59",
    "description":
"description",
    "status_id": "10001",
    "event_time":
"2014-03-12T20:47:41.295Z",
    "closed": false,
    "deleted": false,
    "tracker_id": "10200-2",
    "priority": "Major",
    "assigned_to": "admin",
    "creation_time":
"2014-03-29T00:26:20.091Z",
    "tags": [
      "Tag1",
      "Tag2",
      "Tag3"
    ]
  }
}'
# Event loop
EventMachine.run do
  connection =
  AMQP.connect('amqp://
  guest:guest@example-mq')

```

```

        # Set up our RabbitMQ
information
        channel =
AMQP::Channel.new(connection)
        queue =
channel.queue(Queue, :auto_delete
=> false, durable: true)
        exchange =
channel.direct('')

        # Publish the work item and
exit the loop
        exchange.publish
WORK_ITEM, :routing_key =>
queue.name do
            connection.disconnect
{EventMachine.stop}
        end
    end
end

```

## Extensible Data Sources (XDS) walk-through

This walk-through is intended to provide a detailed step-by-step implementation example of the Extensible Data Sources (or XDS) feature. Please read the parent overview document first: [Extending TeamForge EventQ](#) on page 884

Setting the stage: EventQ does not ship with a stock activity API for deployment tools, so our fictional example lifecycle includes a process to deploy software into various environments. EventQ users would like to see these deployment activities bubble up to activity streams and any deployment associations in the context of other activities. So let's define an XDS Schema modeling the "deployment" domain and then send activities to the newly registered XDS Activity API for deploy.

**Define the Schema Attributes** - The first step is craft an XDS Schema to adequately describe our deployment domain space and the events we expect to send to EventQ. The submission of the data type schema is itself governed by an API, which is documented at [XDS Schema API](#).

### 1. API version

The version number of the XDS Schema API (not your new schema's version number).

### 2. Source association key

This value is obtained by creating a new "XDS Source" in EventQ's web client. Supplying a source association key is optional during schema registration. If provided, the source will be associated with the schema when the schema is received. Otherwise, the XDS source will be associated to a registered schema when the first activity message is received. There is no "preferred" option; this flexibility exists purely to suit various types of integration models.

### 3. Schema name

Should be a human-friendly identifier of the data source.

### 4. Schema ID

Takes a UUID as the value and the value should be unique per schema. If UUID generation is not possible, characters are limited to alphanumerics and the dash symbol.

### 5. Schema version

Allows for the same schema to be updated without sending a separate schema ID. This is an integer field, best practice is to start at 1 and increment.

### 6. Event time

The time at which your newly crafted schema is submitted to EventQ. Requires UTC RFC3339 format.

The schema ID is user-defined; we suggest a long "UUID" style string. The schema name is also user-defined. In our deploy example, we can add a schema ID and name as follows:

```
// Example Deploy XDS Schema
{
  "api_version": "1",
  "source_association_key" : "8578c900-f8df-0131-84ff-3c07547a48b0",
  "custom_schema": {
    "name": "Faux Schema",
    "schema_id": "fc358dd0-11dc-0132-b9fe-3c07547a48b0",
    "schema_version": 2,
    "event_time": "2015-02-02T17:15:32.320Z",
  }
}
```

Note, the source association key value was obtained from our instance of EventQ. It is an attribute of the parent XDS Collection with which you intend to register the new schema.

**Defining the fields expected by your schema** - Since you are defining an API at this stage, the next step is to define the list of fields expected by your API. Fields are the attributes you wish to supply with each event message sent by an adapter. We'll address optional vs. required fields in the next section, but for now we want a comprehensive list of all fields represented in your new activity message format. For instance, in our deploy example you may want to supply deploy environment, time, duration of deploy activities, etc.

The XDS Activity API categorizes fields into three classes: mandatory fields, stock fields, and custom fields.

"Mandatory fields" are required in all registered schemas and thus must be present in all activity messages. "Stock fields" are a set of non-mandatory fields that have special relevance to the EventQ application, such as appearing in specific places in the user interface. "Custom fields" are truly arbitrary and may be used to capture desired data not encompassed by mandatory or stock fields.

**7. Mandatory fields** - all of these fields are required in every XDS schema.

a) schema\_id

Activity messages will need to send this to identify themselves as governed by a specific schema.

b) schema\_version

Should match the same schema\_version as sent by the schema. This makes sure that the message is displayed with the correct version of the schema.

c) event\_time

The UTC RFC3339 time of the event.

d) remote\_id

Represents the source tool's identifier for the activity. For instance, if we were modeling Subversion commits, the remote\_id would be the revision number. Messages with the same remote\_id are treated as part of the same activity and are aggregated as "sub-events".

**8. Stock fields** - Optional fields with special relevance in the EventQ application.

a) associated\_remote\_id

If your activity has an association this field is used to specify one or more IDs for purposes of association.

b) created\_by

The user who instantiated the activity.

c) summary

The primary heading or title of the activity.

d) description

A detailed description of the activity.

e) status\_name

The current status of the activity, corresponds to the "status\_type" field.

f) status\_type

The current status class, either "open", "success", or "fail" each of which maps to status icons displayed in the user interface.

g) link\_to

Creates a link back to the activity on the source system. The link\_to field can take an array with a link label along with the URL to customize the display text in the user interface. The default label is "See More".

Following our example, one could display "Deployment Details" in the user interface in place of the default "See More".

Let's return to our deployment example and see how fields are defined in the Data Type Schema:

```
// Example Deploy XDS Schema
{
  "api_version": "1",
  "source_association_key" : "8578c900-f8df-0131-84ff-3c07547a48b0",
  "custom_schema": {
    "name": "Faux Schema",
    "schema_id": "fc358dd0-11dc-0132-b9fe-3c07547a48b0",
    "schema_version": 2,
    "event_time": "2015-02-02T17:15:32.320Z",
    "fields": ["schema_id", "schema_version", "event_time",
"remote_id",
"associated_remote_id", "summary", "status_name", "status_type",
"link_to", "rule", "deploy_url", "deploy_message"],
  }
}
```

Note that we've included the four mandatory fields, a couple of "stock fields", and some custom fields modeling our "deploy" specific domain data.

**Requiring specific fields** - You may impose required fields. Remember that some fields are always required, so this section is not optional for that reason. But you can extend the set of required fields beyond those required by the API.

## 9. required

Pass in an array of required fields. Any activity message without all of these fields will be discarded.

Let's return to our deploy example and see how fields are defined in the Data Type Schema:

```
// Example Deploy XDS Schema
{
  "api_version": "1",
  "source_association_key" : "8578c900-f8df-0131-84ff-3c07547a48b0",
  "custom_schema": {
    "name": "Faux Schema",
    "schema_id": "fc358dd0-11dc-0132-b9fe-3c07547a48b0",
    "schema_version": 2,
    "event_time": "2015-02-02T17:15:32.320Z",
    "fields": ["schema_id", "schema_version", "event_time", "remote_id",
"associated_remote_id", "summary", "status_name", "status_type",
"link_to", "rule", "deploy_url", "deploy_message"],
    "required": ["schema_id", "schema_version", "event_time",
"remote_id", "rule"],
  }
}
```

The four fields required by the API are represented, along with one additional field we added above named "rule". Any activity message that fails to supply all five of these fields will be discarded.

**Select "important" fields that result in activity stream posts** - Not every event that triggers an activity message is necessarily worthy of a post in the activity stream. List only those fields that you deem worth of an activity stream update in the "important" field.

## 10. important

Array of field names. Only changes to these fields will trigger a new entry in the activity stream.

In this case, we only want changes to the "summary" or "status\_name" field to trigger activity stream posts:

```
// Example Deploy XDS Schema
{
  "api_version": "1",
  "source_association_key" : "8578c900-f8df-0131-84ff-3c07547a48b0",
  "custom_schema": {
    "name": "Faux Schema",
    "schema_id": "fc358dd0-11dc-0132-b9fe-3c07547a48b0",
    "schema_version": 2,
    "event_time": "2015-02-02T17:15:32.320Z",
    "fields": ["schema_id", "schema_version", "event_time", "remote_id",
      "associated_remote_id", "summary", "status_name", "status_type",
      "link_to", "rule", "deploy_url", "deploy_message"],
    "required": ["schema_id", "schema_version", "event_time",
      "remote_id", "rule"],
    "important": ["summary", "status_name"],
  }
}
```

**Parse certain fields for TeamForge object references** - Say you have a field that often contains a reference to a tracker artifact and you want EventQ to automatically convert those into links. An association will also be created if EventQ is aware of the referenced object.

In addition, any URL reference will be converted to a HTML link in the EventQ user interface. Say, [www.example.com](http://www.example.com) is referenced in a "parse" field, then it will be converted to a HTML link, e.g. [www.example.com](http://www.example.com)

#### 11. parse

Array of field names. Converts any TeamForge object references to links, converts URL references to links.

```
// Example Deploy XDS Schema
{
  "api_version": "1",
  "source_association_key" : "8578c900-f8df-0131-84ff-3c07547a48b0",
  "custom_schema": {
    "name": "Faux Schema",
    "schema_id": "fc358dd0-11dc-0132-b9fe-3c07547a48b0",
    "schema_version": 2,
    "event_time": "2015-02-02T17:15:32.320Z",
    "fields": ["schema_id", "schema_version", "event_time", "remote_id",
      "associated_remote_id", "summary", "status_name", "status_type",
      "link_to", "rule", "deploy_url", "deploy_message"],
    "required": ["schema_id", "schema_version", "event_time",
      "remote_id", "rule"],
    "important": ["summary", "status_name"],
    "parse": ["deploy_message"],
  }
}
```

By adding "deploy\_message" to the "parsed" field, any CTF object will be linked and possibly associated. Any web addresses referenced will also be converted to a full HTML link.

**Register your schema with EventQ** - Now that your new schema is formulated, register your schema with your EventQ instance using either:

- `eventq.custom` message queue, or
- **XDS Schema HTTP API** (see `/api/1/docs/http_api_xds_schema`)

There is no preferred method for schema registration; pick the method that works best for your integration scenario.

Registered schemas may be used site-wide. In our deploy example, perhaps we've written an extension to our deploy product and have registered its schema. The same schema may be reused by multiple sources (e.g., multiple instances of the deploy product) across the site.

Note also that successive submissions of the same schema ID and version will be ignored and it is a safe practice to send the same schema multiple times if it is programmatically convenient to do so.

To modify a schema, increment the `schema_version` value and re-register the schema with EventQ using the methods described above.

## 12. Success

You can use the **XDS Schema HTTP API** (see [/api/1/docs/http\\_api\\_xds\\_schema](/api/1/docs/http_api_xds_schema)) to inspect registered schemas and their versions.

## 13. Troubleshooting

Check the server log "custom\_queue.log" to see if the schema was rejected and why it was rejected.

**Send Activity messages** - Now that you've defined an API for your tool/domain you can start supplying activity messages to EventQ that coincide with your new schema. For a technical reference on crafting activity messages, see the [XDS Schema API](#) documentation.

## 14. Submit XDS Activity messages to eventq.custom queue

Note that the activity stream will only reflect new entries when an "important" field is modified or set in an activity message.

```
// Example Deploy XDS Activity Message
{
  "api_version": "1",
  "source_association_key" : "8578c900-f8df-0131-84ff-3c07547a48b0",
  "custom_data": {
    //required fields
    "name": "Faux Schema",
    "schema_id": "fc358dd0-11dc-0132-b9fe-3c07547a48b0",
    "schema_version": 2,
    "event_time": "2015-02-02T17:15:32.320Z",
    "remote_id": "deploy-1234",
    "rule": "1 comes before 2",
    // stock fields
    "summary": "Deploy at 4:52pm on a Friday",
    "deploy_message": "Deploying to fix [artf12345] and [artf12346] as
critical bugs.",
    "associated_remote_id": "b4c8e50c908ade5c6f053e1ef6422f88",
    "link_to": ["Deploy Details", "http://www.example.com/deploy/
deploy-1234"],
    "status_type": "open",
    // other custom fields
    "environment": "Dev Stage 1"
  }
}
```

**Optional Features** - The XDS API offers some optional features to help pretty things up and make life a bit easier.

## 15. Interpolate TeamForge URLs

It may be desirable to create URL references to TeamForge pages or functions (like SSO authentication). There are two interpolation functions in XDS:

Reference TeamForge base URLs: `%CTF_BASE_URL%`

Reference TeamForge "Go" URLs: `%CTF_GO_URL%`

In our example, we may wish to reference a TeamForge Go URL like so:

```
// Example Deploy XDS Activity Message
...
"summary": "Deploy at 4:52pm on a Friday. See %CTF_GO_URL%/page1001",
```

## 16. Control field visibility

Say you wanted to load specific data into EventQ's database for reporting purposes, but you don't want to expose all of those fields in the user interface. The solution is to add those fields to the hidden array in the XDS Schema message. The hidden array takes only custom fields (optional, non-stock). Custom fields are visible by default.

Continuing our example, the newly introduced "environment" and "authorized\_by" custom fields will be stored in the EventQ database, but not exposed in the user interface:

```
// Example Deploy XDS Schema
{
  "api_version": "1",
  "source_association_key" : "8578c900-f8df-0131-84ff-3c07547a48b0",
  "custom_schema": {
    "name": "Faux Schema",
    "schema_id": "fc358dd0-11dc-0132-b9fe-3c07547a48b0",
    "schema_version": 2,
    "event_time": "2015-02-02T17:15:32.320Z",
    "fields": ["schema_id", "schema_version", "event_time", "remote_id",
"associated_remote_id",
    "summary", "status_name", "status_type", "rule", "deploy_url",
"deploy_message",
    "environment", "authorized_by"],
    "required": ["schema_id", "schema_version", "event_time",
"remote_id", "rule"],
    "important": ["summary", "status_name"],
    "parse": ["deploy_message"],
    "link_to": ["deploy_url"],
    "hidden": ["environment", "authorized_by"]
  }
}
```

## TeamForge EventQ XDS Schema API

### Submission parameters

#### Submission queue

- **queue\_name**: eventq.custom
- **auto\_delete**: false
- **durable**: true

#### Top-level

Parameter	Required/ Optional	Description
<b>api_version</b>	Required	A string that matches the API version this document is for.
<b>source_association_key</b>	Optional	A key generated by TeamForge EventQ that links incoming work items with the appropriate source server. (Optional)

Example: "source\_association\_key":  
"kLM56uWtKAqpeLLm5qJF"



**Objects**

Parameter	Required/ Optional	Description
<i>custom_schema</i>	Required	a schema for the XDS events.

*custom\_schema* - The fields specified below define the schema that activities will conform to.

Field	Required/ Optional	Description
<b>name</b>	Required	A visible name that will allow pipeline-creators to recognize this schema when populating a pipeline.
<b>schema_id</b>	Required	A unique id defining this schema. Every activity sent by a single adapter should reference the same schema id. If this changes, the data source is considered a new data source, and data will not populate into pipelines that listened for the old one. If the schema changes, you can use the version to show that it's intended to be parsed differently, but continue to deliver messages to existing pipelines.  Example: "schema_id": "fc358dd0-11dc-0132-b9fe-3c07547a48b0"
<b>schema_version</b>	Required	A string to uniquely identify the work item.  Example: "schema_version": 1

Field	Required/ Optional	Description
<b>event_time</b>	Required	A string containing a timestamp in UTC timezone and RFC 3339 format.  Example: "event_time": "2012-10-02T17:15:32.320Z"
<b>fields</b>	Required	A string to summarize the work item.  Example: "fields": ["schema_id", "schema_version", "event_time", "remote_id", "assigned_to", "hours_played", "severity", "canary"]
<b>required</b>	Required	A longer string that can be used to further describe the work item.  Example: "required": ["schema_id", "schema_version", "event_time", "remote_id"]
<b>important</b>	Optional	A string to uniquely identify the status. This value must match the identifier within the work item configuration message.  Example: "important": ["assigned_to", "severity"]
<b>parsed</b>	Optional	A string to uniquely identify the tracker. This value must match the identifier within the work item configuration message.  Example: "parsed": ["description"]
<b>hidden</b>	Optional	A boolean for the open or closed state of the work

Field	Required/ Optional	Description
		item. True if the work item is closed.
		Example: "hidden": ["comment"]

## Examples

The following code examples are Copyright 2015 CollabNet, Inc., licensed under the Apache License, Version 2.0 (the "License"); you may not use this code except in compliance with the License. You may obtain a copy of the License at [<http://www.apache.org/licenses/LICENSE-2.0>]

### Sample Submission JSON

```
{
  "api_version": "1",
  "source_association_key" :
  "8578c900-
f8df-0131-84ff-3c07547a48b0",
  "custom_schema": {
    "name": "Faux Schema",
    "schema_id":
    "fc358dd0-11dc-0132-
b9fe-3c07547a48b0",
    "schema_version": 2,
    "event_time":
    "2012-10-02T17:15:32.320Z",
    "fields": ["schema_id",
"schema_version",
"event_time", "remote_id",
"associated_remote_id",
"title", "rule", "deploy_url",
"deploy_message", "comment"],
    "required": ["schema_id",
"schema_version", "event_time",
"remote_id", "rule"],
    "important": ["rule"],
    "parsed": ["deploy_message"],
    "hidden": ["comment"]
  }
}
```

## TeamForge EventQ XDS Activity API

### Submission parameters

#### Submission queue

- **queue\_name**: eventq.custom
- **auto\_delete**: false
- **durable**: true

#### Top-level

Parameter	Required/ Optional	Description
<b>api_version</b>	Required	A string that matches the API version this document is for.

Parameter	Required/ Optional	Description
<b>source_association_key</b>	Required	A key generated by TeamForge EventQ that links incoming work items with the appropriate source server.  Example: "source_association_key": "kLM56uWtKAqpeLLm5qJF"
<b>custom_data</b>	Required	The object containing custom activity data conforming to the referenced schema/version.

## Objects

*custom\_data* - The fields specified below conform to the schema referenced, and provide activity information to EventQ:

Field	Required/ Optional	Description
<b>schema_id</b>	Required	The identifier of the XDS schema this message conforms to. The schema needs to be sent before an activity referencing that schema is sent.
<b>schema_version</b>	Required	The version of the referenced schema this message conforms to.
<b>event_time</b>	Required	The identifier of this activity; this is used to collapse related activities together.
<b>remote_id</b>	Required	A string to summarize the work item.
<b>Stock Fields</b>	Optional	The following field names are treated specially by EventQ, allowing a developer to show

Field	Required/ Optional	Description
		<p>specific data in our UI where it can fit with our display patterns. All the fields are optional.</p> <ul style="list-style-type: none"> <li>• <b>associated_remote_id</b> - Specify the singular ID or array of IDs that your event should associate with, these IDs will be searched for in the source chosen in the pipeline custom source and shown in the Association section of the Details page.</li> <li>• <b>created_by</b> - The username or email address of the person who took the action. This is shown in the Details page.</li> <li>• <b>updated_by</b> - The username or email address of the person who took an action. This is shown in the Details page.</li> <li>• <b>summary</b> - The main title of the event. Shown in bold at the top of the Details page and also in the Activity Stream.</li> <li>• <b>description</b> - Detailed description of the event. Shown in full in the Details page, a truncated version is shown in the Activity Stream.</li> <li>• <b>status_name</b> - The name of the current status. This is a text field that is displayed in the Details page and Activity Stream. The value should be representational of the current state of the activity and match with the</li> <li>• <b>status_type</b>. status_type - The type of status icon to show, containing values of "<b>open</b>", "<b>success</b>", "<b>failed</b>", each of which maps to a status icon displayed in the user interface.</li> <li>• <b>link_to</b> - The link back to the event on the source system. It take an array or single string. Array takes a URL along with the link label to customize the display text in the user interface. If only a URL string is passed without the label, the default label is "See More".</li> </ul>

Field	Required/ Optional	Description
Custom Fields	Optional	These custom fields can only contain String data and will be displayed in a key value list in the Details page.

## Examples

The following code examples are Copyright 2015 CollabNet, Inc., licensed under the Apache License, Version 2.0 (the "License"); you may not use this code except in compliance with the License. You may obtain a copy of the License at [<http://www.apache.org/licenses/LICENSE-2.0>]

### Sample Submission JSON 1: A deploy example with string value for link\_to field

```
{
  "api_version": "1",
  "source_association_key" :
  "8578c900-f8df-0131-84ff-3c07547a48b0",
  "custom_data": {
    //required fields
    "schema_id":
    "fc358dd0-11dc-0132-b9fe-3c07547a48b0",
    "schema_version": 2,
    "event_time":
    "2012-10-02T17:15:32.320Z",
    "remote_id": "deploy-1234",
    // stock fields
    "created_by": "Build-bot",
    "updated_by": "Deploy-bot",
    "summary": "Deploy at 4:52pm
on a Friday",
    "description": "Deploying
to fix [FOO-245] and [BAR-900] as
critical bugs.",
    "associated_remote_id":
    "b4c8e50c908ade5c6f053e1ef6422f88",
    "link_to": "http://
www.example.com/deploy/
b4c8e50c908ade5c6f053e1ef6422f88",
    "status_name": "Deploying",
    "status_type": "open",
    // custom fields
    "environment": "production",
    "branch": "master"
  }
}
```

### Sample Submission JSON 2: A build example with array value for the link\_to field

```
{
  "api_version": "1",
  "source_association_key" :
  "1ebbbef0-2025-0132-8f68-3c15c2d51622",
  "custom_data": {
    //required fields
```

```

        "schema_id":
"5193ca40-2025-0132-8f68-3c15c2d51622",
        "schema_version": 1,
        "event_time":
"2012-10-02T17:15:32.320Z",
        "remote_id": "build-1234",
        // stock fields
        "created_by": "Build-bot",
        "updated_by": "Deploy-bot",
        "summary": "Build
successful!",
        "description": "Build
containing commits 109 through
112",
        "associated_remote_id":
["109", "110", "111", "112"],
        "link_to": ["http://
www.example.com/build/build-1234",
"Build-1234"],
        "status_name": "Success",
        "status_type": "success",
        // custom fields
        "tests_passed": "1400",
        "tests_failed": "0",
        "tests_pending": "23"
    }
}

```

## TeamForge EventQ Reporting API

### Overview

- TeamForge EventQ Reporting API provides a means to query activity data in TeamForge EventQ. This API is for data retrieval only.
- Submission URL: **https://<eventq\_hostname>/orc/api/<version>/reporting**
- Method: **POST (only)**
- Authentication: Requires a valid TeamForge session id.
- Authorization: Requires the TeamForge EventQ "EventQ READ", "REPORTING API" permission, or site admin permissions.

### Convention

API requests are constructed from two building-blocks: commands with parameter(s) that can be chained one after another, followed by a single return type with its parameter(s). As such, API requests take the form:


```

command parameter(s)
command parameter(s)
command parameter(s) ...
return type parameter(s)

```

- Each chained command is treated as an "AND" query filter
- Use separate queries to approximate "OR" requests
- Requests are limited to a single return type which must be supplied after the last command
- A query itself is a String constructed by chaining commands each separated from the next (or return) by a new line '\n'.
  - This '\n' is very important since the parser looks for this when tokenizing the query

## Authentication and Authorization

 **Note:** In order to use the reporting api, your reporting client must have access to the endpoint of your TeamForge instance.

When sending query requests to the reporting api, a custom header named 'X-EventQ-Session' must be added to each http post request, the value of which is the 'session id' returned from a successful login call to the TeamForge endpoint. The TeamForge endpoint is found here: [http\(s\)://TeamForgeHost/ce-soap60/services/CollabNet](http(s)://TeamForgeHost/ce-soap60/services/CollabNet)

To gain access to the Reporting API the user must have one either of the below permissions set in TeamForge:

- TeamForge SITE ADMINISTRATOR
  - This will provided access to all active projects in EventQ
- REPORTING API EventQ Permissions on at least one project
  - This limits the projects to query against to only those which the user has REPORTING API permissions
  - If the user is both SITE ADMIN and REPORTING API, SITE ADMIN will be used. There is no way to limit a SITE ADMIN's access

See the Examples below for a simple Java client.

## Commands and Parameters

Command	Description
<code>restrict_to (string)</code>	Filter query results for certain classes of activity objects. <b>Example:</b> <code>restrict_to reviews</code> <b>Accepted values:</b> builds, commits, reviews, work_items, custom.
<code>associated_to (array)</code>	Takes one or more object IDs and returns a list of activity summaries with associations to the object IDs specified. <b>Example:</b> <code>associated_to ["artf1002", "artf1002", "artf1003"]</code>
<code>associations (none)</code>	Used in command chains to modify query results to return associations. No parameters. <b>Example:</b> <code>associations</code>
<code>in_project (array)</code>	Restrict query results to activities in a given project. Takes array of project IDs as a parameter. <b>Example:</b> <code>in_project ["proj1123", "proj1134"]</code>
<code>sort_by (array)</code>	Sort the result set by the fields of the objects which are supplied as a parameter, ascending or descending order. <b>Example:</b> <code>sort_by ["status", "asc"]</code> <b>Accepted values:</b> any object field; see MQ API docs for object descriptions or XDS schema for custom objects. For instance, builds may be sorted by <code>remote_id</code> , <code>status</code> , <code>test_results</code> , <code>duration</code> , <code>created_by</code> , or <code>time_created</code> . <b>Accepted values:</b> asc, desc
<code>has_any (string)</code>	Filter query results to those activities that have matching class associations. <b>Example:</b> <code>has_any "builds"</code>



Command	Description
<code>has_none (string)</code>	<p><b>Accepted values:</b> builds, commits, review, work_items, custom.</p> <p>Filter query results to those activities that do not have matching class associations.</p> <p><b>Example:</b> <code>has_none "reviews"</code></p> <p><b>Accepted values:</b> builds, commits, review, work_items, custom.</p>
<code>limit (integer)</code>	<p>Limit query results to the specified integer for pagination of results and performance. May be used in conjunction with offset.</p> <p><b>Example:</b> <code>limit 100</code></p> <p><b>Accepted values:</b> positive integers.</p>
<code>offset (integer)</code>	<p>Offset query results to specified integer for pagination of results and performance. May be used in conjunction with limit.</p> <p><b>Example:</b> <code>offset 10</code></p> <p><b>Accepted values:</b> positive integers.</p>
<code>before (date/time)</code>	<p>Filter results set to those activity summaries with events that occurred prior to the supplied date and time. May be used in conjunction with after.</p> <ul style="list-style-type: none"> <li>• <b>Example:</b> <code>before 2014-10-02T17:15:45.320Z</code></li> <li>• <b>Example:</b> <code>before 2.weeks.ago</code></li> </ul> <p><b>Accepted values:</b> RFC 3339 format date times in UTC and Ruby-style date references (minutes, hours, days, weeks, months, years).</p>
<code>after (date/time)</code>	<p>Filter results set to those activity summaries with events that occurred after the supplied date and time. May be used in conjunction with before.</p> <ul style="list-style-type: none"> <li>• <b>Example:</b> <code>after 2014-11-02T12:45:45.320Z</code></li> <li>• <b>Example:</b> <code>before 1.day.ago</code></li> </ul> <p><b>Accepted values:</b> RFC 3339 format date times in UTC and Ruby-style date references (minutes, hours, days, weeks, months, years).</p>
<code>tagged_as (array)</code>	<p>Filter results set to those activities from sources labeled with the supplied tag. Takes multiple tags to create an OR condition. Two or more <code>tagged_as</code> commands may be chained which constructs an AND condition, returning only those activities with tag A and tag B.</p> <p><b>Example:</b> <code>tagged_as ["deploy","binary repository"]</code>  <code>tagged_as ["production"]</code></p>

### Return types

Conclude your query with a single "return" method to transform the data and return the results. All results are returned wrapped in a top-level json hash containing the key items, which contains the results of the query. The return

may also contain the optional keys `source_display_data`, containing a hash of display metadata for the associated sources keyed by source id, and `work_item_urls`, containing urls to external work items keyed by id (if applicable).

### Return example

```
{
  "items": [
    {
      "_type": "WorkItemSummary",
      "assigned_to": "dev-bot",

      "associated_activity_summary_ids":
      [

      ],
      "associated_ctf_ids": [

      ],
      "closed": true,
      "comments": [

      ],
      "created_at":
      "2015-04-22T20:54:26.063Z",
      "created_by": "itembot!",
      "creation_time":
      "2014-03-29T00:26:20.091Z",
      "deleted": false,
      "description": "Autogenerated
      workitem description",
      "id":
      "55380a82109a01555d00000c",
      "opaque_id": "6kkd",
      "priority": "mecha-critical",
      "remote_id": "AUTO-6",
      "remote_status_id": "3",
      "settled": true,
      "source_id":
      "5537ed65109a012cec000004",
      "summary": "Release the
      integration",
      "tags": [
        "up",
        "strange"
      ],
      "time_created":
      "2015-04-22T20:54:26.000Z",
      "updated_at":
      "2015-04-23T16:49:19.657Z"
    }
  ],
  "source_display_data": {
    "5537ed65109a012cec000004": {
      "_type": "WorkItemSource",
      "_links": {
        "type_icon": {
          "href": "http://
          ctf.example.com/orc/assets/icons/
          work-item.png",
          "type": "image/png"
        },
        "source_icon": {
          "href": "http://
          external.example.com/images/icons/
          issuetypes/bug.png",

```

```

        "type": "image/png"
      }
    },
    "name": "External Bug
Tracker",
    "background_color": "#f2511b"
  }
},
"work_item_urls": {
  "55380a82109a01555d00000c":
  "http://tracker.example.com/
a_tracker_source/AUTO-6"
}
}

```

## Return Keys

Return Key	Description
<b>items</b> ( <i>array or hash</i> )	Query-dependent results.
<b>source_display_data</b> ( <i>hash</i> )	Display metadata for sources references in results.
<b>work_item_urls</b> ( <i>hash</i> )	URLs for referenced work items. This key is present in the return structure only if 'items' contains WorkItemSummary objects.

## Query Return Data Types

Data item	Description
<b>fields</b> ( <i>array or hash</i> )	<p>Returns json formatted results set limited to the fields supplied as parameters.</p> <p>Example: fields ["time_created", "remote_id"]</p> <p><b>Accepted values:</b> any set of object fields; see MQ API docs for object descriptions or XDS schema for custom objects. For instance, builds may take: remote_id, status, test_results, duration, created_by, or time_created.</p>
<b>field_contains</b> ( <i>nested array</i> )	<p>Returns json formatted results set limited to exact value supplied as parameters.</p> <p><b>Example:</b> field_contains [ "remote_id", ["artf1001", "art1002"] ]</p>

Data item	Description
<b>group_by_field</b> ( <i>string</i> )	<p><b>Accepted values:</b> any object field with matching value set; see MQ API docs for object descriptions or XDS schema for custom objects. For instance, builds may take: remote_id, status, test_results, duration, created_by, or time_created.</p> <p>Returns json formatted results set grouped by the field supplied as parameter.</p> <p><b>Example:</b> group_by_field "status"</p>
<b>count</b> ( <i>none</i> )	<p><b>Accepted values:</b> created_at, remote_id, status</p> <p>Returns a json structure with a count of all matching activities.</p>
<b>timeseries</b> ( <i>date/time</i> )	<p><b>Example:</b> count</p> <p>Useful for charting applications, the output will be a hash, where the keys are the starting points of the time series, and the value is the array of objects that fall into that series. It takes a Ruby-style date reference parameter to specify the results granularity. Requires before and after in the same query.</p> <p><b>Example:</b> timeseries 5.minutes</p> <p><b>Accepted values:</b> Ruby-style date references (minutes, hours, days, weeks, months, years).</p>

## Examples

The following code examples are Copyright 2015 CollabNet, Inc., licensed under the Apache License, Version 2.0 (the "License"); you may not use this code except in compliance with the License. You may obtain a copy of the License at [<http://www.apache.org/licenses/LICENSE-2.0>]

## Use Cases and Query Examples

Use case: "Find the number of all failed builds across all projects"

```
#Query:
  restrict_to builds
  field_contains ["status_type",
["FAILURE" ] ]
  count
```

Use case: "Given a set of artfs as inputs, give me all the associated activities"

```
#Query:
  associated_to: ["artf1234",
"artf1235", "artf1236"]
```

Use case: "Across two specific projects, give me the builds that pass vs fail"

```
#Query:
  in_project ["proj1001",
"proj1002"]
  restrict_to builds
  group_by_field "status_type"
```

Use case: "Given a set of artfs as inputs (say, in a planning folder) give me build stats (total #, pass, fail, etc)"

```
#Query:
  associated_to ["artf1001",
"artf1002", "artf1003"]
  associations
  restrict_to builds
  group_by_field "status_type"
```

Use case: "For a given user, give me a list of commits over X timeline, associated build stats"

```
#Query:
  restrict_to commits
  after 2.weeks.ago
  before 1.week.ago
  field_contains [ "created_by",
["jdoe" ] ]
  associations
  restrict_to builds
  group_by_field "status_type"
```

Use case: "Give me code review stats on a whole project (total #, total reviewed commits, total un-reviewed)"

```
#Query 1 to find un-reviewed commits
  in_project ["proj1706"]
  restrict_to commits
  has_none "reviews"
  count

#Query 2 to find reviewed commits:
  in_project ["proj1706"]
```

```
restrict_to commits
has_any "reviews"
count
```

Use case: "Given a set of artfs, give me associated build results starting 1 week ago ending 5 minutes ago, grouped into 4 hour intervals."

```
#Query:
  associated_to ["artf1001",
"artf1002", "artf1003"]
  associations
  restrict_to builds
  after 1.week.ago
  before 5.minutes.ago
  timeseries 4.hours
```

Use case: "Give me all binary artifact activities grouped by status"

```
#To query XDS sources like binary
artifacts, restrict_to "custom"
source type, then
  #filter to narrow the results
using tagged_as
  #Query:
  restrict_to custom
  tagged_as ["binary artifact"]
  group_by_field "status_type"
```

## Code samples

Create a Java client using Apache Axis & Apache HttpComponents

Get an X-EventQ-Session value from TeamForge.

```
/**
 * @param ctfEndpoint
 * @param ctfUserName
 * @param ctfUserPassword
 * @return sessionId (X-EventQ-
Session)
 */
private static String
getSessionIdFromCTF(String
ctfEndpoint, String ctfUserName,
String ctfUserPassword) {
  try {
    Service service = new Service();
    Call call = (Call)
service.createCall();
    call.setTargetEndpointAddress(new
java.net.URL(ctfEndpoint));
    call.setOperationName(new
QName("http://
schema.open.collab.net/sfee50/
soap60/service", "login"));
    String sessionId = (String)
call.invoke(new Object[]
{ctfUserName, ctfUserPassword});
    return sessionId;
  } catch (Exception e) {
    System.err.println(e.toString());
  }
}
```

```

    return e.toString();
  }
}

```

Post a request to the Reporting endpoint containing the custom header and a query.

```

/**
 * @param url
 * @param token
 * @param query
 * @return queryResults
 * @throws Exception
 */
String sendPostHttpClient(String url, String token, String query)
throws Exception {
    String queryResults;
    CloseableHttpClient client =
    HttpClients.custom().build();

    HttpUriRequest request =
    RequestBuilder.post().setUri(url)

    .setHeader(HttpHeaders.CONTENT_TYPE,
    "application/x-www-form-
    urlencoded")
    .setHeader("X-EventQ-Session",
    token)
    .setEntity(new
    StringEntity(query))
    .build();

    CloseableHttpResponse response =
    client.execute(request);
    queryResults =
    EntityUtils.toString(response.getEntity(),
    "UTF-8");
    return queryResults;
}

```

Build your client using the above methods.

```

public class ReportingClient {

    /* EventQ End Point represents the
    EventQ instance which you will be
    querying against*/
    static String eventqEndPoint =
    "http(s)://EventQHost/orc/api/1/
    reporting/";

    /* teamForgeEndPoint represents
    the soap endpoint for the TeamForge
    instance which your EventQ
    * server is bound to*/
    static String teamForgeEndPoint =
    "http(s)://TeamforgeHost/ce-soap60/
    services/CollabNet";

    /* orcReportingQuery is your
    query in String format. Each LINE
    represents a different command

```

```

    * so in this example every
    distinct command must be separated
    by a newline character '\n'*/
    static String orcReportingQuery
    = "restrict to reviews
    \ngroup_by_field \"status_type\"";

    /* TeamForge username and password
    for the account which the queries
    will be run under. This account
    * must have REPORTING_API
    permissions on at least one project
    in EventQ or must be a SITE ADMIN
    * in Teamforge. If the user is NOT
    a TeamForge SITE ADMIN, only those
    projects for which the user has
    * REPORTING_API permissions will
    be considered in the query.*/
    static String teamForgeUserName =
    "userName";
    static String
    teamForgeUserPassword =
    "userPassword";

    public static void main(String...
    args) {
        EventQReportingClient
        sampleClient = new
        EventQReportingClient();
        String xEventQSession =
        sampleClient.getSessionIdFromCTF(teamForgeEndPoint,
        teamForgeUserName,
        teamForgeUserPassword);
        try {
            String serverResponse =
            sampleClient.sendPostHttpClient(eventqEndPoint,
            xEventQSession,
            orcReportingQuery);

            System.out.println(serverResponse);
        } catch (Exception e) {

            System.out.println(e.toString());
        }

    }
}

```

### Efficient querying with the EventQ Reporting API

Here's a few guidelines for building efficient EventQ Reporting API queries.

#### Cutting vs Building

Unlike a SQL query, which connects separate tables together to form a result, the EventQ Reporting API starts with the full set of all activities, and is cut down to the size of the result by successively applying query restrictions. This difference makes query optimization quite different between SQL and NoSQL.

#### Ordering guidelines

Some query commands will result in larger reductions to the than others. This of course also depends on the range applied to the query. In example: Restricting to a week of activities using “after 1.week.ago” will make a very small



set of activities to work on, compared to “after 1.year.ago” which will include a large number of activities. In general however, the commands here are listed by most efficient to least efficient in scoping the query result:

- **Good:**

```
:before , type : :time
:after , type : :time
:restrict_to , type : :choice , choices : %w(builds commits reviews
work_items custom)
:in_source , type : :json
:in_repository , type : :string
:in_project , type : :json
```

- **Mid:**

```
:associated_to , type : :json
:tagged_as , type : :json
:offset , type : :integer
```

- **Poor:**

```
:associations
:field_contains , type : :json
:sort_by , type : :json
:has_any , type : :string
:has_none , type : :string
```

- **Returning command (must be run last)**

```
:limit , type : :integer
:timeseries , type : :duration
:group_by_field , type : :string
:fields , type : :json
:count
```

## Use cases and examples

```
:in_project ["proj1211"]
after 14.days.ago
restrict_to commits
field_contains ["settled",[true]]
has_any "reviews"
```

## Exceptions in ordering and other exceptional behaviours to note

**Limit:** The limit command is best used at the end of the set. Using it before the end can result in more than the limit of items being returned at the end of the query.

**Tagged\_as:** Is specifically for custom activities.

## TeamForge EventQ HTTP API

### TeamForge EventQ HTTP API for XDS Schema Creation and Viewing

#### Overview

- TeamForge EventQ HTTP API for XDS Schema provides a means to create an XDS schema or view a schema and its versions. The API supports both data retrieval and submission.
- Authentication: requires a valid TeamForge session id, passed with the 'X-EventQ-Session' header field. One way to acquire this session id is to programmatically log into TeamForge using TeamForge SOAP API.

- Authorization: requires the TeamForge EventQ "EventQ CREATE" permission, or site admin permissions for creating schema. No permission required for viewing schema.

### Convention

API requests for viewing and creating schemas are constructed as detailed below.

- **List all schemas**
  - **Request Type:** GET
  - **Request URL:** `https://<eventq_hostname>/api/1/schemas`
  - **Return Format:** Array of JSON objects. Each of the JSON object is a schema.
- List all schemas with a same schema\_id
  - **Request Type:** GET
  - **Request URL:** `https://<eventq_hostname>/api/1/schemas/SCHEMA_ID`
    - 👉 **Note:** SCHEMA\_ID needs to be encoded (escaped).
  - **Requires Parameter:** SCHEMA\_ID
  - **Return Format:** Array of one or more JSON object(s). Each of the JSON object is a schema with the given schema\_id and all the versions.
- Get one schema
  - **Request Type:** GET
  - **Request URL:** `https://<eventq_hostname>/api/1/schemas/SCHEMA_ID/VERSION`
    - 👉 **Note:** SCHEMA\_ID needs to be encoded (escaped).
  - **Requires Parameter:** SCHEMA\_ID, VERSION
  - **Return Format:** Array with one JSON object. The JSON object is a schema with the given schema\_id and given version.
- Create a schema
  - **Request Type:** POST
  - **Request URL:** `https://<eventq_hostname>/api/1/schemas`
  - **Return Format:** An encoded (escaped) URL. This URL can be used to view the schema once it has been processed and saved.

### TeamForge EventQ HTTP API for Tool Creation and Viewing

#### Overview

- TeamForge EventQ HTTP API for Tool provides a means to create an Tool or view a tool. The API supports both data retrieval and submission.
  - Authentication: requires a valid TeamForge session id, passed with the 'X-EventQ-Session' header field. One way to acquire this session id is to programmatically log into TeamForge using TeamForge SOAP API.
  - Authorization: requires the TeamForge EventQ "EventQ CREATE" permission, project admin, or site admin permissions.
- 👉 **Note:** PROJECT\_URL\_NAME is auto generated by TeamForge with project name initially. This should be found in last path of URL at project home page.

### Convention

API requests for viewing and creating tools are constructed as detailed below.

- **List all tools**
  - **Request Type:**
  - **Request URL:** `https://<eventq_hostname>/api/stable/projects/PROJECT_URL_NAME/tools`

- **Return Format:**

```
{
  "tools": [
    {
      "id": "unique_tool_id_1",
      "name": "Sample tool name One",
      "service_name": "Sample service name One",
      "project_id": "project id ",
      "active": true,
      "source_ids": [

    ]
  },
  {
    "id": "unique_tool_id_2",
    "name": "Sample tool name two",
    "service_name": "Sample service name two",
    "project_id": "project id ",
    "active": true,
    "source_ids": [
      "SourceID1",
      "SourceID2"
    ]
  }
 ]
}
```

- **Return Codes**

- **Success** : 200
- **Error** : 404, 503

- **Get one tool**

- **Request Type:** GET
- **Request URL:** [https://<eventq\\_hostname>/api/stable/projects/PROJECT\\_URL\\_NAME/tools/TOOL\\_ID](https://<eventq_hostname>/api/stable/projects/PROJECT_URL_NAME/tools/TOOL_ID)
- **Return Format:**

```
{
  "tool": {
    "id": "TOOL_ID",
    "name": "Sample tool name two",
    "service_name": "Sample service name two",
    "project_id": "project id ",
    "active": true,
    "source_ids": [
      "SourceID1",
      "SourceID2"
    ]
  }
}
```

- **Return Codes**

- **Success** : 200
- **Error** : 404, 503

- **Create a tool**

- **Request Type:** POST
- **Request URL:** [https://<eventq\\_hostname>/api/stable/projects/PROJECT\\_URL\\_NAME/tools](https://<eventq_hostname>/api/stable/projects/PROJECT_URL_NAME/tools)
- **Requires Parameter:** PROJECT\_URL\_NAME

- **Request body:**

```
{
  "tool": {
    "name": "Sample tool name two",
    "source_ids": [
      "SourceID1",
      "SourceID2"
    ]
  }
}
```

- **Return Format:**

```
{
  "tool": {
    "id": "unique_tool_id_x",
    "name": "Sample tool name two",
    "project_id": "project id ",
    "active": true,
    "source_ids": [
      "SourceID1",
      "SourceID2"
    ]
  },
  "errors": {
    "base": [
      "error msg"
    ],
    "attribute_key": [
      "error msg"
    ]
  }
}
```

- **Return Codes**

- **Success** : 200
- **Error** : 404, 503, 422

- **Update a tool**

- **Request Type:** PUT
- **Request URL:** [https://<eventq\\_hostname>/api/stable/projects/PROJECT\\_URL\\_NAME/tools/TOOL\\_ID](https://<eventq_hostname>/api/stable/projects/PROJECT_URL_NAME/tools/TOOL_ID)
- **Requires Parameter:** PROJECT\_URL\_NAME, TOOL\_ID
- **Request body:**

```
{
  "name": "Tool New Name",
  "source_ids": [
    "sourceID3"
  ]
}
```

- **Return Format:**

```
{
  "errors": {
    "base": [
      "error msg"
    ],
    "attribute_key": [
      "error msg"
    ]
  }
}
```

```

    ]
  }
}

```

- **Return Codes**

- **Success** : 200
- **Error** : 404, 503, 422

- **Activate a tool**

- **Request Type:** PUT
- **Request URL:** https://<eventq\_hostname>/api/stable/projects/PROJECT\_URL\_NAME/tools/TOOL\_ID/activate
- **Requires Parameter:** PROJECT\_URL\_NAME, TOOL\_ID
- **Return Format:**

```

{
  "errors": {
    "base": [
      "error msg"
    ],
    "attribute_key": [
      "error msg"
    ]
  }
}

```

- **Return Codes**

- **Success** : 200
- **Error** : 404, 503, 422

- **Deactivate a tool**

- **Request Type:** PUT
- **Request URL:** https://<eventq\_hostname>/api/stable/projects/PROJECT\_URL\_NAME/tools/TOOL\_ID/deactivate
- **Requires Parameter:** PROJECT\_URL\_NAME, TOOL\_ID
- **Return Format:**

```

{
  "errors": {
    "base": [
      "error msg"
    ],
    "attribute_key": [
      "error msg"
    ]
  }
}

```

- **Return Codes**


- **Success** : 200
- **Error** : 404, 503, 422

## TeamForge EventQ HTTP API for Source CRUD

### Overview

- TeamForge EventQ HTTP API for Sources provides a means to create or view a source. The API supports both data retrieval and submission.

- Authentication: requires a valid TeamForge session id, passed with the 'X-EventQ-Session' header field. One way to acquire this session id is to programmatically log into TeamForge using TeamForge SOAP API.
- Authorization: requires the EventQ "EventQ CREATE" permission or TeamForge's project-admin / site-admin permissions.

 **Note:** PROJECT\_URL\_NAME is auto generated by TeamForge with project name initially. This should be the last value in the URL at project home page.

## Convention

API requests for viewing and creating sources are constructed as detailed below.

- **List all sources by Project**

- **Request Type:** GET
- **Request URL:** https://<ctf\_hostname>/api/stable/projects/PROJECT\_URL\_NAME/sources
- **Return Format:**

```
{
  "sources": [
    {
      "id": "5578cd1c0fbb6ad771000001",
      "_type": "Build",
      "tool_id": "552c3736c69d3b3dfb000004",
      "active": true,
      "display_name": "Sample_Build_Source",
      "queue_username": "4d4eaf60-f1f9-0132-4687-185aeb8b25b8",
      "queue_password":
"d70bc954b64e368f1f650bfeaba2e2ff6e5525774ae42f9a959f4ccf2314e9cb",
      "queue_server": "amqp://rabbitmqserver.example.com:5672",
      "source_association_key": "29bf0c90-3b40-0130-ae2d-dddd5893",
      "associated_commit_server_id": "551d8560c69d3b474b000102"
    }
  ]
}
```

- **Return Codes**

- **Success :** 200
- **Error :** 404, 503

- **List all sources by Tool**

- **Request Type:** GET
- **Request URL:** https://<ctf\_hostname>/api/stable/projects/PROJECT\_URL\_NAME/tools/TOOL\_ID/sources
- **Return Format:**

```
{
  "sources": [
    {
      "id": "5578cd1c0fbb6ad771000001",
      "_type": "Build",
      "tool_id": "552c3736c69d3b3dfb000004",
      "active": true,
      "display_name": "Sample_Build_Source",
      "queue_username": "4d4eaf60-f1f9-0132-4687-185aeb8b25b8",
      "queue_password":
"d70bc954b64e368f1f650bfeaba2e2ff6e5525774ae42f9a959f4ccf2314e9cb",
      "queue_server": "amqp://rabbitmqserver.example.com:5672",
      "source_association_key": "29bf0c90-3b40-0130-ae2d-dddd5893",
      "associated_commit_server_id": "551d8560c69d3b474b000102"
    }
  ]
}
```

```
}

```

- **Create a source**

- **Request Type:** POST
- **Request URL:** [https://<ctf\\_hostname>/api/stable/projects/PROJECT\\_URL\\_NAME/sources](https://<ctf_hostname>/api/stable/projects/PROJECT_URL_NAME/sources)
- **Request body:**

Required: `_type`, `display_name`

```
{
  "source": {
    "_type": "Build",
    "tool_id": "552c3736c69d3b3dfb000004",
    "display_name": "Sample_Build_Source",
    "associated_commit_server_id": "551d8560c69d3b474b000102"
  }
}
```

- **Return Format:**

```
{
  "source": {
    "id": "5578cd1c0fbb6ad771000001",
    "_type": "Build",
    "tool_id": "552c3736c69d3b3dfb000004",
    "active": true,
    "display_name": "Sample_Build_Source",
    "queue_username": "4d4eaf60-f1f9-0132-4687-185aeb8b25b8",
    "queue_password":
"d70bc954b64e368f1f650bfeaba2e2ff6e5525774ae42f9a959f4ccf2314e9cb",
    "queue_server": "amqp://rabbitmqserver.example.com:5672",
    "source_association_key": "29bf0c90-3b40-0130-ae2d-dddd5893",
    "associated_commit_server_id": "551d8560c69d3b474b000102"
  },
  "errors": {
    "base": [
      "error msg"
    ],
    "attribute_key": [
      "error msg"
    ]
  }
}
```

- **Return Codes**

- **Success :** 200
- **Error :** 404, 503, 422

- **Get one source**

- **Request Type:** GET
- **Request URL:** [https://<ctf\\_hostname>/api/stable/projects/PROJECT\\_URL\\_NAME/sources/SOURCE\\_ID](https://<ctf_hostname>/api/stable/projects/PROJECT_URL_NAME/sources/SOURCE_ID)
- **Return Format:**

```
{
  "source": {
    "id": "5578cd1c0fbb6ad771000001",
    "_type": "Build",
    "tool_id": "552c3736c69d3b3dfb000004",
    "active": true,
    "display_name": "Sample_Build_Source",
    "queue_username": "4d4eaf60-f1f9-0132-4687-185aeb8b25b8",

```

```

    "queue_password":
      "d70bc954b64e368f1f650bfeaba2e2ff6e5525774ae42f9a959f4ccf2314e9cb",
    "queue_server": "amqp://rabbitmqserver.example.com:5672",
    "source_association_key": "29bf0c90-3b40-0130-ae2d-dddd5893",
    "associated_commit_server_id": "551d8560c69d3b474b000102"
  }
}

```

- **Return Codes**

- **Success** : 200
- **Error** : 404, 503

- **Update one source**

- **Request Type**: PUT
- **Request URL**: [https://<ctf\\_hostname>/api/stable/projects/PROJECT\\_URL\\_NAME/sources/SOURCE\\_ID](https://<ctf_hostname>/api/stable/projects/PROJECT_URL_NAME/sources/SOURCE_ID)
- **Request body**:

```

{
  "tool_id": "552c3736c69d3b3dfb000004",
  "display_name": "Sample_Build_Source",
  "associated_commit_server_id": "551d8560c69d3b474b000102"
}

```

- **Return Format**:

```

{
  "errors": {
    "base": [
      "error msg"
    ],
    "attribute_key": [
      "error msg"
    ]
  }
}

```

- **Return Codes**

- **Success** : 200
- **Error** : 404, 503, 422

- **Activate one source**

- **Request Type**: PUT
- **Request URL**: [https://<ctf\\_hostname>/api/stable/projects/PROJECT\\_URL\\_NAME/sources/SOURCE\\_ID/activate](https://<ctf_hostname>/api/stable/projects/PROJECT_URL_NAME/sources/SOURCE_ID/activate)
- **Return Format**:

```

{
  "errors": {
    "base": [
      "error msg"
    ],
    "attribute_key": [
      "error msg"
    ]
  }
}

```

- **Return Codes**

- **Success** : 200
- **Error** : 404, 503, 422



- **Deactivate one source**
  - **Request Type:** PUT
  - **Request URL:** `https://<ctf_hostname>/api/stable/projects/PROJECT_URL_NAME/sources/SOURCE_ID/deactivate`
  - **Return Format:**

```
{
  "errors": {
    "base": [
      "error msg"
    ],
    "attribute_key": [
      "error msg"
    ]
  }
}
```

- **Return Codes**
  - **Success :** 200
  - **Error :** 404, 503, 422

## Return Types

The source objects have a set of common fields, while some of the fields depend on the `_type` of the source. The examples below show the type of source and the expected structure of the returned objects.

- **Source Type: Build**

```
{
  "source": {
    "id": "5578cd1c0fbb6ad771000001",
    "_type": "Build",
    "tool_id": "552c3736c69d3b3dfb000004",
    "active": true,
    "display_name": "Sample_Build_Source",
    "queue_username": "4d4eaf60-f1f9-0132-4687-185aeb8b25b8",
    "queue_password":
"d70bc954b64e368f1f650bfeaba2e2ff6e5525774ae42f9a959f4ccf2314e9cb",
    "queue_server": "amqp://rabbitmqserver.example.com:5672",
    "source_association_key": "29bf0c90-3b40-0130-ae2d-dddd5893",
    "associated_commit_server_id": "551d8560c69d3b474b000102"
  }
}
```

- **Source Type: Commit**

```
{
  "source": {
    "id": "5578cd1c0fbb6ad771000001",
    "_type": "Commit",
    "tool_id": "552c3736c69d3b3dfb000004",
    "active": true,
    "display_name": "Sample_ExternalRepository",
    "queue_username": "4d4eaf60-f1f9-0132-4687-185aeb8b25b8",
    "queue_password":
"d70bc954b64e368f1f650bfeaba2e2ff6e5525774ae42f9a959f4ccf2314e9cb",
    "queue_server": "amqp://rabbitmqserver.example.com:5672",
    "source_association_key": "29bf0c90-3b40-0130-ae2d-dddd5893",
    "repository_uri": "ssh://sample_gitserver/sample_project"
  }
}
```

```
}

```

- **Source Type: CommitCTF**

```
{
  "source": {
    "id": "5578cd1c0fbb6ad771000001",
    "_type": "CommitCTF",
    "tool_id": "552c3736c69d3b3dfb000004",
    "active": true,
    "ctf_id": "cmmt1234",
    "display_name": "Sample_ExternalRepository",
    "queue_username": "4d4eaf60-f1f9-0132-4687-185aeb8b25b8",
    "queue_password":
"d70bc954b64e368f1f650bfeaba2e2ff6e5525774ae42f9a959f4ccf2314e9cb",
    "queue_server": "amqp://rabbitmqserver.example.com:5672",
    "source_association_key": "29bf0c90-3b40-0130-ae2d-dddd5893",
    "repository_uri": "ssh://sample_gitserver/sample_project"
  }
}
```

- **Source Type: XDS**

```
{
  "source": {
    "id": "5578cd1c0fbb6ad771000001",
    "_type": "XDS",
    "tool_id": "552c3736c69d3b3dfb000004",
    "active": true,
    "display_name": "Sample Custom Activity Name One",
    "queue_username": "4d4eaf60-f1f9-0132-4687-185aeb8b25b8",
    "queue_password":
"d70bc954b64e368f1f650bfeaba2e2ff6e5525774ae42f9a959f4ccf2314e9cb",
    "queue_server": "amqp://rabbitmqserver.example.com:5672",
    "source_association_key": "29bf0c90-3b40-0130-ae2d-dddd5893",
    "associated_activity_source_id": "55429d045667ba223e000010",
    "activity_source_tag_names": ["Exclusive", "Limited", "Sample"],
    "background_color_value": "#61c0e5",
    "icon_uri": "/sample_icons/custom_source/
Sample_Icon_White_12.png"
    "custom_schema_store_id": object.custom_schema_store_id
  }
}
```

- **Source Type: Review**

```
{
  "source": {
    "id": "5578cd1c0fbb6ad771000001",
    "_type": "Review",
    "tool_id": "552c3736c69d3b3dfb000004",
    "active": true,
    "display_name": "Sample Review server Name One",
    "queue_username": "4d4eaf60-f1f9-0132-4687-185aeb8b25b8",
    "queue_password":
"d70bc954b64e368f1f650bfeaba2e2ff6e5525774ae42f9a959f4ccf2314e9cb",
    "queue_server": "amqp://rabbitmqserver.example.com:5672",
    "source_association_key": "29bf0c90-3b40-0130-ae2d-dddd5893",
    "commit_association_prefix": "Review"
  }
}
```

- **Source Type:** WorkItemSource

```
{
  "source": {
    "id": "5578cd1c0fbb6ad771000001",
    "_type": "Review",
    "tool_id": "552c3736c69d3b3dfb000004",
    "active": true,
    "display_name": "Sample Work Item Name One",
    "queue_username": "4d4eaf60-f1f9-0132-4687-185aeb8b25b8",
    "queue_password":
"d70bc954b64e368f1f650bfeaba2e2ff6e5525774ae42f9a959f4ccf2314e9cb",
    "queue_server": "amqp://rabbitmqserver.example.com:5672",
    "source_association_key": "29bf0c90-3b40-0130-ae2d-dddd5893",
    "associated_commit_server_id": "551d8560c69d3b474b000102"
  }
}
```

## Validations

The following constitute the validations that are performed for the API request-response parameters.

- **\_type**
  - type: string
  - present for all sources
  - must be one of: *Commit*, *CommitCTF*, *WorkItem*, *CustomActivity*, *Review*, *Build*, *XDS*
- **active**
  - type: boolean
  - present in all sources
  - must be true or false
  - default value: true
- **associated\_activity\_source\_id**
  - type: string
  - attribute in XDS, need not be present
  - references an existing activity source if present
- **commit\_association\_prefix**
  - type: string
  - attribute in Review, must be present
  - must be unique
  - key length must be less than 101
- **activity\_source\_tag\_names**
  - type: array with string elements
  - attribute in XDS, need not be present
- **background\_color\_value**
  - type: string
  - attribute of XDS, must be present
  - must be a valid RGB value or a HTML5 color name
- **associated\_commit\_server\_id**
  - type: string
  - attribute in Build and WorkItem, need not be present
  - must be associated with a valid commit server if present
- **custom\_schema\_source\_id**

- type: string
- attribute in XDS, need not be present
- must be associated to an existing `custom_schema_source` if present
- **display\_type**
  - type: string
  - must be present in all sources
  - maximum length less than 101
- **icon\_uri**
  - type: string
  - attribute of XDS, need not be present
  - has the default value: `"/assets/icons/custom_source/PNG_White/CN_Icon_White_12.png"`
  - must be a valid uri if present
- **repository\_uri**
  - type: string
  - must be present in Commit and CommitCTF
  - must be a valid uri if present
  - must be a valid protocol (*http, https, ssh, git*) if present
- **ctf\_id**
  - type: string
  - must be present in CommitCTF
  - must be a valid ctf id if present
- **source\_association\_key**
  - type: string in all sources
  - must be present
- **tool\_id**
  - type: string
  - need not be present or present in any source
  - must be associated with a valid tool if present

# Frequently asked questions about CollabNet TeamForge 17.4

---

Here is some background knowledge to help you make the most of TeamForge.

## TeamForge installation/upgrade/configuration/administration FAQs

---

To install CollabNet TeamForge, you download the software, make decisions about how you want the site to work, and set up data for the site to work with.

Before you begin the installation you will want to consider hardware requirements and other factors like supported software. For more information about installation see [Platform specification for TeamForge 17.4](#).


### How many servers do I need to run a TeamForge site?

You can run TeamForge on one server or distribute its services on multiple servers.

TeamForge functionality is delivered by several discrete services. Each service can run on its own server or share a server with one or more other services. You assign specific services to specific boxes when you customize your TeamForge installation by editing the `site-options.conf` file. See [TeamForge installation requirements](#) on page 441 for more information.

#### TeamForge core functionality

This is known internally as the `app` server. It implements JBoss and Apache services. One and only one instance of this application must be present.

 **Tip:** TeamForge 17.4 supports an integration with the Review Board, a source code review tool supported by TeamForge, if required, must be installed on the TeamForge `app` server.

#### Database

The `database` application handles site users' data. You set the type of database by setting the value of the `DATABASE_TYPE` token to `oracle` or `postgresql`. One and only one instance of this application must be present. TeamForge's operational `database` can run on the TeamForge `app` server or on a separate server.

#### Subversion

Subversion can be used to provide source control functionality. It uses the Tomcat and Apache services. A site can have zero, one, or more than one instances of the `svn` application, running on an arbitrary number of servers.

#### CVS

CVS can be used to provide source control functionality. It uses the Tomcat service. A site can have zero, one, or more than one instances of the `cv`s application, running on an arbitrary number of servers.

#### etl

The reporting service.

#### datamart

A mirror of your site database for the reporting engine to work with.

#### EventQ

TeamForge EventQ is a TeamForge capability that provides traceability for product life cycle activities such

as work items, SCM commits, continuous integration (CI) builds, and code reviews. While you can run both EventQ and TeamForge on the same server, CollabNet recommends such an approach only for testing purposes. It's always recommended to run EventQ on a separate server for optimal scalability. See [EventQ installation requirements](#) for more information.

## Git

TeamForge 17.4 supports an integration with the Git distributed version control tool powered by Gerrit. TeamForge 17.4 supports Git integration based on Gerrit 2.13.x. For more information, see [Set up Git](#) on page 591.

In principle, services can be combined in any configuration, with some constraints, such as:

- Only one site database, datamart and reporting service.
- Multiple source control integration services, but only one per box.

In practice, CollabNet has identified few configurations as the most useful for a wide variety of site. You can follow the instructions here to set up your site in one of these configurations, or you can adapt one of them to your own conditions.

## Should I upgrade to TeamForge 17.4 on a new server?

You can upgrade TeamForge on new hardware, or you can stick with the same server your current site is running on. To choose, consider how your members use your site and how you maintain it.

### Upgrade on new hardware

Upgrading on new hardware is a little more complex, but provides more flexibility.

Do it this way if you want to:

- Change to a different operating system.
- Stage your upgraded site before moving it into production.
- Serve the site from a different data center.
- Support an expanded user base.
- Minimize downtime associated with the upgrade.

### Upgrade on the same hardware


If you are OK with your current hardware setup and you don't want to rock the boat, you can run your new TeamForge site on the same box where your current site is running.

Do this if you want to:


- Keep your upgrade as simple as possible.
- Stick with a tested hardware setup.
- Use virtual hardware.
- Support a stable user base.

## Should I move my TeamForge database to its own server?

If you expect your site to have heavy user traffic, you may want to run the site's database on its own server.

 **Note:** Before moving your database to its own server, make sure you have access to someone with advanced skills in the database service you are using.

The advantage of hosting a service on a separate server is that it does not share CPU, RAM or I/O bandwidth with the server that is hosting the main TeamForge application.

 **Note:** Each TeamForge site can have only one database server.

To help decide whether you need a separate database server, consider these approximate values:

	Shared TeamForge-database server	Standalone database server
Daily users	Fewer than 1000	More than 1000
Daily discussion forum entries	Fewer than 1000	More than 1000

### Should I move my source control application to its own server?

If you anticipate heavy source code check-in and check-out traffic, consider setting up the source code application on its own server.

To host your source control services on their own server, you must set up a source code repository server and integrate it with TeamForge. You can integrate any number of source code servers with your TeamForge site.

The advantage of hosting a service on a separate server is that it does not share CPU, RAM or I/O bandwidth with the server that is hosting the main TeamForge application.

 **Note:** If you need to move a source code integration, contact your CollabNet representative for help.


To help decide whether you need a separate source control server, consider these approximate values:

	Shared TeamForge-SCM server	Standalone SCM server
Daily source code commits	Fewer than 1000	More than 1000

### What are the right PostgreSQL settings for my site?

Your site's PostgreSQL settings depend on the conditions your site is operating under, especially the number and size of projects and the number of users.

The default values in the `site-options.conf` file are designed for a TeamForge site running on a system with 8 GB of RAM. This table contains recommended values for systems with various amounts of RAM, based on testing carried out in CollabNet's performance lab. Use your discretion in selecting the right values for your environment.

 **Note:** Remember to recreate the runtime environment after changing any value in the `site-options.conf` file.

#### Recommended values if PostgreSQL and TeamForge are on the same server

site-options.conf tokens	8GB RAM	16GB RAM	32GB RAM	64GB RAM	128GB RAM
PGSQL_EFFECTIVE_CACHE_SIZE	6GB	12GB	24GB	48GB	96GB
PGSQL_SHARED_BUFFERS	2GB	4GB	8GB	16GB	32GB
PGSQL_WORK_MEM	64MB	64MB	64MB	64MB	64MB
PGSQL_WAL_BUFFER_SIZE	32MB	32MB	32MB	32MB	32MB
PGSQL_MAINTENANCE_WORK_MEM	615MB	615MB	615MB	615MB	615MB

### Recommended values if PostgreSQL is on a separate server

site-options.conf tokens	8GB RAM	16GB RAM	32GB RAM	64GB RAM	128GB RAM
PGSQL_EFFECTIVE_CACHE_SIZE=	1GB	2GB	24GB	48GB	96GB
PGSQL_SHARED_BUFFERS=	2GB	4GB	8GB	8GB	8GB
PGSQL_WORK_MEM=	4MB	64MB	64MB	64MB	64MB
PGSQL_WAL_BUFFERS=	16MB	32MB	32MB	32MB	32MB
PGSQL_MAINTENANCE_WORK_MEM=	16MB	64MB	615MB	615MB	615MB

### Is it possible to change artifact prefixes in TeamForge?

Artifact prefixes in Teamforge cannot be changed or customized, all artifacts across the site will have an artifact ID of artfXXXXXX.

### I have made changes to the James config, how can I enable them without restarting TeamForge?

You can restart James without restarting all of TeamForge with the following command: `/opt/collabnet/teamforge/bin/teamforge restart -s phoenix`. Please keep in mind that any changes made to the James config that are not stored in the TeamForge configuration may be overwritten when doing future upgrades.

### Why am I getting an error about NoEmailShouldEverBeSentToNoBody@localhost?

This error is caused by the 'nobody' user somehow being set as monitoring one or more items in TeamForge.

You can see if this is the case with your TeamForge install by running the following against your TeamForge database:

```
select count(*) from mntr_subscription where user_id='user1002';
```

If the above returns a count greater than zero, then this is the problem. To correct this, do the following:

- Completely stop TeamForge
- Log into the TeamForge database
- Execute the following code: `delete from mntr_subscription where user_id=user1002;`
- Exit your database session
- Start TeamForge

**Note:** You should make a backup of your database before attempting this procedure.

### How does TeamForge handle third-party applications?

TeamForge relies on many third-party applications to augment or enhance functionality.

TeamForge integrates with additional third party applications, such as Microsoft Office. Support will always make an effort to provide assistance in using third party applications. However, for complete, end-to-end support, customers should consult the application vendor, as the vendor is best equipped to provide the support necessary to use their products.

### CVS

Technical Support provides best-effort support for Subversion and CVS client usage issues. TeamForge is not shipped with CVS source control functionality. For best results, contact the vendor for assistance.



The CVS RPM that ships with RedHat Linux Enterprise Server 3 and RedHat Advanced Server 2.1 has a known bug that prevents users who have access to 32 or more CVS repositories from accessing the repositories that are alphabetically after the 31st. This is currently RedHat bug #131124 ([https://bugzilla.redhat.com/bugzilla/show\\_bug.cgi?id=131124](https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=131124)). Customers are advised to contact Red Hat for a solution to this bug for any case where users are members of 32 or more CVS repositories on a TeamForge CVS server.

### Discussion forum threading

For TeamForge discussion forums to properly thread posts sent in via email, the email message must include either the `References` or `In-Reply-To` header. Email received without both of those headers cannot be threaded accurately and will most likely be treated as a new topic or thread in the discussion. While the lack of either of the headers is not an explicit RFC violation, the inclusion of such headers is considered compliance with section 3.6.4 of [RFC 2822](#).

Microsoft Outlook and Lotus Notes are prone to sending mail without at least one of the required headers. There is evidence that Lotus Notes versions 6.5 and newer are capable of sending email that includes at least one of the two required headers. However, older versions of Lotus Notes either do not include the headers, or require special reconfiguration in order to do so.


Microsoft Outlook on its own does include the `In-Reply-To` header. However, mail is sent through a Microsoft Exchange server, that header is stripped off. There are no known versions of Microsoft Exchange server that do not strip these RFC headers from outbound email, and therefore there are no known workarounds.

Contact your IT group or the vendor of your email client with questions or concerns.

## When do I run the initial load job?

You can run the initial load job any time after the site is upgraded to TeamForge 17.4. We recommend that you run it before you hand over the site to the users.

You can run the job when the site is in use. On an average, there might be a ten percent degradation in the response time when the job is running. If the job is not run as a part of the upgrade, we recommend that you run it at a time when the site usage is relatively low. The initial load job may consume more time to complete and this depends on the data available in the site. So this job should be triggered when the site is stable after the upgrade and the chances of site shut down is remote.

 **Note:** The initial load job can be run immediately in case of a fresh installation.

### How do I list all the ETL jobs?

You can list all the ETL jobs by running the following command.

```
./etl-client.py -a
```

### How do I run the initial load jobs?

Run the scripts specified below to run the initial load jobs.

```
$ [RUNTIME_DIR]/scripts/etl-client.py -r SCMCommitInitialJob
$ [RUNTIME_DIR]/scripts/etl-client.py -r TrackerInitialJob
```

### What happens if I run an initial load job that has already run successfully?

Re-running the initial load job does no harm. The system performs an internal check and the job is aborted.

### How do I check the status of the job?

Site administrators will be notified through mails if there is a job failure. The following command will help you identify whether a job is currently running or not.

```
$ [RUNTIME_DIR]/scripts/etl-client.py -a
```

**What happens if there is a job failure?**

An email will be sent to site administrators in case of a failure. The job should be re-triggered manually. Data will be harvested from the last completed milestone prior to the failure.

**Is it mandatory to run the initial load job for a new site?**

Yes, it is mandatory to run the initial load job for a new site.

**What happens if I miss running the initial load during the installation or upgrade?**

You can run this any time after the site is upgraded to TeamForge 17.4. However, the incremental data harvesting is blocked until the initial load is run. This is true for an existing site that is upgraded to TeamForge 17.4 as well as for a new site.

**Why does the incremental\_etl job fail when the initial load job is not triggered?**

The incremental load is made to fail if the initial load is not run. You can ignore this.

**Why do ETL jobs fail post TeamForge upgrade?**

This could be due to incompatible database and JDBC driver versions. See [Why do ETL jobs fail post TeamForge upgrade?](#) on page 962

**When do I run the Tracker initial load job?**

You can run the Tracker initial load job any time after the site is upgraded to TeamForge 17.4. We recommend that you run it before you hand over the site to the users. This task might take few hours depending on the amount of artifact data in the site.

You can run the job when the site is in use. On an average, there is a ten percent degradation in the response time when the job is running. If the job is not run as a part of the upgrade, we recommend that you run it at a time when the site usage is relatively low. The initial load job may consume more time to complete and this depends on the data available in the site. So this job should be triggered when the site is stable after the upgrade and the chances of site outage is remote immediately.

**How do I run the job?**

```
$ [RUNTIME_DIR]/scripts/etl-client.py -r TrackerInitialJob
```

**How do I list all the ETL jobs?**

You can list all the ETL jobs by running the following command.

```
./etl-client.py -a
```

**How do I check the status of the job?**

Site administrators will be notified through mails if there is a job failure. The following command will help you identify whether a job is currently running or not.

```
$ [RUNTIME_DIR]/scripts/etl-client.py -a
```

**What happens if there is a job failure?**

An email will be sent to site administrators in case of a failure. The job should be re-triggered manually. Data will be harvested from the last completed milestone prior to the failure.

**Is it mandatory to run the initial load job for a new site?**

Yes, it is mandatory to run the initial load job for a new site.

**What happens if I miss running the initial load during the installation or upgrade?**

You can run this any time after the site is upgraded to TeamForge 17.4. However, the Tracker incremental data harvesting is blocked until the initial load is run. This is true for an existing site that is upgraded to TeamForge 17.4 as well as for a new site.

**Why does the tracker\_incremental\_etl job fail?**

The incremental load is made to fail if the initial load is not run. You can ignore this.

**Why do I get an error during tracker incremental load job?**

An error (as illustrated below) shows up during the tracker incremental load job in case there is no data to harvest for certain fields.

You can ignore such errors that have no impact on the incremental load job's success or failure.

```
2015-11-02 02:30:05,547 ERROR [TrackerIncrementalJob] SetVar - We can not
field
[date_to_harvest_teams] in the previous job entry result!
2015-11-02 02:30:05,940 ERROR [TrackerIncrementalJob] SetVar - We can not
field
[date_to_harvest_dependency] in the previous job entry result!
2015-11-02 02:30:06,399 ERROR [TrackerIncrementalJob] SetVar - We can not
field
[date_to_harvest] in the previous job entry result!
2015-11-02 02:30:11,278 ERROR [TrackerIncrementalJob] SetVar - We can not
field
[date_created] in the previous job entry result!
2015-11-02 02:30:11,493 ERROR [TrackerIncrementalJob] SetVar - We can not
field
[date_membership_removed] in the previous job entry result!
```

**Can I run other java applications in the same JBoss instance as CollabNet TeamForge?**

It is recommended to create a separate instance of JBoss or tomcat and deploy your applications there. While it is possible to deploy other applications alongside CTF, you may encounter some errors.

In order to deploy an application into the CTF JBoss instance, you must place the war files into the deploy directory beneath the runtime directory. All directories beneath the runtime directory are recreated each time the application runtime is rebuilt, deleting your applications.

If your applications include the CTF SDK, you may receive class conflicts and errors such as "java.lang.reflect.InvocationTargetException" when attempting to connect to CTF.

**Users invoked via su command in TeamForge**

TeamForge invokes the super user (su) command during service start or stop.

The following are the operating system users that are invoked via the su command.

Users	TeamForge component that invokes the user	Purpose
sf-admin	TeamForge integration server (Tomcat)	During service start (collabnet start tomcat)
sf-admin	TeamForge ETL server (Tomcat)	During service start (collabnet start etl)

Users	TeamForge component that invokes the user	Purpose
sf-admin	TeamForge Application Server (JBoss)	During service start (collabnet start jboss)
PostgreSQL	TeamForge database server (Postgres)	During service start/stop (postgresql-9.6 start/stop) and service deploy ("teamforge" script)
bds-codesight	Code search database server (Postgres)	During service start/stop (bds-codesight-postgresql start/stop) and service deploy ("teamforge" script)
gerrit	Gerrit integration server (Jetty)	During service start (collabnet start gerrit)

### Why does the the Yum installer display a warning message on CentOS 6?

The yum installer displays a warning message on Centos 6 if you are not using the latest version update of CentOS or if the new GNU Privacy Guard (GPG) keys are not in place.

On the Centos system, the new GPG keys are generally sent as an update before the expiry of the old keys. You must perform the version update regularly to receive the new GPG keys.

To avoid the warning message, use the "-nogpgcheck option" in the yum install command.

```
yum install --nogpgcheck teamforge
```

### Why am I getting a Yum repository filename conflict?

This means that you have created a repository filename starting with `collabnet`.

The repository file for the TeamForge installer begins with "collabnet-\*.repo". If you have a customized repository with the same name, then the TeamForge installer renames the repository file with ".backup" as an extension and installs the repository file. It is better not to have any repository file name that begins with `collabnet` apart from the one that is installed by the TeamForge installer.

### Why is a user restricted and cannot be changed?

If using the TeamForge integration with CVS, the machine running CVS must be using the default Linux authentication method (/etc/passwd, /etc/group) and not a network service (NIS, MS AD, MS winbind, etc).

### How do TeamForge licenses work?

When you purchase a TeamForge license, you get the right to assign licenses to a specified number of users.

TeamForge supports a more flexible and granular approach to tool instantiation. TeamForge's license model consists of the following license types: ALM, SCM, DevOps, Version Control, Collaboration and Trackers. These license types are tailored to suit the needs of specific set of users that need access to certain tools and functions. A TeamForge user's license type determines the tools available to the user in TeamForge.



**Attention:** While TeamForge supports more selective tool options with these new license changes, there's no impact on customers, both new and existing, requiring all the tools that TeamForge supports.

The following table lists the license types and the tools that go with them (refer to the table at the end of this topic for a complete list of tools and functions).

License type	Available tools
ALM	Offers the full range of ALM tools and functions to users.
SCM	Offers core Source Code Management tools and functions to users. Includes all the ALM tools and functions but Tracker and Documents component.

License type	Available tools
Version Control	Offers Source Code Management, File Releases and Review tools and functions to users.
DevOps	Package Management (Application & Environment) and File Releases tools and functions to users.
Collaboration	Offers a range of collaboration tools such as Documents, Wiki and Discussion Forums to users.
Trackers	Offers TeamForge's Tracker capability (Trackers, Planning Folder, Teams, Planning Board, Task Board and Kanban Board) to users. Also includes File Releases.

In addition to the tools offering, the Reporting framework is also controlled by the licenses you have. Meaning, you can view and generate reports based on the license types assigned to you. For example, you must have SCM license to view or generate SCM activity reports.

Check with your CollabNet representative if you aren't sure how many licenses or what kind of licenses you want/have.

- Your license key contains a few important numbers:
  - The number of users eligible to use specific licenses such as ALM, SCM, DevOps, Version Control, Collaboration and Trackers.
  - The IP address of the machine that your site runs on.

For example, if your organization has 80 users who will use only the source code management features, 10 users who will use DevOps and 100 users who need the Application Lifecycle Management features, your license key string may look like this:

```
ALM=100:SCM=80:DEVOPS=10:supervillaininc:144.16.116.25.:302D02150080D7853DB3E5C6F67EAB
```

- Your license key only works for the IP address (or range of addresses) of the hardware that your CollabNet TeamForge is running on, as specified in your order form. If your site uses a separate server for its database or source code repositories, make sure your license key reflects the IP addresses of all the necessary servers. If any of these addresses change, ask your CollabNet representative to generate a new key.
- When you create a user account, you can assign the user with available licenses.
- You can purchase a TeamForge license for as many years as you want. The validity period is encoded into your license when it is generated by your CollabNet representative.
- How many users your site can support depends on the type of license. Check with your CollabNet representative if you aren't sure what kind of licenses you have.
- Your license key is attached to the IP address of the server where your site runs. You can get a license key for a single IP address or for a range of IP addresses.
- Your service year starts the first time you log into your site, or the first time you create or edit any item on your site, such as a tracker artifact or a document. Whichever of these events comes first starts the clock.
- The expiration date of your license is shown on the **License Info** page. (Go to **My Workspace > Admin** and select **Projects > License Info**).
- When your service year expires, you can still see the project data on your site, but you cannot make any changes to it. However, you can still carry out some critical maintenance functions for your site:
  - Enter a new license key.
  - Disable or delete users.
  - Change user passwords.
  - Get forgotten user passwords.
- Except deleted users, all other users in TeamForge such as active users, pending users, disabled users, and expired users continue to consume licenses.

**Table 3: Tools Availability Matrix**

Tools	ALM	SCM	Version Control	DevOps	Collaboration	Tracker
Agile Task and Planning Boards	✓					✓
Tracker	✓					✓
Git/SVN Repository Management and Replication	✓	✓	✓			
Code Review	✓	✓	✓			
Build Automation	✓	✓	✓	✓		
Test Management	✓					✓
File Releases	✓	✓	✓	✓		✓
Binary Repository Management	✓	✓	✓			
EventQ/Activity Stream	✓	✓	✓	✓		✓
Access Controls	✓	✓	✓	✓	✓	✓
Project Work Spaces	✓	✓	✓	✓	✓	✓
Wiki and Discussion Forums	✓	✓			✓	
Document Management	✓				✓	
User Management	✓	✓	✓	✓	✓	✓
Flexible Process and Toolchain Templates	✓	✓	✓		✓	✓
Reusable Dashboards	✓	✓	✓	✓	✓	✓
Categories and Groups	✓	✓	✓		✓	✓

Tools	ALM	SCM	Version Control	DevOps	Collaboration	Tracker
Cross Project Visibility	✓	✓	✓			✓
Cross Project Reporting and Dashboards	✓	✓	✓	✓		✓
Cross Project Search	✓	✓	✓	✓	✓	✓
Site-wide Administration	✓	✓	✓	✓	✓	✓
Custom Branding and Custom Integrations	✓	✓	✓	✓	✓	✓
Security Architecture	Single Tenant	Single Tenant	Single Tenant	Single Tenant	Single Tenant	Single Tenant
Onsite and Hosted Provisioning	✓	✓	✓		✓	✓
SVN Auto Updates	✓	✓	✓			
SVN Repository Backup and Monitoring	✓	✓	✓			

## What happens when my site license expires?

When your service year expires, you can still see the project data on your site, but you cannot make any changes to it.

Your service year starts the first time you log into your site, or the first time you create or edit any item on your site, such as a tracker artifact or a document. Whichever of these events comes first starts the clock.

The expiration date of your license is shown on the **License Keys** page. (See the *About CollabNet TeamForge* page for the company name associated with the license.)

If your site has switched to read-only mode, you can still carry out some critical maintenance functions:

- Enter a new license key.
- Change users' license types.
- Deactivate or delete users.
- Change user passwords.
- Get forgotten user passwords.
- Use the comprehensive help on [help.collab.net](http://help.collab.net).

## How to reinstall a deleted installation directory?

You can reinstall an installation directory if it has been deleted inadvertently.

Before attempting to reinstall the deleted installation directory the remnants must be wiped out completely. Rather, remove the `collabnet-local.repo` file under `/etc/yum.repos.d` before reinstallation. This is mandatory for a clean and complete installation.

## How do I configure the user session timeout?

Use the `SESSION_TIMEOUT` token to set the user session timeout duration for newly created Application Server/Integration Server sessions.

The default value of the `SESSION_TIMEOUT` token is set as 30 minutes (for security reasons). You may change this to any value in the range of 1-1440 (minutes). However, you must create runtime for the changes to take effect.

## Where do I configure my client proxy settings?

Configure proxy settings in the servers file (created when Subversion is installed on your system).

The server file is created when you install TortoiseSVN, Eclipse or command-line Subversion. Use the appropriate path from the installation folder to configure proxy settings:

CLI Unix/Linux	/home/<username>/subversion/servers
CLI Windows	C:\Documents and Settings\<username>\Application Data\Subversion\servers
Eclipse	Window > Preferences > General > Network Connections
IE (6/7)	Tools > Options > Advanced Network > Connection Settings
Firefox	Tools > Internet Options > Connections tab > LAN Settings
TortoiseSVN	Windows Explorer > File > TortoiseSVN > Settings > Network (by default, TSVN uses browser proxy settings)

## How do I configure the timeout for Apache in TeamForge?

This can be changed by editing the setting `WWW_SERVER_TIMEOUT`.

The timeout for Apache (httpd) can be configured in seconds. Default is 300 seconds. You must re-create the runtime for the change to take effect.

## How do I make TeamForge work the same when the IP address of the server changes?

Create runtime whenever the IP address changes.

```
teamforge provision
```

## How do I change the JBoss log level in TeamForge 17.4?

Use the **Configure Logging** page to change the JBoss log level in TeamForge 17.4.

For more information, see [Configure your site's log level](#) on page 365.



### Note:

- The default JBoss log level is **INFO**.
- JBoss restart is not required after changing the log level using the **Configure Logging** page.
- Once you change the log level in the **Configure Logging** page, the changes you may do to the `standalone-full.xml` file to change the log level have no effect. In other words, the log level set using the **Configure Logging** page is persisted in the database and therefore overrides the log level in the `standalone-full.xml` file.

## Upgrade to TeamForge 8.0 and later fails if the field 'team' already exists. How do I resolve this?

While upgrading to TeamForge 8.0 (or later) you may encounter an error stating that the field `team` already exists for the tracker(s).



This is because if the trackers in your project already have a field with the name `team`, it conflicts with the Team feature introduced in TeamForge 8.0. This error occurs while running the migration script `migrate.py`. To resolve the issue, do the following and then run the `migrate.py` script again:

1. Identify the tracker(s) where the `team` field is present by running the following script:

- For the `folder_layout` table:

```
select folder_layout.id, field_name, title, folder_id tracker_id from
folder_layout, folder where field_name = 'team' and folder_id like
'tracker%' and folder.id = folder_layout.folder_id;
```

- For the `field` table:

```
select field.id, name, title, folder_id from field, folder where
name='team' and folder_id like 'tracker%' and field.folder_id =
folder.id;
```

2. Rename the field `team` as `original_team` by running the following script:

- For the `folder_layout` table:

```
update folder_layout set field_name='original' || '_team' where
field_name='team';
```

- For the `field` table:

```
update field set name='original' || '_team' where name='team';
```

3. Verify the newly renamed fields by running the following script:

- For the `folder_layout` table:

```
select folder_layout.id, field_name, title, folder_id tracker_id from
folder_layout, folder where field_name like '%team' and folder_id like
'tracker%' and folder.id = folder_layout.folder_id;
```

- For the `field` table:

```
select field.id, name, title, folder_id from field, folder where
name like '%team' and folder_id like 'tracker%' and field.folder_id =
folder.id;
```

4. Re-run `migrate.py`.



**Important:** Before re-running this script, make sure you delete the `ctf72_ctf80` folder (which was created the first time you ran the migration script), because every time you run the script, this folder gets created.

## Troubleshooting TeamForge errors

---

Common errors and solutions in TeamForge.s

Concept definition.

### I deleted the email that verifies my CVS password. How can I retrieve my password?

You can sync up your TeamForge login and CVS passwords by changing your password in TeamForge.

To change your password

1. Log in and click on My Workspace.
2. Click on My Settings.
3. Click on **Change Password**.
4. Enter your old password and new password.
5. Click **Submit**.



**Note:** The old and new passwords can be the same if you do not wish to actually change it.

## Not all my team members appear in the "assigned to" field. Why is that?

In order to appear in the Assigned To field, the user must belong to a role giving them edit/view permissions on the tracker.

## Why don't the branding repo changes get rendered into UI?

It may be due to the property 'subversion\_branding.repository\_base' pointing to /sf-svnroot instead of the /svnroot directory, which is used by the scm-integration of the csfe installation.

First, check the location of the branding repository in subversion\_branding.repository\_base=/sf-svnroot' in /opt/collabnet/teamforge/runtime/conf/sourceforge.properties.

If it has to be /svnroot, then add an entry that states SUBVERSION\_BRANDING\_REPOSITORY\_BASE=/svnroot  
Then re-create a runtime and restart TeamForge.

## Why am I getting "Could not connect" status for my email and search server?

On the System Tools page, when you see "Could not connect status for search and email servers," you must stop and start your phoenix.sh process.

You may also need to set the *JAVA\_HOME* environment variable to the location of your JDK.

The stop/start Phoenix commands:

```
sh /opt/collabnet/teamforge/runtime/scripts/phoenix.sh stop
sh /opt/collabnet/teamforge/runtime/scripts/phoenix.sh start
```

## Why am I not able to see the charts for tracker metrics?

You may not be able to see the charts for the tracker metrics if the tracker initial load is not running correctly.

The incremental data collection is disabled until the initial load is run. You can check if the initial load is completed successfully by executing the query below from the Ad-hoc reporting page against the datamart.

```
select status from etl_job where job_name='tracker_initial_etl';
```

You must get the status value as 1 if the initial load is completed successfully. Otherwise, you must trigger the job manually by executing the command: [RUNTIME\_DIR]/scripts/etl-client.py -r TrackerInitialJob

## Post install fails for Git. What should I do?

Sometimes, the post installation fails during the GIT setup. In such a scenario, you need to manually kill the process using the following command and then do a post install again:

```
kill -9 <ProcessID>
```

## Why am I getting a 'Not running' message when the Datamart service is stopped?

When TeamForge and Datamart are running in a single instance, the TeamForge database is stopped when you stop the Postgres services. The message, 'Not running' is displayed when you stop the Datamart service. You can ignore this message.

## Why are the dynamic images that TeamForge creates broken?

If you have a fresh install of TeamForge and you're noticing that the dynamic images are not correct, you may be missing a library that is needed to create the images.

The easiest way to find this is to check and see if you have the xorg-x11-depreciated-libs rpm installed:

```
rpm -qva | grep xorg-x11-depreciated
```

Watch for the results. If you see that you have the xorg-x11-depreciated-libs rpm installed, and after a server restart you're still not seeing the images, please open a support request. If you do not have the xorg-x11-depreciated-libs rpm installed, it can usually be installed by performing a simple: `up2date xorg-x11-depreciated-libs` and restarting TeamForge.

## Due to firewall restrictions I cannot send email from James. How can I resolve this?

If James is unable to send email directly due to firewall restrictions, or mail being rejected from the application servers IP address, you may have to configure it to use a gateway mail server to send outgoing messages through.

To do this, you will need to add the following to the `<mailet match="All" class="?RemoteDelivery">` directive in the james config file at `/opt/collabnet/teamforge/james/james-<version>/apps/james/SAR-INF/config.xml`:

```
<gateway>smtp.example.com</gateway>
<gatewayPort>25</gatewayPort>
```

You should find these commented out on line 362 of the config file. If your gateway mail server requires authentication to send email, you may also add the following directives:

```
<username>username</username>
<password>password</password>
```

## Why is my email taking a long time to arrive?

TeamForge uses the James MTA to send and parse all email coming to and from the system. In this case, the best course of action is to look in the james mailet logfile.

Your logfile will help you to determine what is going on with the emails that are being sent from your system. Your logfiles will look very similar to this:

```
07/02/07 07:54:43
INFO James.Mailet: ?RemoteDelivery:
Attempting delivery of Mail1170135534355-39088-to-domain.invalid to host
domain.invalid at 192.168.0.1 to
addresses [invalid.user@domain.invalid] 07/02/07 07:55:43
INFO James.Mailet: ?RemoteDelivery: Could not connect to SMTP host:
192.168.0.1, port: 25; nested exception
is: java.net.ConnectException: connection to 192.168.0.1 timed out 07/02/07
07:58:43
INFO James.Mailet: ?RemoteDelivery: Storing message
Mail1170135534355-39088-to-domain.invalid into outgoing after 7 retries
07/02/07
07:58:43 INFO James.Mailet: ?RemoteDelivery: Attempting delivery of
Mail1170831482124-2756-to-company.com to host mx.company.com. at 127.0.0.1
to addresses
[good.user@company.com]
```

As you can see, the James MTA stores outgoing emails to resend at a later time. These files can be located in this directory: `/opt/collabnet/teamforge/james/james-<ver>/apps/james/var/mail/outgoing/`

When you do a directory listing of the files, you will see a listing of files very similar to this:

```
4D61696C3131373038333323131343031322.Repository.?FileObjectStore
4D61696C3131373038333323131343031322.Repository.?FileStreamStore
```

The FileObjectStore is a binary file, but, the FileStreamStore can be viewed with an editor or your favorite paging program in order to determine the contents. Sometimes, the directory can grow to a large number, where you will not be able to use a standard bash expander to delete all of the files. In that case, use the following shell script to remove all of the objects from the outgoing directory:

```
for i in * ; do /bin/rm $i; done
```

## How do I enable the Postgresql log files archiving when the services are not started using the CollabNet startup script?

The Postgresql log files archiving can be enabled by running a simple command.

It is recommended you use the teamforge script to start and stop the Postgres services. However, if you use the Postgresql init script to start or stop the Postgres services, the postgresql log files are not archived by default.

To enable the Postgresql log files archiving, run the following command:

```
/opt/collabnet/teamforge/bin/teamforge -s ctf-postgresql-9.6 start
```

## Why am I not getting any error messages when executing the Subversion upgrade script?

Error messages may come when Subversion is installed with a dependent package from an unknown source.

The Subversion working copy script assumes that Subversion is installed with the dependent packages from a proper source repository(RHEL/CollabNet). If you install any dependent packages from any unknown source that is not authorized by RHEL/CollabNet, it will result in inconsistency and this cannot be handled by the Subversion working copy script.

## Why do we have errors creating or altering repositories and adding or removing users?

The TeamForge SCM Integration server runs an instance of Tomcat and then launches TeamForge inside the Tomcat container.

If you are experiencing issues creating or altering repositories or adding and removing users from repository access, and the other TeamForge integration logs are not providing any clues, you may wish to review the Tomcat log at: `/opt/collabnet/teamforge/log/integration/catalina.out`

Sometimes, OS-level errors will be flagged into this log and not others. In our experience, it is pretty rare to find something in this log that is not logged elsewhere.

## Why am I getting an email specifying that the ETL job has failed?

You are getting this email because one of the Extract Transformation and Load (ETL) jobs has failed during the run.

You can see the `etl.log` for more details to find out the reason for the job failure.

The ETL job failure may happen because of the following reasons:

- Out of memory error.
- No response from the database.

If the ETL job failure is happening for the first time, you can restart the ETL (`[RUNTIME_DIR]/scripts/collabnet restart etl`) and check if the problem is occurring again. You can increase the JVM heap size by specifying the same in `ETL JAVA_OPTS` if the problem keeps recurring. The default value is `-Xms160m -Xmx256m`. You can increase the heap size depending on the memory available in the box.

Check if both TeamForge and Datamart are up and responding to queries if there is no response from the database. Restart the ETL (`[RUNTIME_DIR]/scripts/collabnet restart etl`).

Contact CollabNet support if the problem persists.

## My custom event handler is not working! What is wrong?

If you get a message in server.log (or events.log) that your event handler could not be parsed, these are some points you might check.

<b>Did you change the package of your class containing handler code?</b>	This is not supported.
<b>Did you move event.xml to another directory?</b>	This is not supported.
<b>Did you compile your java code with a Java version that is not supported by the JVM running on the TeamForge server?</b>	
<b>Did you use a very long file name for your JAR file?</b>	Only 31 characters are supported.
<b>Did you include libraries in your JAR file that have already been part of the default TeamForge class path?</b>	In this case, your libraries will not be picked up.
<b>Did you try to reference internal TF classes from packages other than the SOAP namespace?</b>	In this case, the TeamForge security classloader will reject your JAR file.
<b>Did you change the DTD location referenced in events.xml?</b>	This will not work.
<b>Did you try to trigger a follow up action in a synchronous event handler?</b>	This will most likely result in a time out exception, since your code is already running in a transaction that locks the resources you are trying to obtain a lock for.
<b>Did you trigger a TeamForge SOAP call with the session ID passed to your handler which threw an exception?</b>	See <a href="#">Using an asynchronous event handler: trigger follow-up events</a> on page 827 for an explanation why this does not work.

## Why do I get a JBoss error - "failed to start in 240 seconds, giving up now" - while installing TeamForge?

You get this error when the system's RAM is less than the minimum recommended value of 4GB. However, it's most likely that JBoss will start within a few minutes.

To make sure that JBoss starts up, check the service.log file using this command:

```
tail -f /opt/collabnet/teamforge/log/apps/service.log
```

If you see messages like the following, the TeamForge application will start in a few minutes.

```
Check Port Available PASSED: Port 4444 on localhost is available
Check Port Available PASSED: Port 4445 on localhost is available
Waiting for application server to start up.. this can take a few
minutes.
```

## JBoss crashed with out of memory error, how do I prevent this?

This can indicate that the JVM heap size is set too small.

You can adjust this by changing the -Xms and -Xmx settings of the JBOSS\_JAVA\_OPTS token in site-options.conf and rebuilding runtime.

This will appear if the JBoss application server has crashed and you find this error in the server.log:

```
INFO [STDOUT] java.lang.OutOfMemoryError: Java heap space
```

The default maximum heap size of 640MB can cause issues on a heavily used site. If the CTF application is the only thing running on the server, you can increase this to half of the total physical ram on the machine. This should still allow enough memory for the OS and other necessary processes. If you are also running the app, database and scm on the same machine a maximum heap size of 1/4 or the total ram maybe a better setting. Determining the right JVM settings for your install will require testing with your particular usage patterns and database.

You can view the current memory usage under the JVM Environment section of the JBoss webconsole at `http://<CTF_SERVER>:8080/web-console/`. You will need to log in using the CTF admin password.

## JBoss status is in 'starting' for a long time. How to have JBoss started successfully?

If JBoss is not started successfully (status remains 'starting' for a long time), you may have to wait until it starts up successfully or you can kill JBoss process (using its PID) and restart it again.

The following error messages show up when you try to start or stop JBoss respectively while its status is still 'starting':

```
Cannot execute action as another process is holding the lock.
```

```
Cannot stop service 'jboss' while in status 'starting'
```

To kill the JBoss process:

```
kill 9 <JBoss PID>
```

## Why would some users not get email?

Check your dnsserver-<date and time>.log.

James records all errors related to resolving DNS for outbound mail to: `/opt/collabnet/teamforge/james/james-<version>/apps/james/logs/dnsserver-.log`. If you find that some of your TeamForge users are receiving email, but significant groups of others are not, you should consult this log to determine if James is experiencing difficulties in resolving their domain or MX records.

## Why do I get 'Name is Invalid' error when trying to create a project using createProject() method via the SOAP API?

You need to pass certain default parameters to createProject() method.

The default parameters that can be passed to createProject() method are:

```
(sessionId, name, title, description)
where, name = URL Name [in UI]
title = Project Name [in UI]
```

In TeamForge, URL Name ['name'] accepts only lower case and the numeric characters. If the value of 'name' contains any other character than the aforesaid, user will end up seeing the 'Name is Invalid' error and the project creation would fail.

## Why is the password and login shell changed for users on my cvs/svn server?

TeamForge uses local user accounts on the SCM server to provide access to CVS repositories via ssh.


If any local user accounts on the SCM server match user names within TeamForge they will be changed.

If you are planning to use CVS in the SCM server, you should ensure that the local accounts do not conflict with the TeamForge user accounts. Adding a prefix to the local user accounts (local\_username) would be one way to resolve this and prevent the usernames from conflicting. Alternatively, if you are not using CVS repositories, the CVS integration can be removed altogether.

## Why am I not able to see the status of the Postgres in the collabnet startup script?

You may not be able to see the status of the Postgres if the host name of the `HOST_` token is set to localhost in a SaaS multibox setup.

The Teamforge installer fails to add the IP address of the database box to the listen address in the `postgresql.conf` file if the host name of the `HOST_` token is set to localhost in a SaaS multibox setup.

 **Note:** You must add the IP address of the database box to the listen address in the `postgresql.conf` file.

## What does the "psql: could not connect to the server: No such file or directory" error message mean?

This error indicates that the PostgreSQL database server is not running. You need to restart the server.

Use the following command to start the server:

```
service postgresql start
```

## Why do I get a proxy timeout when I try to view certain SCM pages?

If you are getting a proxy timeout error when you try to view a SCM page, you may need to configure the *Apache 2.2 Proxy Timeout* to 300 or less in the `httpd.conf` file.

If you get the following error while attempting to view a SCM page in SFEE:

```
The proxy server received an invalid response from an upstream server.
The proxy server could not handle the request
GET /integration/viewcvs/viewcvs.cgi/ibe-rules/tags/phases/ibe-
rules_09.02.0-Ph-200902_test_20090105/
Reason: Error reading from remote server
```

Configure your Apache 2.2 Proxy Timeout to 300 or less in the `httpd.conf` file.

## Why do search and email server show "Could not connect"?

Typically this means the tomcat container for James and the search service are not running.

You can restart this with the commands shown below. You may need to set the `JAVA_HOME` environment variable to the location of your JDK.

```
sh /opt/collabnet/teamforge/dist/james/james-2.2.0/bin/phoenix.sh stop
sh /opt/collabnet/teamforge/dist/james/james-2.2.0/bin/phoenix.sh start
```

## Why do I get a server status error when I perform a search?

Occasionally, an exceedingly large or complex document causes the search indexing service to abort.

This is typically when all searches in TeamForge return an exid to the user.

Check the server status page and see if the search server is listed as anything other than OK. If it is not OK, then you should restart the search service by logging into the TeamForge application server as root and issuing the following command:

```
/opt/collabnet/teamforge/james/james-/bin/phoenix.sh restart
```

Check the server status page again in TeamForge and ensure that it shows a status of OK. If it shows OK, then searches should now work, and the site will slowly catch up on any indexing requests that were logged while the service was down. If you continue to get exids returned for all searches even with an OK status, then you probably have corrupt search index files.

## Why can't TeamForge send my outbound mail?

If you are unable to send email directly due to *firewall* restrictions, or if mail is being rejected by the application server's IP address, configure TeamForge to send outgoing messages through a gateway mail server.

Configure TeamForge to send outgoing message through a gateway mail server by adding the following to the `<mailet match="All" class="RemoteDelivery">` directive in the configuration file at `/opt/collabnet/teamforge/runtime/james/apps/james/SAR-INF/config.xml`:

```
<gateway>smtp.example.com</gateway>
<gatewayPort>25</gatewayPort>
```

If your gateway mail server requires authentication to send email, you may also add the following directives:

```
<username>username</username>
<password>password</password>
```

## Why does the SOAP service show "could not connect" on the Server Status page when everything else appears to work?

This can be caused by an incorrect host name in `/etc/sourceforge.properties`. Rebuilding runtime will correct this, assuming the hostname is set correctly in the `site-options.conf` file.

This issue can occur when using the `restore.py` script to restore data from a TeamForge instance with a different hostname.

## Why does startup fail or produce errors?

If TeamForge fails to start up, or is starting but is throwing errors on every page, then typically something went wrong during the JBoss bootstrap process.

Fortunately, JBoss logs this process to: `/opt/collabnet/teamforge/jboss/jboss-<version>/server/default/log/boot.log`.

TeamForge writes its startup and shutdown info, as well as any system errors to `/usr/local/sourceforge/log/server.log`. If you encounter a system error while using TeamForge, it is logged here. Additionally, if you see an 'exid' string in the application, the Java stack trace for that exid will be logged in this file.

## Why do I get a TeamForge system error in the project template creation page?

This may be because of a few stale permissions in the project in which you are trying to create the template.

You can resolve this by identifying and deleting the stale records using this SQL.

```
select role_id from role_operation ro left outer join ia_project_association
ia
on (ro.resource_value = ia.id) where ia.id is null and resource_value like
'%prpl%' ;
```

```
select role_id from operation_cluster ro left outer join
ia_project_association ia
on (ro.resource_value = ia.id) where ia.id is null and resource_value like
'%prpl%' ;
```

## How can a user who deleted a document get it back?

It is possible to recover the latest version of a document if someone has accidentally deleted it from TeamForge.

However, note that you can recover the raw document but not undelete it from the app. The user will have to add the document back into TeamForge, and the history of the document will not be recovered. If you need a true undelete so that the history is available, please contact Technical Support and ask for us to schedule our Professional Services group for you.



Let's assume you have a project in TeamForge called Test, and the URL in your browser when you load this project is: `http://your_sfee_server/sf/projects/Test`. Please note the word after projects, this is your project name.

Now, let's assume the file to be retrieved was in a folder called 'foo' in the document manager. Further, let's assume the document was called 'example'. On the database, run the following command:

```
select stored_file.raw_file_id, document.id, item date_last_modified,
       item.title,
       stored_file.file_name, sfuser.full_name from item, folder,
       document, document_version,
       stored_file, sfuser where item.title like '%example%' and
       item.folder_id = folder.id and
       folder.path = 'docman.root.b' and document.id =
       document_version.document_id and
       document_version.stored_file_id = stored_file.id and
       item.is_deleted = 't' and
       item.last_modified_by_id = sfuser.id ;
```

This will return something like:

```
raw_file_id | id | date_last_modified | title | file_name | full_name
```

Examine the rows of output to determine which of the returned results is the document you need, from the title and the name of the file that was uploaded. Once you've determined this, look for the file in `/opt/collabnet/teamforge/var/deleted_files`.

```
yujr0C98B989011083DBFF4F534F | doc1001 | 2007-04-11 08:49:19-07 | example |
foo.txt | Joe User bxma0C98B98901115B9CE99E82C2 | doc1002 | 2007-04-11
08:49:19-07 |
example | test.txt | Joe User rpyf0C98B98901116DD1E40F9F8B | doc1003 |
2007-04-11
08:49:19-07 | example | findme.txt | Joe User ltrq0C98B9890111D76E74920A0E |
doc1004
| 2007-04-11 08:49:19-07 | example | project_logo.jpg | Joe User
```

Let's say the fourth document above was the one you wanted. You can find it at: `/opt/collabnet/teamforge/var/deleted_files/l/lt/ltr/ltrq/ltrq0C98B9890111D76E74920A0E`. Simply copy that file off the system, rename it to `project_logo.jpg` and return it to the user.

## Why do I get a URL "not found" or "moved permanently" error after applying a patch/upgrade?

If you are experiencing a URL "NOT FOUND" or "MOVED PERMANENTLY" error after applying a patch or upgrade, then set Apache ProxyPreserveHost token to ON in the `httpd.conf` file.

If you have applied a patch or upgrade and are now receiving the following error:

```
<The document has moved <a href="https://www.<site>/sf/global/jsp/
buildtime.html"
       format="html" scope="external">here</a>.</p>
<hr> <address>Apache/2.2.3 (Red Hat) Server at www.<site>.com Port 80</
address> </body></html>
Not Found
```

```
The requested URL /sf/sfmain/do/userPicker/projects.pftool//sfmain/do/
listMonitoringUsers/projects.pftool/discussion.announcements was not found
on this server
```

Or if you are trying to add users to a monitoring list, and are receiving the following error:

```
Not Found
The requested URL
  /sf/sfmain/do/userPicker/projects.pftool//sfmain/do/listMonitoringUsers/
projects.pftool/discussion.announcements
was not found on this server.
```

Set the `ProxyPreserveHost` token to ON in the `httpd.conf` file.

## Unable to embed Jenkins page into an iframe. How to fix it?

From the Jenkins version 1.532 and later, the Jenkins page is not displayed in an iframe.

Do you need a way to permit Jenkins to be visible in selected iframes? Here you go.

### To configure X-Frame-Options:

If you try to load content into an iframe and the **X-Frame-Options** header is set to **SAMEORIGIN**, then the content loading is prevented. To override the default **X-Frame-Options** header settings, you have to perform the following:

- You need to have the Jenkins plugin installed to configure the **X-Frame-Options**. You can download the plugin from <https://wiki.jenkins-ci.org/display/JENKINS/XFrame+Filter+Plugin>.
- You must have Jenkins administrative privileges to do this. On the **Jenkins Configuration** page, in the **XFrame Filter Configuration** section, enter the value for **X-Frame-Options Options** as **ALLOW-FROM <ctf domain name>**.

*For example, on the <http://jenkins.maa.relenghub.collab.net/jenkins/configure> page, enter **ALLOW-FROM http://forge.collab.net** for **X-Frame-Options Options**.*

Now, the page can only be displayed in a frame on the specified origin.

## The TeamForge-Binary integration is down after enabling SSL. What should I do?

TeamForge-Binary integration is found to be down soon after enabling SSL on sites which initially had SSL disabled. You must update the base URL, go URL and end point URLs stored in Postgres to access your Binary servers.

Run the following query with the `psql-wrapper` script (`/opt/collabnet/teamforge/runtime/scripts/psql-wrapper`) to update the base URL, go URL and end point URLs.

```
update integrated_application set base_url='https://<url>',
go_url='https://<url>', end_point='https://<url>' where name='Binaries';
```

- 👉 **Note:** In the above query, replace `<url>` with valid base URL, go URL and end point URL for your site. The URLs must use `https` as illustrated above.

## Why do ETL jobs fail post TeamForge upgrade?

ETL jobs can fail due to reasons such as incompatibility between the database and JDBC driver versions and ETL jobs not being able to connect to the Datamart. Try the following solutions.

Pentaho, used by TeamForge for data integration and transformation jobs, recommends using compatible JDBC drivers meant for specific database versions. See Pentaho's [JDBC Drivers Reference](#) for more information.

### If ETL jobs fail post TeamForge upgrade due to incompatibility between the database and JDBC driver versions:

1. Refer to Pentaho's [JDBC Drivers Reference](#) page.
2. Click the JDBC driver reference URL corresponding to your database, Oracle or PostgreSQL.
3. Identify and download the compatible JDBC driver for your database.
4. Replace the JDBC driver found in the following directories with the one you downloaded. (The TeamForge ETL process refers to the JDBC driver available in these directories.)

- /opt/collabnet/teamforge/dist/tomcat/commonlib/
- /opt/collabnet/teamforge/runtime/tomcat\_etl/webapps/etl/WEB-INF/lib



**Note:** You can also refer to [this page](#) for more information about Pentaho-special database issues and resolutions.

### If ETL jobs fail due to unavailable connections to the PostgreSQL Datamart:

Make sure that the following error message is found in `etl.log`:  
 Invalid JNDI connection  
 java:comp/env/jdbc/ReportsDS :  
 FATAL: remaining connection slots  
 are reserved for non-replication  
 superuser connections

If yes, restart the ETL service and restart the failed ETL jobs manually using `./etl-client.py` script in the `/opt/collabnet/teamforge/runtime/scripts/` directory. The ETL jobs should be able to connect to the PostgreSQL Datamart after the restart.



**Note:** If the problem persists even after restarting, contact CollabNet Support.

### Why are statistics charts broken on a wiki page?

Statistics charts such as artifact statistics, task statistics, FRS statistics and document statistics, when added to a wiki page, do not get displayed because the relevant markup formats `[sf:artifactStatistics]`, `[sf:taskStatistics]`, `[sf:frsStatistics]` and `[sf:documentStatistics]` have been deprecated from TeamForge 8.1 and later.

So, if you try to create a wiki page with these values, instead of displaying the charts, the following error messages are displayed:

- For task charts, "The Task Statistics are discarded." is displayed.
- For all the other statistics charts, "The Statistics charts are deprecated viewing from Wiki. It can be accessible from Project Home Page." is displayed.

### Error while modifying custom reports via branding repository. How to fix?

When you create a custom report, change its category via the branding repository and then attempt to view the original report that you created (before modifying the category), you will encounter an error.

Run the following query to update the report:

```
update report set category=<category_name> where id = <report_id> and
surrogate_id = <surrogate_id>;
```

In the above query `category_name` is the ID from the `report_category` table, which you can get by running the following query:

```
select id from report_category;
```

Ensure that the report type in the report belongs to the new category using the `report_type_meta_data` table by running the following query:

```
select type, category from report_type_meta_data;
```

## Manage security

---

Questions about managing security.

Concept definition.

### After switching to ADS authentication, why did the Create button disappear from the user admin section?

When using external authentication such as LDAP, creating users from within the application is disabled. All users must be created via LDAP.

See [login-config.xml](#) for more information.


### Can I block project data from public search engines?

Yes, edit the robots.txt file to specify the pages or directories that should not be indexed by search engines.

If your site has some content that you may not want to be publicly searched for, or if some search engine hits are causing the site to slow down, you can make this change.

Only one robots.txt file can be created for a site. This file can contain a list of url patterns that should not be indexed, against the web crawler name.

The default robots.txt is available in the `SITE_DIR/var` folder. The robots.txt file is accessible without logging into CollabNet TeamForge, via the `domain/robots.txt` URL. You can update and commit the robots.txt file into the branding root directory, and the system then places the file in the `SITE_DIR/var/overrides` folder.

 **Tip:** You must have the commit permissions to create a robots.txt file under the branding repository in the Look Project.

### How can I enforce strong passwords?

You can configure the application to reject passwords that do not meet your security criteria.

To enforce password requirements, place the following lines: :

```
system.password.min-length=5
password.requiresNumber=true
password.requiresNonAlphaNum=true
password.requiresMixedCase=true
```

into:

```
/opt/collabnet/teamforge/sourceforge_home/etc/
sourceforge_configuration.properties
```

Once these lines are in place, restart TeamForge for them to take effect. The above example would require a password of at least 5 characters that must include at least one (1) mixed case letter, at least one (1) number, and at least one non-alphabetic character, e.g. Us3r!

**NOTE:** These settings apply only to new passwords. Anyone in the system currently will be able to continue to use their existing, potentially weak, password. You should force all users to change their passwords after changing these.

### How do I configure Subversion to authenticate against multiple LDAP domains?

For some configurations, a Subversion server may need to be authenticated against multiple LDAP domains. This is possible by modifying the Apache configuration.

This is now possible due to the `mod_authn_alias` module for Apache. The external link for the module contains multiple usage scenarios. You will need to confirm that your Apache has been compiled with the module enabled. (This is the case for CollabNet Subversion binary packages since 1.5.4). If it is compiled as a module, make sure it is enabled via the `LoadModule` directive in your Apache configuration.

Example for configuration usage for authentication against three LDAP servers :

```
<AuthnProviderAlias ldap ldap-US>
  AuthLDAPBindDN cn=ldapuser,o=company
  AuthLDAPBindPassword password
  AuthLDAPURL ldap://ldap-us.company.local/ou=Developers,o=company?sub?
(objectClass=*)
</AuthnProviderAlias>

<AuthnProviderAlias ldap ldap-EU>
  AuthLDAPBindDN cn=ldapuser,o=company
  AuthLDAPBindPassword password
  AuthLDAPURL ldap://ldap-EU.company.local/ou=Developers,o=company?sub?
(objectClass=*)
</AuthnProviderAlias>

<AuthnProviderAlias ldap ldap-IN>
  AuthLDAPBindDN cn=ldapuser,o=company
  AuthLDAPBindPassword password
  AuthLDAPURL ldap://ldap-in.company.local/ou=Developers,o=company?sub?
(objectClass=*)
</AuthnProviderAlias>

<Location /svn>
  DAV svn
  SVNParentPath /opt/subversion/repos
  AuthType Basic
  AuthName "Subversion Repository"
  AuthBasicProvider ldap-US ldap-EU ldap-IN
  AuthzLDAPAuthoritative off
  Require valid-user
</Location>
```

## How do I authenticate multiple LDAP via Apache?

If you need to add multiple OU= values in the LDAP url you must have separate LDAP urls and utilize AuthnProviderAlias to check both LDAP searches.

Use the following AuthnProviderAlias to check LDAP searches.

```
LoadModule authn_alias_module
modules/mod_authn_alias.so

<AuthnProviderAlias ldap ldap-alias1>
AuthLDAPBindDN cn=youruser,o=ctx
AuthLDAPBindPassword yourpassword
AuthLDAPURL ldap://ldap.host/o=ctx
</AuthnProviderAlias>

<AuthnProviderAlias ldap ldap-other-alias>
AuthLDAPBindDN cn=yourotheruser,o=dev
AuthLDAPBindPassword yourotherpassword
AuthLDAPURL ldap://other.ldap.host/o=dev?cn
</AuthnProviderAlias>

Alias /secure /webpages/secure
<Directory /webpages/secure>
Order deny,allow
Allow from all

AuthBasicProvider ldap-other-alias ldap-alias1

AuthType Basic
```

```
AuthName LDAP_Protected_Place
AuthzLDAPAuth_ritative_off
Require valid-user
</Directory>
```

## How does TeamForge authenticate CVS users?

CVS is treated as a special case when managed by a TeamForge site. It is not authenticated in the same way as SOAP API clients.


CVS relies on the Linux operating system to provide access and security. This includes permissions on individual repositories and access to the server itself. To add users, change passwords, create repositories, etc., the TeamForge integration simply changes the appropriate settings in the Linux operating system.

Users access CVS via an encrypted SSH session. To support this, TeamForge creates accounts on the Linux server that hosts the CVS repository. A typical CVS repository is created in the `/cvsroot` directory and is owned by root, with a group assigned by TeamForge. To gain access to a repository, TeamForge will add a user to the appropriate system group.

When TeamForge detects that a user's password has changed, it changes the password for that user on the Linux server too.

Users also have the option to use SSH keys or Kerberos tickets.

When a user is added to the Linux server, the login shell is `cvsssh`, which limits their activities to CVS commands.

 **Important:** Do not expose `cvspserver` (the TCP protocol over port 2401) either internally or to the Internet if there is any way you can avoid it.

### Password changes under internal authentication

To set the password for the user at the operating system level, TeamForge needs to have the clear-text version of the users password. The only time TeamForge has this is when the user uses the Change Password form in the Web UI. This is because the database-stored version, as an MD5 Password Hash, is a one-way encryption and can't be decrypted.

On a successful password update, TeamForge makes a SOAP call to the integration server that manages CVS. For this reason, the integration server must be SSL-enabled.

### Password changes under external authentication

When a password change happens in an external authentication system, TeamForge does not immediately know that the password has changed. TeamForge needs a way to detect that the password has changed.

To accomplish this it keeps a copy of the last password the user successfully logged in with as an MD5 Password Hash in the same database table and field that it normally uses for Internal Authentication.

Now that TeamForge has a reference point, it still needs a clear-text copy of the password to make the change at the Linux operating system level. The only time this is available is when the user logs into TeamForge via the Web UI or SOAP API. So upon a successful login TeamForge compares the password to the encrypted one. If it is different it tells the Linux operating system to change the users password and then saves it in the database (as an MD5 Password hash).

Until the user logs into TeamForge, the CVS server will still have and accept the old password. There is no CVS server-side way to trigger a password update, unless an alternative method is used, such as LDAP or Kerberos.

### Alternative Authentication

Because users use SSH to access a TeamForge-managed CVS server, it is possible to configure SSH to accept other authentication features such as SSH keys and Kerberos tickets. It is even possible to disable the use of passwords and require the use of other alternative methods.

- TeamForge supports SSH Keys natively. The user uploads their public key into their profile under “My Settings” in the TeamForge Web UI. The key will automatically be copied to each CVS server that TeamForge manages.
- If TeamForge is using External Authentication and the method is Kerberos, then SSH can be configured to use the same Kerberos server. This allows users to use Kerberos tickets for CVS operations.


## LDAP

Linux supports LDAP as an authentication source. If TeamForge is using External Authentication and the source is LDAP, then SSH can also use that same source. When you do this, passwords and user account status are observed in real time instead of as a mirror of TeamForge.

## Can the users be forced to change their passwords at first login?

Yes, as a site administrator you can configure the CollabNet TeamForge site options to force the users to change their passwords at first login.

Setting the `REQUIRE_USER_PASSWORD_CHANGE` attribute as `true` in the `site-options.conf` file enforces password change on first login into CollabNet TeamForge.

 **Tip:** You can not force password change on a user who had self-created the user account, or if a password-request had been raised for the user or if an administrator had reset the login password for that user.

## Does TeamForge work with LDAP?

Yes, you can have your TeamForge installation authenticate against an LDAP server.

This is handy when users want to use a variety of different resources without having to maintain credentials for each one separately.

### Overview


CollabNet TeamForge is a JBoss2 based application and relies on the JBoss JAAS service for user authentication. This enables a TeamForge site to authenticate users internally or externally.

#### Internal user authentication

Out of the box, TeamForge relies on its local database to manage user accounts. This includes username, password, full name, email address and a variety of other meta data values. Passwords are stored in the database using the standard MD5 Password hashing algorithm<sup>1</sup>. The database is only accessible by the application itself and a user with root access to the physical server. While running in this default configuration users are allowed to change their passwords in TeamForge, and any user with site administration privileges can create and approve new user accounts.

#### External user authentication

The JAAS service comes with several standard providers that allow TeamForge to be integrated with services such as LDAP, Active Directory and Kerberos. The JAAS service allows more than one source to be configured in the event several sources are needed.

 **Note:** It is possible to use both types of authentication with a single TeamForge installation. See your CollabNet representative for details.

To ensure that you are not locked out of your site, the site administrator account is always validated by TeamForge, not by LDAP.

LDAP accounts must conform to the TeamForge rules for user names and passwords. For example:

- If a password is used in LDAP that is shorter than the minimum allowable password length in TeamForge, you cannot create the user in TeamForge.

- A user name that starts with a special character, such as an underscore, will not be accepted by TeamForge, even if it is valid in LDAP.

(For detailed TeamForge user name and password rules, see [Create a new user account](#) on page 336.)

### How is life different for the user under external authentication?

- When you turn external integration on, every user account (except the site administrator account) must have a matching LDAP entry to log in. This may require changing some existing accounts to match their corresponding LDAP records. (Accounts created after LDAP is in place are validated with the LDAP server when they are created, so you don't have to worry about this.)
- Every login attempt (Web UI and SOAP access) is passed to the external provider. This means that any changes to the user status in the external system take effect immediately. Users who have already logged in and have valid sessions are not affected.
- When TeamForge is using internal authentication, a site administrator can change a user's password. This is disabled for external authentication.
- Under external authentication, passwords can't be changed in the TeamForge web UI. Users have to use the interface provided by the third-party authentication source to change their password. Such password changes are available immediately to TeamForge for the next login attempt.
- Site administrators can no longer create user accounts. The end user must create their own account by logging into TeamForge just like a user who already has an account. At that point TeamForge detects that a new account needs to be created and presents the new user with a registration form, which requests the user's password in the external authentication system. On submit, TeamForge verifies the user account with the external system, and only if the username/password is verified does TeamForge create the new account.
- Once a new user has created their account, TeamForge can optionally be configured to put every new account in a pending status so that a site administrator can approve the new account. By default, new users will have immediate access to the system.

### LDAP for source control

LDAP is integrated into your TeamForge source control services.

- For Subversion, the integration server queries TeamForge as needed.
- CVS authentication is not managed directly by LDAP, but each TeamForge user's SCM password is synchronized automatically with the user's LDAP password upon logging into TeamForge.

### What can go wrong?

When TeamForge is configured to authenticate against an LDAP server and the LDAP server is down, all TeamForge authentication is disabled until the LDAP server is restored.

If a user does not exist on the LDAP server, or is deleted from the server, that user cannot log into TeamForge.

## Why do I get the "Invalid command 'AuthLDAPAuthoritative'" error when I try to set LDAP for SVN users?

The invalid command `AuthLDAPAuthoritative` error may occur if you need to upgrade Apache from version 2.0 to 2.2.

CollabNet Subversion 1.5 is bundled with the latest version of Apache (currently 2.2.x). It includes the module `mod_authnz_ldap` and does not include `mod_auth_ldap`. Hence compatibility issues arise due to missing directives. Upgrade your Apache version to 2.2 if you get the following error when trying to install CollabNet SVN:

```
bash-3.00# /etc/init.d/collabnet_subversion start
Starting CollabNet Subversion:
Syntax error on line 29 of
/etc/opt/CollabNet_Subversion/conf/collabnet_subversion_httpd.conf:
Invalid command 'AuthLDAPAuthoritative',
```



```
perhaps misspelled or defined by a module not included in the server
configuration
FAILED
```

## How does TeamForge handle multiple redundant LDAP servers?

When configuring LDAP authentication for a TeamForge instance, there may be a business need for using multiple LDAP servers.

Follow the guidelines below for configuring.

The additional LDAP servers can be added to the `java.naming.provider.url` option in `login-config.xml`:

```
login-config.xml:
  <module-option name="java.naming.provider.url">
    ldap://primary/ ldap://secondary/</module-option>
```

Once the primary and secondary servers have been defined, they will be consulted in order of definition for every authentication request. First the primary, and if the primary fails, then the secondary. This prevents specifying multiple servers for round-robin handling of authentication, but it can still be used for redundancy needs.

## What user activities are tracked?

In case of a data security compromise, a record of who is performing what activities will help resolve some of the security issues.

Typically web servers log every page (or URL) being accessed, including the IP address of the user, date and time of access, etc. These logs are very useful in tracking the source of any security violations that may occur.

CollabNet TeamForge auditing tools are a powerful way to track unwanted and/or unauthorized changes within the system.

## What is a CERT advisory?

CollabNet Product Support monitors the CERT coordination center (<http://www.cert.org/>) for notification of vulnerabilities or exploits against applications that CollabNet TeamForge provides.

If CollabNet Technical Support identifies an advisory that may indicate potential challenges for users who have deployed CollabNet TeamForge, Support proactively releases a notification and a statement of action. CollabNet will provide product updates as it deems appropriate or necessary.

## J2EE Architecture and security

CollabNet TeamForge is a J2EE application that employs three-tier architecture to provide a secure environment for mission-critical data.

In a multi-tier architecture, access to each tier is restricted to the tier above it, effectively securing the tiers behind the firewall. For example, while clients (users accessing the system through a web) access the web server, they neither have access to the application and backend servers nor are they aware of their existence.

Similarly, the web server itself does not have access to the backend servers (database, SCM, mail etc.)

Exceptions to this rule include:

- Direct client access provided to the SCM servers. SCM servers are accessed across the firewall typically through SSH protocol (for CVS), or HTTP or HTTPS (for Subversion). SCM server data is also accessible in a view only mode through the web interface.
- Clients must have access to the mail server for posting messages to mailing lists.
- Mail server must have access to deliver messages across the firewall.

Clients can also access the SOAP APIs through the web server. The web server in turn forwards SOAP requests to the application server for processing.

## With self signed certificates in place, what is the recommended protocol (SSH or HTTPS) to use while cloning a Git repository from a TeamForge Git server running on RHEL/CentOS?

When using a self-signed certificate on a Teamforge Git server, you cannot clone a repository using the standard git client on RHEL/CentOS.

In RHEL/CentOS, the Linux certificates used by git and other tools are stored in the `/etc/pki/tls/certs/ca-bundle.trust.crt` file. This file is managed by the RPM package system. If a certificate is added to the `ca-bundle.trust.crt` file, Trusted Root Certification Authority updates are not installed automatically which in turn leaves the system vulnerable to potential attacks.

An alternative way for adding trusted certificates is available at <http://stackoverflow.com/questions/9072376/configure-git-to-accept-a-particular-self-signed-server-certificate-for-a-partic>. However, it appears that some tools do not read files outside of the main bundle.

Therefore, it is recommended to use SSH protocol while cloning a git repository from a TeamForge Git server running on RHEL/CentOS.

## Roles and permissions

---

Questions about setting permissions and using roles in TeamForge.

Concept definition.

### Can I block a user's access to source code?

From time to time, you may need to prevent a user from using a source code repository.

SCM access can be withdrawn from a user by doing one of the following:

- Change the user's access to specific paths in the repository. (Subversion only.)
- Remove the role providing the desired source code permissions from the user.
- Remove the user from the project.

How this works depends in part on whether the repository is managed by TeamForge.


- If the source control server that hosts the repository is managed by TeamForge, the user's access is removed immediately.
- In the case of a CVS server that is not managed by TeamForge, all TeamForge administrators are sent an email requesting that the change be made manually. The request appears in the **SCM Access Requests** section of the **Integrations** page for approval.

Both the project administrator and the user receive email notifications when the change is made.

### Can I assign a role to all users of the site at once?

If you are a site administrator in CollabNet TeamForge, you can assign a role to all the system users (non-site administrators) at one go.

Now, CollabNet TeamForge provides a system created user group called "All Users", including all site users. All role assignments that can be made to a user group can be made even for the "All Users" user group.

 **Note:** The "All Users" group members can not be manually added, edited, viewed or deleted. This group is automatically updated each time a user is added or deleted from the site.

### Can I request a role?

If you are a project member in CollabNet TeamForge, you can request a role in your project.

You can use the **Project Home > My Roles** page to request for a role. Your request is submitted to the project administrator for approval. You will receive an email notification when your request is either approved or denied.



**Note:** Based on your project administrator's discretion, some role requests may be granted immediately, while the other role requests may need approval.

## Why don't I have access to the "Reported in Release" information in my artifact?

If you cannot see the "Reported in Release" information in your artifact, ask your project owner to add the role required to obtain access to this information.

Users with tracker artifact submit/view/edit permissions will not be able to see "Reported in Release" information in an artifact, as it is directly associated with the File Releases section.

You can request that the project owner add a role permission of View Only for All Packages in File releases for you to obtain access.

## How do user permissions in CollabNet Enterprise Edition map to permissions in TeamForge 17.4?

Permissions in TeamForge 17.4 are different from permissions in CollabNet Enterprise Edition.

### Permission mapping details

Care has been taken to map the permissions of a CollabNet Enterprise Edition role to an equivalent role in CTF. However, after migration it may be similar, but not same. See these details to get a better understanding of the mapping between CollabNet Enterprise Edition and TeamForge 17.4 permissions: [User permissions in CEE and CTF 17.4 projects](#) on page 388



**Note:** Site-wide roles in TeamForge 17.4 are equivalents of Domain Roles in CollabNet Enterprise Edition. Global project roles in TeamForge 17.4 are the equivalents of Project roles with "All projects" visibility in CollabNet Enterprise Edition.

### Exception reports

In cases where it is not possible to retain the permissions mapped to a role in CollabNet Enterprise Edition, the exceptions are reported and the exception report is provided for your reference in TeamForge 17.4. Each project will have an exception report with the change details at the project level. The exception report for the project is placed in the "Documents" tool under the `CollabNet Enterprise Edition_migration_exception_report` folder. The report file is named as `<project-name>-REPORT.txt`.

The Look project contains another exception report with site-level details. This would also have exception report of individual projects in the system. The site level report is placed in the look project under the same folder as `GLOBAL-REPORT-CONSOLIDATED.txt`.

Please go through the global exception report for the site, and see that you are okay with it, or make necessary adjustments as needed. To use the exception report effectively, see [How to read an exception report?](#)

### Resource Patterns

Resource patterns associated with the CollabNet Enterprise Edition roles are not migrated. This means that the users belonging to these roles will have access to the complete repository after migration, depending on their access level in CollabNet Enterprise Edition. For example, if a role had "VersionControl - Modify" access to only "trunk/src", post-migration, the role will have commit access to the entire repository. Similarly, if the role in CollabNet Enterprise Edition had "VersionControl - Read" access to only "trunk/docs", post-migration, the role will have view access to the entire repository.

TeamForge 17.4 uses Path-based permissions (PBP) as an alternative. You can transfer the regular expressions to PBP to control the repository access. Setting PBPs is easier when compared to using the Resource Patterns; however, PBP has some limitations compared to CollabNet Enterprise Edition's regular expressions.

PBP limitations include:


1. PBPs are not inherited.
2. Site-wide and global project roles do not allow PBP associations with permissions.

See these details to get a better understanding of the mapping between CollabNet Enterprise Edition Resource Patterns and TeamForge 17.4 permissions: [TeamForge 17.4 Permissions for Resource Patterns](#) on page 399

The exception report contains the permission-resource pattern mapping that was dropped, against the affected role.

### Role associations with **\*\*Host/\*Domain users**

If your project had any role associations with \*Domain users and \*\*Host user's user groups, they are not retained upon migration. At this point, TeamForge 17.4 doesn't have an alternate mechanism for this.

 **Note:** If any such associations are dropped, role association of default user groups in Domain and project level are captured in domain and project level exception reports respectively. The default user group's subscription to discussions are captured in respective project's exception reports.

### Registered User role and permission of a logged in user


Permission of a CollabNet Enterprise Edition's Registered user role are mapped to Default Access Permissions of individual projects. The difference here is that, there is no central place that controls the access rights but it has the advantage that each project can specify different permissions.

In cases where the CollabNet Enterprise Edition permissions are not retained, they are exception reported. Please refer the global exception report for details.

### Anonymous access

You may skip this section if guest access was not enabled in your CollabNet Enterprise Edition site.

There is no central control in TeamForge 17.4 to enable/disable guest access. Each project can specify what an anonymous user can access through the DAP of the project. Permissions of CollabNet Enterprise Edition's Anonymous Guest role are mapped to Default Access Permissions of individual projects. If guest access was enabled in your CollabNet Enterprise Edition site, please look at the global exception report to see whether there are permissions that changed during migration.

 **Note:** It is recommended that you try out this in stage before doing in production.

## Who can access an application?


Application permissions help you minimize the need to create and assign many similar roles for individual users. Instead, you can permit or restrict access to individual applications within the project for whole classes of users.

Application permissions supplement role-based access control (RBAC.) For each application's concepts, documents, file releases, trackers, and discussion forums, you can assign permissions globally based on user type.

For example, if you know that you want all project members to be able to view and submit to all project trackers, you can specify this application permissions. You need to configure these settings only once. All current and future project members will be able to view and submit to all trackers without having the tracker view/submit permission assigned to them individually via a role.

Before you do this, you should have identified your project as private, gated community, or public. Configuring permissions is a finer-grained level of control that operates within this hierarchy of project types.

Some applications may be invisible to some users based on the roles you assign. If you give a user a role that does not grant access to a particular application, that user cannot see the button for that application in the Web interface. (However, if that user also has some other role that does grant access to that application, the user can see the application button.

 **Remember:** A user's license type also influences what the user can see and do on your site. A user's license type supersedes any role assignments. Ask your site administrator how many licenses of each kind are available for your users. For more information, see [How do TeamForge licenses work?](#) on page 948.

## Exceptions

- If a user has an SCM license, that user can see only the tools that support the core source control functions of the site. The Tracker, Documents, and Tasks tools are not visible.
- If a user has any level of permission on a tracker or a task folder, that user can see the **Reports** button.
- If a user has tracker administration permission on any tracker, or task administration in any task folder, that user can see the **Project Admin** button.

## User classes

These are the classes of users to which you can assign application permissions:

User class	Description
All users with Role Permissions	Only project members with appropriate RBAC permissions.
All project members	All project members.
All project members and unrestricted users	All unrestricted users, whether or not they are project members, plus all project members.
All logged-in users	All restricted and unrestricted users (all logged-in users,) whether or not they are project members.
All users	All users, whether or not they are logged in or have CollabNet TeamForge accounts.

## Restrictions by type of site

On some types of sites, you can't assign application permissions to certain classes of users. In such cases, you must use role-based access control (RBAC) permissions.

- On a private site, you cannot set application permissions for these classes of users:
  - All project members and unrestricted users
  - All logged-in users
  - All users
- On a "Gated Community" site, you cannot set application permissions for these classes of users:
  - All logged-in users
  - All users

## Can I delete an item if I have 'Administer' permission?

No. You cannot delete an item if you have the administer permission.

The administer permission does not allow you to delete an item. You can create, edit, submit, and view items if you have the administer permission. For deletion, you must have the delete permission.

## Can I disable creation of user accounts?

As of TeamForge 4.1 SP3, it is possible to disable the creation of new accounts by users so that only a 'site admin' can create new users.

To enable this mode of operations, simply add the following line to `/opt/collabnet/teamforge/sourceforge_home/etc/sourceforge_configuration.properties`:

```
sf.disableUserSelfCreation=true
```

Once this line is in place, restart TeamForge for it to take effect.


## Why do I see the project home page though I lack the necessary permission?

This has been set so that when users are provided access to a project, they do not see an empty page.

Irrespective of whether the view permission on project pages is checked, users will be able to see the contents in the project home page.

## Why do you need the authentication and authorization plugin for Hudson?

The Authentication and Authorization plugin allows you to set up your Hudson installation to authenticate against a CollabNet server and specify access control for CollabNet users.

 **Note:** Authentication and authorization are independent actions. You could set up your Hudson installation to authenticate against a SourceForge or TeamForge site, but not use that CollabNet site for authorizing users. Authorization is available only with CollabNet TeamForge 5.2, but authentication is possible with earlier CollabNet SourceForge Enterprise versions as well.

### Authentication

Authentication determines user names and passwords. You establish user credentials when you enable your Hudson site to use the CollabNet security realm.

### Authorization

Authorization determines what users can do on the Hudson site.

Your Hudson server can be shared between many CollabNet projects, and you can grant TeamForge users permissions at the site level. You specify site-level permissions when you configure the site to use a CollabNet server for authorization.

A job on your Hudson server may be involved with more than one TeamForge project. For example, a job that builds software can pull in source code from multiple project repositories. For the purpose of authorization, however, each job is associated with one CollabNet project, and you can give users project-level roles.

When your Hudson site is set up to use CollabNet authorization, TeamForge project administrators can assign these roles to project members:

- Hudson Build/Cancel
- Hudson Configure
- Hudson Delete
- Hudson Promote
- Hudson Read

Users get the highest permissions they are entitled to. For example, if a user is part of a group that has administration permissions for the Hudson site, that supersedes any Hudson-related role the user might be assigned within a specific CollabNet project.

## Can I control user access to an integrated application?

TeamForge can integrate the permissions scheme of a separate application into the TeamForge role-based access control system.

To look at how this works, we'll use the Pebble blogging tool as an example. Pebble is an application that you can quickly integrate with TeamForge.

Pebble brings with it a set of pre-determined roles that you can assign to project users. The roles are defined in the XML application configuration file.

### Blog Reader


You can only read blogs and make comments, the comments are sent for moderation.

### Blog contributor

You can add blog posts, but they will be sent for moderation.

<b>Blog publisher</b>	You can add blog posts, moderate comments and blog posts.
<b>Blog owner</b>	You can do all that a Blog publisher does as well as change the blog properties and security options.

Any site user with one or more of these roles can see the **Pebble Blog** button in their project toolbar. Clicking that button allows them to operate Pebble according to their access rights.

 **Note:** Site Administrators don't need any specific permissions; they have all permissions on all projects on the site.

## How does inheritance work?


Site or project administrators can create hierarchical relationships between projects so that one project can inherit members, roles and permissions from a parent project.

When you define users, user groups and roles with specific permissions in one project, they can be inherited in one or more subprojects. This helps you avoid duplicating the effort of defining users, user groups and roles across projects.

You can still, if required, create roles specifically for this project and add direct members to any project. The direct members can be assigned inherited roles. The inherited members can be made direct members and/or assigned direct/inherited roles too.

While inheriting roles, only the permissions associated with all (top-level) folders are inherited.

In a subproject, you can select the inherited project members from the **Assigned to** lists or from the user picker, as the case may be. Inherited users may not be part of these lists if their role inheritance was prevented in the parent project.

 **Note:** At the time of role creation, you can choose to allow or disallow the role inheritance into private subprojects.

## How does inheritance work for project groups?

Site or project administrators can create project groups and add member projects to the group, to manage several projects as a single unit.

Moreover, hierarchical relationships may exist between projects so that one project can inherit members, roles and permissions from its parent project.

Quite like a project can inherit roles and permissions from its parent project; projects can inherit only permissions via the project group.

Project Groups permissions when granted to users via site-wide roles, also allows the users to access project groups in all cases. Prevent inheritance option for roles does not apply in those cases.


The project in which a user gets permissions through role assignment in project groups appears on the **My Workspace > Projects > All Projects** page.

The roles created specifically for a project group are not made available as inherited roles in the group's member projects for assignment. Project group roles can only be assigned directly in the project group context. These roles are not requestable.

## Can I create new user accounts as "unrestricted?"

You can make all new accounts be "unrestricted" by default instead of the more secure "restricted."

This is controlled by the `USER_ACCOUNT_RESTRICTED` variable in the `site-options.conf` file. Change that value to false and restart runtime.

 **Note:** This does not change any existing accounts. If someone was restricted before this flag was turned on, they remain restricted until the site admin edits their account.

## Why can't Oracle connect to my TeamForge installation?

The simplest way to correct this is to overwrite the `.jar` included with TeamForge with the one from `$ORACLE_HOME`.

TeamForge uses the thick Oracle JDBC driver, which has two parts. One of these is provided by TeamForge, the other is in `$ORACLE_HOME`. If these two components are incompatible, TeamForge will be unable to make a connection to the database.

Follow these steps to overwrite the `.jar` included with TeamForge with the one from `$ORACLE_HOME`:

```
cp $ORACLE_HOME/jdbc/lib/ojdbc14.jar
   /opt/collabnet/teamforge/jboss/jboss-3.2.6/server/default/lib/
```

A restart of the application will be required to use the new `.jar`.

## Why can't I move these artifacts into this planning folder?

To avoid confusion, you must observe a few basic rules when you assign artifacts to planning folders.

When an artifact in a planning folder has a child artifact, the child artifact can only be assigned to the same planning folder as the parent artifact, or a sub-folder of that folder. For example:

- Before you assign an artifact to a planning folder, make sure the artifact and all its children are assigned to a single planning folder and its sub-folders.
- Don't assign an artifact to a planning folder that is above the artifact's parent artifact in the planning folder hierarchy.
- Before promoting a planning folder to a higher level in the hierarchy, make sure its member artifacts (and its members' parent artifacts) are all assigned to the same planning folder or a sub-folder of it.

An artifact can only be assigned to a planning folder in its own project.

As a best practice, consider instead creating a matching artifact in this project and associate the two artifacts.

Before you assign an artifact to a planning folder, be sure you have permission to edit every artifact affected by the move.

## Who can access a project?

You control access to your CollabNet TeamForge project by a combination of project settings, membership rules and user restrictions.


## Who should I allow into my project?

To decide how to control access to your project, think about what the project is for and who will be using it.

Consider these elements:

- The project's access setting.

A project can be public, private, or gated community.

 **Note:** The site administrator can change the default project access permissions.

- The project's membership.
- Each site user's user type.
- Each site user's license type.
- Each TeamForge user's user type.
- Any parent projects from which members, user groups or roles are inherited.
- Any subprojects that inherit members, user groups or roles from this project.

### User type

Users can be restricted or unrestricted.



- Restricted users can access only public projects and projects of which they are members.
- Unrestricted users can access all projects except private projects of which they are not members.

### License type

Users can have an ALM license or an SCM license.

- An ALM license enables the user who holds it to use the full range of TeamForge features: both the core source-code management tools and the extended application life cycle management functionality.
- An SCM license enables the user who holds it to use the core TeamForge source-code management tools.

License type supersedes user type. For example, if you give a user an SCM license, and then declare that user an unrestricted user, the user can see only the core source code management tools in any project they can access.

### Project access setting

A project can be private, gated, or public.

- **Private** - Private projects can only be accessed by project members. Private projects do not appear on the **Home** page, in the **All Projects** list, or in search or reporting results to users who are not project members.


Create a private project when you want to strictly limit project access.

- **Gated** - Gated community projects can be accessed by project members and by unrestricted users. As with private projects, gated projects are not visible to users who are not allowed to access them.

Create a gated community project when you want to exclude restricted users, but do not need to exclude other, unrestricted users. For example, your organization might wish to designate all contractors or outsourced staff as restricted users. They will not be able to see any gated community projects, but all of your full-time, regular staff will have access.

- **Public** - Public projects can be accessed by all users.

Create a public project when you have no need to restrict access.

-  **Note:** Project access is not the same as project membership. Project access allows a user to see the project in the **All Projects** list, visit the project home page, and browse selected project data. A restricted user may be able to access a project without being a project member.

This table shows which user types can access projects with each project access setting.

Project access type	Project member (Unrestricted user)	Project member (Restricted user)	Non-project member (Unrestricted user)	Non-project member (Restricted user)
Private	Yes	Yes	No	No
Gated	Yes	Yes	Yes	No
Public	Yes	Yes	Yes	Yes

### How do I give read-only anonymous access to cvs/svn repository?

Set the default access permissions of the project to public and allow Source Code to view permission for all or specific repositories to All users while restricting other permissions.

To give read-only anonymous access to cvs/svn repository within a project while still restricting write access, you could set the default access permissions of the project to public with Source Code view permission open to all users, while restricting other permissions to specific user classes.

The steps to set this up are as follows:

1. Go to 'Project Admin' tab in the required project.
2. Click Permissions in the left nav.
3. Click on Default Access Permissions tab.

4. Set Project Access Permissions as Public.
5. Under Application Permissions, choose All users from the drop-down for Source Code View permission.
6. For other application permissions, choose a user class based on your access requirement.

**Note:** In Step 5 above, it is possible for you to choose whether you want to give View access to All users for all repositories or a specific repository.

## Who can access source code?

Project administrators can give project members specific kinds of access to a whole Subversion repository or any path within that repository.

### Overview

You can specify access permissions for users with a given role at the repository level, or at the path level within a repository.

- When you set a permission for a role on any directory in the repository, all directories and files under that directory get the same permission.
- When you set a permission on an individual file in the repository, there is no effect on the permissions assigned to paths above the level of that file. (A file is just a path that ends with a file name.)

How you use path-based permissions will depend on whether you view permissions primarily as a way to grant access or as a way to restrict access.

#### Full access, with exceptions

Give your company's own employees commit access to your whole source code base, while allowing developers from contracting firms to commit only to those parts of the code base that they are expected to work on.

#### No access, with exceptions

Assign all developers "No access" by default, then assign each type of developer access to certain directories and files according to their responsibilities.



**Note:** You can control access to a path or to an individual file. This is different from normal Subversion checkout and commit operations, which are performed on directories but not individual files.

### No Access

When you deny all access to a repository for a role, users with that role cannot see that the repository exists, except if:

- The role has "View and Commit" access to some directory within the repository. In this case, users with this role can see the directories that contain the directory they have access to.
- The user has another role that grants access to some part of the repository.



**Note:** An individual user can have multiple roles. When two roles have permissions that conflict with each other, the role with the more expansive permissions takes precedence.

### View Only

Users with a role that has "View Only" access to a path can browse the contents of the repository on the Web site or by connecting directly to the repository from a client, such as Tortoise, CollabNet Desktop for Eclipse, or CollabNet Desktop for Visual Studio.

### View and Commit

Users with a role that has "View and Commit" permission to a path can browse and download code, and can also check code into the repository.

## Wildcard-based access control and path-based permissions

You can create rules that use wildcards to control access to specific paths in a repository. See [Wildcard-based access control and path-based permissions in TeamForge](#) for more background information about wildcard and path based permissions.

### Scenario 1

Controlling access to a specific path in a branch using authz rules can be very time consuming and inefficient. Assume you have three branches, [/branches/foo/build], [/branches/bar/build], and [/branches/baz/build]. If you create authz rules, you may have to write a separate rule for every branch in the repository, let alone the fact that you need to write such rules for a branch you may create in the future.

```
[/branches/foo/build]
@build = rw
@dev = r
```

```
[/branches/bar/build]
@build = rw
@dev = r
```

```
[/branches/baz/build]
@build = rw
@dev = r
```

Instead, you can write a rule using wildcards such as:

```
/branches/*/build
```

Later, you can create and assign roles to users such as "build engineers" and "developers" with "read-write" and "read-only" access permissions respectively.

### Scenario 2

In this scenario, assume that there is a particular file or folder pattern that needs access control in a repository. For example, you may want to restrict all users but release managers from modifying .iso files. It is impossible to define such a rule using the authz file.

With TeamForge, you can write a rule that partially uses wildcards in file or folder names such as:

```
/trunk/build/*.iso
```

This rule applies to all files and folders that end with .iso under the path /trunk/build.

### Scenario 3

In this scenario, assume that you want to control access to a particular folder no matter where the folder is in a branch. For example, you may want to control access to the "build" folder wherever it is in a repository. You can write a rule using wildcards such as:

```
/**/build
```

The "\*" matches any number of folder levels in a repository. For example, this rule applies to the following paths:

```
/trunk/build
```

```
/branches/feature1/build
```

```
/trunk/external_module/build
```

```
/build
```

In addition to these scenarios, you can use wildcards to create rules that suit your requirement. A few examples of how you can create wildcard-based rules:

- `/**/*.iso` - Match any `.iso` file anywhere in a repository.
- `/branches/RB*` - Match any branch if the name starts with `RB`.
- `/branches/*/*.txt` - Match all `.txt` files one level under any branch.

### Notes about path-based permissions

If two paths have different permissions, the permissions on the lower-level path take effect. For example, consider a role that has "No Access" set for the path `/branches/version3/users`, but has "View and Commit" access to `/branches/version3/users/vijay`. A user with this role can:

- Check code in and out of the `vijay` directory.
- Click down through all the directories in that path, including `users`. (Directories that are not included in the user's permissions are not shown.)

Users with this role cannot:

- Check files in and out of the `users` directory.
- Monitor commits to `users`.
- Execute Subversion copy, move or delete operations that involve resources in the `users` directory (or any other paths where the user has View Only or No Access).

### Commit messages

When users monitor a repository, they receive commit announcements by email that include the resources that their role permits them to see.

Users with a "No Access" role on a repository cannot monitor that repository to receive commit messages by email.

By default, commit emails provide all the details about the commit that a logged-in user can view in the **Source Code** application. A repository owner can elect to have the notification emails omit most of the detail and provide only the commit ID and the committer's user name.

### Toolbar button

If none of the available permissions (**View Only**, **View and Commit**, or **Path-based Permissions**) is selected for any repository, and none of the options under **Source Code Permissions** is selected, users with this role do not see the **Source Code** toolbar button.

## Can I ensure that only site admins can create projects?

You can do this using Velocity.

Create `/opt/collabnet/teamforge/sourceforge_home/templates/body_footer.vm` and populate it with this:

```
#set($superUser = ${PAGE_INFO.isSuperUser()})
#if($superUser == false)
<script>
```

```

if (document.getElementById("createProject") != null) {
var e = document.getElementById("createProject").parentNode.parentNode;
while (e.firstChild) {
e.removeChild(e.firstChild);
}
}
</script>
#end

```

Save the file, and TeamForge will start hiding the button for non-site admins immediately.

## How do user roles work?


Project administrators can define specific access permissions for individual project members. They do this by using global project roles and/or creating roles and assigning the roles to project members.

Under role-based access control (RBAC), permissions are not assigned directly to an individual user. Instead, each user has the permissions that are attached to any role that is assigned to that user.

A project member can be assigned multiple roles.


In CollabNet TeamForge, site or project administrators assign roles to the site users or project members. Besides this, a project member can submit a role request to the project administrator. The project administrator can approve or reject such requests.

When you define users, user groups and roles with specific permissions in one project, they can be inherited in one or more subprojects. This helps you avoid duplicating the effort of defining users, user groups and roles across projects.

 **Note:** Permissions are cumulative. If a project member is assigned a role that provides a specific permission, and another role that does not, the user has that permission.


A role defines these things:

- The applications that project members with that role can and cannot access.
- The resources on which project members with that role can use the applications.
- The actions that project members can take in each application and on each resource.

 **Note:** If a user has an SCM license, that user can see only the tools that support the core source control functions of the site, even if the user has a role that would otherwise grant access to other resources.

When a user's roles do not include access to an application or resource, that application or resource is not visible to that user. For example, imagine that you are assigning roles to Jason, a software developer. Jason needs to check source code in and out in order to fix bugs, develop features and create software releases. However, Jason does not need access to the wiki. If you set up Jason's roles according to those requirements, Jason's experience is like this:

- On any page in the project site, Jason can see and click the **Trackers**, **Source Code**, and **File Releases** buttons along the top of his screen.
- Jason does not see the **Wiki** button.
- If someone sends Jason a link to a page in the Wiki application and Jason clicks the link, he gets an error message. (The message does not specify whether the page exists or not.)
- When he accesses the project directly from Eclipse or Visual Studio, Jason can expand the project node and browse the **Trackers**, **Source Code**, and **File Releases** nodes, but not the **Wiki** node.

 **Remember:** A user's license type also influences what the user can see and do on your site. A user's license type supersedes any role assignments. Ask your site administrator how many licenses of each kind are available for your users. For more information, see [How do TeamForge licenses work?](#) on page 948.

## Applications

An application is a collection of related features designed to enable a user to collaborate on particular kinds of tasks. For example, the Documents application helps users create documents, share in document reviews, and publish documents, among other things.

In the Web interface, each application is represented by a button in the navigation bar at the top of any project page. A given user can see the buttons corresponding to applications they have access to by virtue of the roles assigned to them.

Applications are also known as "tools."

## Resources

- The tracker application might contain a bugs tracker and a feature request tracker. These are the tracker resources.
- A project can contain multiple SCM repositories. These are the SCM resources.

## Permissions

<b>View</b>	Allows users to view and download items, but not to create or edit items, administer folders, or edit application settings.
<b>Create or Submit</b>	Allows users to create new items, but not to edit items, administer folders, or edit application settings. Users with the create or submit permission also have the view permission.
<b>Edit</b>	Allows users to edit items, but not to administer folders or edit application settings. Users with the edit permission also have the view permissions.
<b>Administer</b>	Allows users to create and edit items, administer folders, and edit application settings. Users with the administer permission also have the edit, create or submit, and view permissions. To delete items, the user needs to have the delete permission.
<b>Delete</b>	Allows users to delete items, but not to administer folders or edit application settings. Users with the delete permission also have view permissions. Without the delete permission, users with the administer permission are not allowed to delete items.

## How does CollabNet TeamForge help protect data access?

Access to data must be strictly controlled to meet the security requirements of the enterprise. Strict data access control is achieved through a combination of firewalls, authentication, and authorization.


### Firewalls and network configuration

A firewall provides the first level of protection by restricting access to the private network from the Internet. Sophisticated firewall configuration can provide strong security for all enterprise resources.

All CollabNet TeamForge application server nor the backend servers should ever be exposed to the Internet.

The CollabNet TeamForge application to function effectively, the following conditions must be met.

- Across the firewall, clients (users) must have access to:
  - The web server through a secure protocol such as HTTPS (port 443). The web server typically handles both the browser requests as well as the SOAP requests and forwards them to the CollabNet TeamForge application server.
  - Send mail to CollabNet TeamForge mail server via SMTP (port 25).
  - The SCM server through a secure protocol such as SSH (port 22).
- The web server must have access to the application server (typically port 8080).

 **Note:** This port is not exposed outside the firewall.

- The web server must have access to the SCM server for repository browsing functionality.
- The application server must have access to the backend (SCM, database, mail, etc.) servers.
- The SCM server must be able to access CollabNet TeamForge for commit notifications.
- The mail server must be able to deliver messages across the firewall.

### Authentication and authorization

To secure sensitive data, CollabNet TeamForge provides access control tools to restrict unauthenticated and non-member access.


User authentication is supported through verification of username and password during login. Project administrators can completely restrict access to authenticated members by marking projects as gated communities or private. A gated community is only accessible to unrestricted users, while a private project is only accessible to its members.

## How does CollabNet TeamForge help protect my data?

Sensitive data must be protected from illegal access at various points in the system. Key areas where security is typically compromised include data transmission and data storage.

### Data transmission

Network traffic is not encrypted by default. The HTTP protocol (non-SSL) does not protect data during transmission. HTTPS provides Strong Encryption using the Secure Socket Layer and Transport Layer Security protocols (SSL/TLS).

 **Note:** The web server employed by a CollabNet TeamForge installation must be reconfigured to employ the HTTPS protocol.

### Data storage

Sensitive data, such as credit card numbers, financial information, etc., must be stored securely. Usually this is done by encryption. In the context of an application like CollabNet TeamForge only stores password digests with an MD5 based cryptographic hash to guarantee adequate data protection.

MD5 is a one-way hash function that is used to verify data integrity through the creation of a 128-bit digest from data input. A one-way hash function is designed in such a way that it is hard to reverse the process, that is, to find a string that hashes to a given value. MD5 is currently a standard, Internet Engineering Task Force (IETF) Request for Comments (RFC) 1321. According to the standard, it is "computationally infeasible" that any two messages that have been input to the MD5 algorithm could have as the output the same message digest, or that a false message could be created through apprehension of the message digest.

## How do site administrator roles work?

Site administrators are default site managers who can create additional site administrators and delegate few site administrative tasks to them.

They can also allow some CollabNet TeamForge users to use one or more CollabNet TeamForge tools across several projects, by creating site-wide roles with specific project permissions, minus site administrative permissions. They can also provide ready-to-use roles as global project roles, creating uniformity across the site.

In CollabNet TeamForge, site or project administrators assign roles to the site users or project members. Besides this, a project member can submit a role request to the project administrator. The project administrator can approve or reject such requests.

A role defines these things:

- The applications that project members with that role can and cannot access.
- The resources on which project members with that role can use the applications.
- The actions that project members can take in each application and on each resource.

When a user's roles do not include access to an application or resource, that application or resource is not visible to that user. For example, imagine that you are assigning roles to Jason, a software developer. Jason needs to check source code in and out in order to fix bugs, develop features and create software releases. However, Jason does not need access to project wiki. If you set up Jason's roles according to those requirements, Jason's experience is like this:

- On any page in the project site, Jason can see and click the **Trackers**, **Source Code**, and **File Releases** buttons along the top of his screen.
- Jason does not see the **Wiki** button.
- If someone sends Jason a link to a page in the Wiki application and Jason clicks the link, he gets an error message. (The message does not specify whether the page exists or not.)
- When he accesses the project directly from Eclipse or Visual Studio, Jason can expand the project node and browse the **Trackers**, **Source Code**, and **File Releases** nodes, but not the **Wiki** node.

## Applications

An application is a collection of related features designed to enable a user to perform tasks and collaborate with other users. For example, the Documents application helps users create documents, share in document reviews, and publish documents, among other things.

In the Web interface, each application is represented by a button in the navigation bar at the top of any project page. A given user can see the buttons corresponding to applications they have access to by virtue of the roles assigned to them.

Applications are also known as "tools."

## Resources

- The tracker application might contain a bugs tracker and a feature request tracker. These are the tracker resources.
- A project can contain multiple SCM repositories. These are the SCM resources.

## Site Administration Responsibilities

The additional site administrators can be granted administrative rights for any of the site administrator responsibilities related to the following:

- Projects (includes project templates)
- Project Groups
- Users
- Groups
- Roles
- Categories
- System Tools
- Integrated Applications

## Site Administration Permissions

### View Only

Allows users to view and download items, but not to create or edit items, administer or edit application settings.

### Create or Submit

Allows users to create or submit and edit items, but not to administer or edit application settings. Users with the create or submit permission also have the edit and view permissions.

### Edit

Allows users to edit items, but not to administer items or edit application settings. Users with the edit permission also have the create / submit and view permissions.



**Administer**

Allows users to create, edit and administer items plus edit application settings, if required. Users with the administer permission also have the edit, create or submit, and view permissions. To delete items, the user needs to have the delete permission.

**Delete**

Allows users to delete items, but not to administer items or edit application settings. Users with the delete permission also have view permissions.

Without the delete permission, users with the administer permission are not allowed to delete items.


**Which role is assigned to me?**

There are several ways in which you could have been assigned certain roles in CollabNet TeamForge. Your access to the projects depends largely on the permissions granted to you via your roles.

You could have been granted access to any CollabNet TeamForge project in any of the following ways:

- Through site-wide roles (assigned directly to the user)
- Through global project roles (assigned directly or inherited via another project, either individually or as a group member)
- Through project roles (assigned directly or inherited via another project, either individually or as a group member)
- Through permission inheritance (via project hierarchy)

There are two ways you can check which roles are assigned to you in any of your projects:

 **Note:** If you are a site administrator, you can view any user's access rights to the system.

**To view the roles assigned directly to you:**


Use the **My Workspace > My Page > Projects** page and click the **My Projects** tab.


 **Note:** You can click the role to view the folder-level permissions assigned to you in a project.

**To view all the roles assigned to you:**

Use the **My Workspace > My Page > My Settings** page and click the **Roles** tab.

You can select **Roles Created For a Project**, **Roles Inherited From a Parent Project** or **Site-wide roles** to view the corresponding roles assigned to you.

 **Note:** A global project role that has been assigned to you in a project where you are a direct member will be listed as a directly assigned role.

 **Note:** You can click the role to view the folder-level permissions assigned to you.

**Who can access project planning information?**

A project manager should consider carefully who is able to change the project plan.

You can specify the people who can change the product plan based on their roles in the project. Only users with the role you select can view, create, modify, and delete artifacts in the planning folder hierarchy, and populate those artifacts with content.

Every project must strike its own unique balance between openness and control. These broad guidelines may help you decide how to set up the appropriate roles for your project.

- Err on the side of visibility. Most of the time you will want all project members to at least see the contents of all planning folders. This helps developers avoid duplicating their efforts, and it can enable team members to volunteer their help to each other when appropriate.

However, occasionally you may want to hide all or part of the planning folder hierarchy from selected people. For example, when you bring on a contract developer for a short period, you may want to restrict that person's view to the specific artifacts they are working on directly.

- Limit the number of people who can modify or populate planning folders. In many cases, only the project manager really needs to do this.

However, it may make sense to allow some people with a direct stake in the project to modify some folders. For example, if you are working with a product owner, it might be useful for that person to be assigning artifacts to future planning folders while your team is working on the current one.

## Who can see a project page?

A user can see a given project page if the page is not hidden and the user's role includes permission to see the page.

- 👉 **Note:** When a page is hidden, it is only visible to users with the Project Admin role, and then only when the user has clicked **Configure: On**.

If a person can see your project, they can see the project home page and any of its subpages. To see any other page, they must be assigned permission explicitly.

Permissions determine two kinds of access:

- Anyone with access to the project can see the project homepage.
- Users whose role includes permission to see a given top-level project page can see all the subpages of that page.

Subpages inherit their permission setting from the top level page they belong to.

## Who can see a project page component?

To see a given project page component, you must have permission to see both the project page where the component is and the tool that the component represents.

Your ability to see and edit the contents of a project page component is controlled by your permissions with regard to the tool the component represents.

For example:

- If you have permission to create documents in a selected folder in the Document Manager tool, you have that same permission when viewing that folder in a Documents component on a project page.
- If you do not have permission to view a folder in the Document Manager tool, you cannot see a document component that contains that folder on a project page.

- 👉 **Note:** Users with the Project Admin role can see project page components even if they do not have access to the related tool.

## Who is an owner of the sub component?

The user who creates the sub component is the owner of the sub component.

## As a project admin, why don't I have permissions to the wiki?

Most likely, your site was installed before TeamForge contained the wiki component.

When the wiki was added to TeamForge, it was decided that there was no way for us to know the security requirements at customer sites, so permissions for the new wiki component were not assigned to project admins by default. As a project admin, you can alter any existing role to grant this permission, or create a new role for this permission and then assign to the appropriate project members as needed.

## Concepts and terms in TeamForge

---

Descriptions of concepts and terms in TeamForge.


Concept definition.

## What is an association?

TeamForge allows users to easily associate, or link together, any objects in the system to simplify knowledge sharing and provide traceability throughout the lifecycle.

For example, a discussion post regarding a customer problem could be linked to a document specifying the feature requirement. An issue might be created to track the defect, the source code commits that fix the issue, and the release that contains the fix. All of these records can be linked together with an association. When any item is associated with another, the link appears on the item's **ASSOCIATIONS** tab.

Associations enable development organizations to improve information sharing, capture institutional knowledge, and simplify regulatory compliance.

 **Note:** When an association is added to or removed from TeamForge objects such as tracker artifacts, tasks, documents, discussions, and file releases, a notification mail is sent to users monitoring these objects.

An option is provided at site level and user level to make sure whether the notification mail has to be sent or not. For more information on this, see [Configure your site's settings](#) on page 365.

## What is change management?

Change management is about monitoring and addressing work items, product defects, feature requirements, and support requests.

## What is a document in CollabNet TeamForge?

For TeamForge purposes, a document is any file that you track with a document artifact.

Any project involves some number of documents that have to be written, edited, reviewed and published. You can use a document artifact to facilitate and standardize the work you do with documents. TeamForge automatically maintains the version history of each document, including the person who posted each version, its status, and any version comments.

If you are the author, use the Review feature to track the process of getting others to evaluate your document.

To make sure everyone who needs to know about a document knows about it, associate it with a task, a code commit, or another kind of artifact.

## What is a feature tree?

A feature tree is a way to look at planned product features hierarchically, so you can quickly understand the relationships among product features you have specified.

A feature tree is what you see when you click the **Tree** button on a tracker overview screen. (You can return to the traditional single-level tracker view by clicking **list**.)

The feature tree shows you the nested descendants of an artifact at whatever level of detail you need, all the way down to the tasks supporting your lowest-level user stories. When you define a user story that depends on several tasks for its completion, and you declare each of those tasks a "child" of the user story, they appear as subordinate nodes under the "parent" user story artifact in the feature tree.

User stories are the most common components of a feature tree, but any kind of artifact can appear. For example, when the completion of a user story depends on fixing a bug, the artifact representing the bug appears in your feature tree as a "child" node under the artifact representing the user story.

For each artifact in the feature tree, look for this key information:

### Priority

A relative measure of how important the feature is, compared with the other features you want in the product.

### Artifact ID

The permanent alphanumeric identifier for the artifact. Use the **Jump to ID** box to quickly find artifacts by their artifact ID.

<b>Title</b>	A short phrase identifying the work that the artifact describes.
<b>Status</b>	The artifact's position in the work flow of its project. Standard status values are "Open," "Under development," "In QA," and the like.
<b>Estimated Effort</b>	A measure of the work required to deliver the work described in the artifact. The units used are up to you -- just make sure they are used consistently across the project, so that you get accurate totals when you need them.
<b>Remaining Effort</b>	A measure of the work that's still to be done on the artifact, in the same units used in <b>Estimated Effort</b> . Project members must update this value whenever they advance the work in the artifact.
<b>Planning Folder</b>	The planning folder that the artifact is assigned to, if any.

## What is a discussion forum?

TeamForge discussions provide workspaces where project members can work together online or by email.

Discussion forums and mailing lists are closely integrated. Forum administrators can choose to enable a mailing list for each project forum. A mailing list extends the discussion forum functionality to allow project members to post messages to the forum using email.

Discussion forums can be public or private, depending on the forum's objective and desired level of access into the forum. Private discussion forums restrict anyone without specific access permissions from viewing or posting into the forum.

In a moderated discussion, messages from anyone except "trusted" users are screened by a moderator before they are posted. Messages posted by trusted users do not require the moderator's approval.

It's a good idea to use a discussion forum instead of direct personal email whenever privacy or security concerns don't prevent it, even when the communication only involves two project members.

- Members not directly involved often can make unexpected contributions if they are aware of the discussion.
- TeamForge archives all news items, forum posts, and mailing list communications, so you can go back and find valuable information later.

## What is an integrated application?

An integrated application is a stand-alone application that can seamlessly integrate into any TeamForge project.


You can use integrated applications to incorporate these types of applications into your TeamForge project:

- Third party applications
- Internally developed applications
- Integrations developed using the TeamForge SOAP APIs
- External websites


When you add an integrated application to your project, an icon is added to your project navigation bar. Clicking the icon displays the integrated application in the main TeamForge project window.

TeamForge site-administrators can register site-wide integrated applications that project administrators can opt to use across projects.

Site administrators or users with site-wide roles with the administration permissions for integrated applications can enable/disable integrated applications.

 **Tip:** Disabling an integrated application restricts it from being added to projects. However, disabling an integrated application does not affect the projects where the integrated application might already be in use.

After your site administrator registers an integrated application on the site level, on adding it to your project, an icon is added to your project navigation bar. Clicking the icon displays the integrated application in the main TeamForge project window.

-  **Note:** You can register and integrate as many applications per project as you wish. However, because each integrated application adds an icon to the project navigation bar, creating a large number of integrated applications can cause horizontal scrolling.

## What is an iteration?

An iteration (also called a sprint) is the period of time it takes to produce something that works well enough to show to the user.

At the end of an iteration, users get to try out the functionality. They can accept it or send it back for more work.

- Acceptance means the user agrees that the acceptance criteria stated in the user story have been met, and the capabilities as implemented support the needs that the user identified.
- Rejection means the product, as implemented, does not support the user's requirements. Sometimes this means the development team has not understood the user story the way the product owner meant it. Sometimes it means the user's initial ideas turned out not to be so helpful in practice. In either case, this leads to more analysis and development work.

Keeping iterations as short as possible helps reduce the time spent on corrections when features don't meet the users' expectations. It is less costly to redo two weeks of development than six or eight weeks.

Shorter, more modest iterations also tend to generate more reliable burndown data for estimating further work.


## What is a linked application?

A linked application is an external application or site that users can get to from inside a TeamForge project.

You can use linked applications to incorporate these types of applications into your TeamForge project:

- Third party applications
- Internally developed applications
- Integrations developed using the TeamForge SOAP APIs
- Company intranet sites
- External websites

After you create a linked application, a button is added to your project navigation bar. Clicking the button displays the linked application in the main TeamForge project window (or in the case of a site-wide linked application, in a separate window).

-  **Note:** You can create as many linked applications per project as you wish. However, because each linked application adds a button to the project navigation bar, creating a large number of linked applications can cause horizontal scrolling.

TeamForge administrators can also create site-wide linked applications that appear in all TeamForge projects.

## What is the look project?

The `look` project contains special files that can override your site's default appearance and content, such as the default icons, fonts, colors, and labels.

Unlike most projects, the `look` project has no members. It is only visible to users with site administration permission. Its only purpose is to control your site's look and feel, including such things as fonts, background colors, icons, and the wording of the onscreen labels that appear throughout your site.

Any project on your TeamForge site can have one or more Subversion repositories associated with it. The `look` project has just one Subversion repository. That repository is named `branding`.

When a user requests a page from your site, TeamForge checks the `branding` repository to see if any files there specify custom fonts, colors or text strings. If such specifications are found, TeamForge displays the page according to those specifications. If not, the page displays according to the default design.

Having your custom look-and-feel specifications in a Subversion repository enables you to roll back changes, track contributions, and use all the other features of a source code versioning system.

## What is a parent project?

A parent project is the base from which a subproject's members, user groups and roles, with their corresponding permissions, are derived. A subproject can inherit project members, user groups and roles from its parent project.

You can create a parent project to track and manage several smaller assignments as subprojects.

When you define users and roles with specific permissions in a project, those users and roles are passed down to any subprojects that belong to that project. This helps you avoid the repeated effort of defining users, user groups and roles across projects.

## What is a patch?

A patch is a package of code that fixes or adds to the functionality of a CollabNet product. Patches are also known as "component upgrades."


### Things to know about patches

- Patches are cumulative. You don't need to apply multiple patches sequentially to get to the desired patch level. You can move up (or down) one or more patch levels with a single operation.
- The Level option (-l) allows you to downgrade or upgrade to any patch level (within the maximum available in the cumulative patch).
- The Rollback option (-r) allows you to revert the site to the previous patch level it was at, before the current patch was applied.
- The Uninstall option (-u) allows you to downgrade the patch level on the site by one.
- When a patch installation fails you can use the Force option (-F) to proceed, without manually uninstalling previous patches.
- The system displays a summary of what happens during the patch installation.
- Before proceeding with the patch installation, you can use the "dry run" mode (-t option) to see the summary of actions that will be performed during the installation.

### Best practices

Before applying a patch, note the following principles.

- The upgrade scripts are usable only with an existing installation.
- No data migration will occur if any changes have been made to the database schema.
- You must use the `sudo` command or have an account that is equivalent to root in order to complete a patch installation successfully.

 **Important:** Before installing a patch, verify that it has been fully tested and qualified.

## What is a planning folder?

A planning folder is a virtual container that helps you organize and plan the work that goes into delivering a product.

You can create a hierarchy of planning folders to organize artifacts by product, release, iteration, etc. You can store artifacts from multiple trackers in a planning folder. This allows you to plan various stages of your project (i.e. releases, iterations, etc.).

Selecting an individual planning folder provides a view of all artifacts from all trackers within the selected planning folder.

The **Planned For** field identifies which product, release, or iteration the artifact is planned for, based on the planning folder that it is assigned to.


For example, in an agile development environment, a project manager breaks down the prospective product into its component parts and looks at what it would take to deliver each one. When all the parts planned for a given iteration are finished, the product is considered complete for that iteration.

Some parts of a product can be developed more or less in isolation, but most depend on other parts. Tracking these relationships is one of the trickiest aspects of product development.

For example, you can only provide a graphical user interface for a shopping cart application if you also come up with a database for the customer's payment data to be stored and accessed. That in turn requires a data storage and backup solution of some kind. And so on.

Use a planning folder to track the dependencies among the parts of your project as each moves toward completion. As you work through the question of what depends on what, you'll move artifacts representing user stories into the appropriate iteration in the appropriate release. As you proceed, you'll find a pattern like this emerging:

- Product 1
  - Release 1
    - Iteration 1
    - Iteration 2
  - Release 2
- Product 2

 **Note:** "Move" is meant figuratively. When you "move" an artifact into a planning folder, it is still a member of the tracker where it lives, and you can still do all the things with it that you can do with an ordinary tracker artifact.

Your planning folder lets you see at a glance the pieces of the work that support other pieces, and the pieces that depend on other pieces. Think of this as a "planning tree." If you are responsible for a development project, you can use this tree view to understand and predict the time and effort required to deliver a given set of features.

## What are Planning, Task and Kanban boards?

When you have set up your planning folders and teams, you have four views available to work with them: **List**, **Plan**, **Track** and **Kanban**.

TeamForge user roles and permissions that are in place for planning folders apply to all the four views.



### Planning Board

The Planning Board is an important tool for your TeamForge project's Agile planning activities. It enables you to plan and monitor the features that are required in each sprint (or iteration), and assign them from the product backlog to specific sprints. The planning board view complements the list view. While the latter offers you capabilities to accomplish various actions such as create, edit, and delete artifacts, planning folders and teams, the former offers product owners (or similar users) the ability to view, rank and move artifacts across the three planning folders (swimlanes) in a physical board-like user interface. In the Planning Board, planning folders are represented as swimlanes. In each swimlane, the tracker artifacts for the selected planning folders are represented as cards. You can also have a team's view of artifacts (backlogs and tasks), which is a swimlane representation of artifact cards for the selected team in the selected planning folder.

### Task Board

The Task Board is an important tool in the Agile process. It helps the team to focus on the work at hand in the current sprint and feed progress data back into the system. Unlike the list view and planning board view, which can be used for agile project planning, the task board view is for tracking tasks in a sprint. TeamForge project administrators can configure the Task Board. Once that's done, projects members can use the Task Board to break down the stories into tasks and then progress the tasks, and the story, towards completion.

The Task Board can have at least two and up to seven swimlanes depending on how your project administrator configures it. Every swimlane in the Task Board represents a task status. Once configured, team members can use the Task Board to view tasks in a selected planning folder or filter tasks for a specific team within the selected planning folder, add new tasks for a backlog item (epic, story, etc) and move tasks from one swimlane to the other as tasks progress from one status to the other.

## Kanban Board

The Kanban Board is an agile project management tool, which gives you a snapshot of the statuses of work items within a planning folder, how your project teams are placed in terms of work distribution and directs you to re-distribute the tasks to ensure optimal resource utilization.

Kanban Board uses the kanban method for project implementation. Kanban is built on two basic concepts: value stream mapping and WIP (work-in-progress) limits.

- **Value stream map:** A value stream, as defined in the *APICS (American Production and Inventory Control Society) dictionary*, comprises "the processes of creating, producing, and delivering a good or service to the market". In software development, a value stream map is an end-to-end mapping of the flow of activities (tasks/work items) from one state to another, starting from conceptualization to delivering the product to the customer.
- **WIP limits:** These are constraints (minimum and maximum) applied on each point or state (Planning, In Progress, etc.) in the value stream to ensure optimal WIP. This defines the minimum and maximum artifact count that ought to be present in each state so that if these constraints are violated, Kanban Board flags the issue for you to fix the bottleneck.

For example, let us assume that you want to view the status of the work items for a planning folder called 'Iteration 1' and assess the distribution of work between teams, Team A and Team B within this planning folder. You want to see a maximum of 8 artifacts in the 'In Progress' state and you configure your Kanban Board accordingly. When you go to your Kanban Board and view the artifacts for the selected planning folder, you see a constraint violation because there are 14 artifacts in this status. In addition, when you see the work split-up between the teams, Team A has 14 artifacts whereas Team B has none. This is clearly an indication that not only is the overall count of artifacts more than the maximum constraint, but even between the teams, Team A is overloaded and Team B is underutilized; this calls for re-distribution of work items to avoid any delay in the development process.

You can configure the Kanban Board based on your project activities and you can configure as many boards as you require.

## List, Plan, Track and Kanban views

You can click the **LIST**, **PLAN**, **TRACK** and **KANBAN** buttons to toggle between these views.

When you make a planning folder or team selection in the respective list view, (either from the planning folder/team tree or through the **Jump to ID** box), the selection is retained when you move to the task board or kanban board view. Similarly, when you select a planning folder or team in the task board or kanban board view, the selection is retained as well when you move to the list view. These selections are retained across browser sessions as well.

- 👉 **Note:** An exception to this case is when you select the root planning folder or team ('Project Teams') in the list view, this selection is not retained when you move to the task or kanban board view; instead, the selection of a specific planning folder or team you had made prior to this one is retained.

In the Planning Board, multiple planning folder and team selections are retained when you navigate to other pages in TeamForge and return to your planning board.

## What is a product owner?

The product owner works with the users, or customers, to define how a product can address their needs and wishes.

As the intermediary between the user and the development team, the product owner defines the product scope and requirements.

The product owner is sometimes called a product manager or a marketing manager. If your organization has a project office, that group may include people who function as product owners.



The product owner's primary focus is on the content of the product. The details of implementation planning and scheduling will affect and be affected by the product owner's work, but they are not the product owner's direct concern.

In TeamForge, the product owner's main focus is the feature tree. This is the part of the interface where discussions of product features are organized and tracked.

## What is a project administrator?

The project leader takes user stories created by the product owner and works with the project team to manage the process of analyzing, scheduling and executing them.

The project leader coordinates the development team.

The project leader is sometimes called a scrum master, a team lead, or some other variation.

The project leader's primary focus is on the development of the product. This work provides feedback that's useful for defining and revising the scope of the product and for estimating the effort involved, but those are not what the project leader is directly concerned with.

In TeamForge, the project leader's main focus is the planning tree. This is the part of the interface where you define and schedule the work items and personnel needed to deliver the required features.

## What is a project?

A project is a workspace where people can use the CollabNet TeamForge applications to collaborate and to create, store, and share data.

All the work you do with CollabNet TeamForge is organized into projects. Any registered CollabNet TeamForge user can create a project, subject to approval by a CollabNet site administrator. After a new project is approved, the project creator can configure project applications, add project members, and create and assign roles to govern each user's individual access permissions and the access permissions of groups of users. A registered CollabNet can also request membership in any CollabNet project. Requests to join projects are submitted to the project's administrators for approval.

How CollabNet TeamForge projects are organized is up to you and your organization. You might choose to create one large, centralized CollabNet TeamForge project in which to manage all of your organization's development work. Or you might choose instead to create a number of smaller projects for each team or sub-project.

Any registered user on the site can create a project, subject to approval by a CollabNet TeamForge site administrator.

You can use a project template to pre-populate new projects with the structure and configuration of an existing project.

When you create a project, it is submitted to the CollabNet TeamForge administrator for approval. You will receive an email notification when the site administrator approves or rejects your project. When your project is approved, you are assigned the Founder Project Admin role and made a project administrator. You can access the project from **My Projects** or **View All Projects** menu option under **My Workspace**.

## What is an agile project team?

As a member of a project team, you implement product features according to user stories created by the product owner and planning input from the project leader.

Project team members must use their experience and judgment to estimate the effort involved in work that's been specified by the product owner and the project leader. These estimates will often lead to changes in the product scope and the project plan. In that sense, project team members are an important part of the agile planning process.

The project team members may be known as developers, programmers, engineers, individual contributors, or something else that's appropriate to the type of work being done.

In TeamForge, the project team member's main focus is the tracker. Project members can manage and track their work by observing and maintaining tracker artifacts such as user stories, defects, or tasks.

## What is in a project template?

When you create a template from an existing project, each project tool contributes its own structure to the template, and its content if you want it.


For each tool, you can include or omit the actual content that was created with that tool.

For example, suppose you have a project for which you have created some tracker artifacts, and these artifacts have proved useful to members of the project.

- You can include those artifacts in your template, so that people who create a new project from that template will have access to the same artifacts that you developed for your original project.
- You can leave those artifacts out of the template and let future users create new artifacts to fit their own needs.

Imagine that you have documented company-wide process standards on wiki pages in an existing project.

- You can include those wiki pages in a project template, so that the manager of a project created from that template won't have to go find those pages and copy them into the new project.
- Or you can leave the wiki tool empty and let the new project's users create new wiki content for themselves.

 **Note:** You can choose not to include content from the original project, but you can't choose not to include a project tool. Every project template you create must include all the project tools, even if the project from which you created the template does not use all the tools. Every project you create from a template will include all the tools.

Tool	Always in Template	Can Be in Template
Tracker	You can have any number of separate trackers in a project. When you create a project template from that project, all the trackers that exist in the original project are copied into the project template. Any default column configurations, select field dependencies or text field validation rules are also included.	You can include the artifacts that were created in the original project, and any parent-child dependencies among them. If users of the original project have shared saved searches, these can be included too.
Planning folders		You can include all the planning folders in the original project, or none of them.
Documents	When you create a project template from an existing project, all the document folders that were created in the original project are copied into the project template.	If users have created documents and attachments in the original project, you can choose to include these documents in the project template.
Tasks	A project can have any number of task folders. All of those folders will be included in any project template that you create from that project.	You can include the tasks that were created in the original project, and any dependencies among them.
Discussions	All discussion forums that have been created in the original project will be included in any project template that you create from that project.	You can choose to include discussion topics, posts and attachments in your template, or leave them out.
File Releases	Any packages or releases that have been created in the original project will be included in the project template. If you have mapped a	


Tool	Always in Template	Can Be in Template
	planning folder to a file release, that mapping is also included.	
Wiki		When you create a template from an existing project, all wiki pages that users have created in the original project can be included in the project template.
Integrated applications	All external applications you have integrated into the existing project become part of the project template.	Which elements are included depends on the application that is integrated. See the <b>Content includes</b> section of the <b>Projects &gt; Templates</b> tab.

## What is a project template?

Project templates enable you to capture and re-use the structure and content of existing projects, including project pages, custom tracker fields, and work flow definitions, to speed new project creation and standardize lifecycle processes.

Project templates allow you to create and configure new projects quickly, enforce organizational standards, and facilitate process improvement.

To create a new project, you can use any project template created by any project administrator.

 **Note:** You can also create a new project without using a template.

The template name and description appear on the **Templates** tab of the **Projects** list, accessible from the navigation bar on your **My Workspace** page. If your site administrator has made it possible to preview the contents of project templates, click the name of a template to see what's in it.

A template can include the structure of the original project without any of the content, or it can include both the structure and the content of the original project.

- "Structure" means the folders and sub-folders in the original project.
- "Project content" means any work items you or any other project member have created as part of the project. For example, when you create a tracker artifact to manage a piece of work, that tracker artifact is part of your project's content. Any documents you upload to the project and any wiki pages you create are part of the project's content.

## What is a release?

A release is a group of one or more files that are published as a unit.

Each release can have a maturity level attribute to describe its state of completeness. Maturity levels are predefined and include development build, alpha, beta, and general availability.


The TeamForge file release system enables users to publish files and groups of files to selected audiences. Using role-based access control, administrators can control which project members can access each package or release.


## What is a story point?

A story point is a measure of effort that expresses the relative difficulty of implementing a user story. You can use story points, also referred to simply as "points", to help estimate how much work can be done in a sprint.

Story points are useful for relative measurement. A story point has no specific value in hours, or lines of code, or anything else. When you are estimating work, use story points only to compare one piece of work with another.

When you record a value in the **POINTS** field of an artifact, that value is added to the total points (story points) in the planning folder that the artifact belongs to. You can then use that figure to support forecasting.

In the parent artifact, you can view the total of story points assigned to each of its children. The calculator icon indicates that the artifact's points is a sum of its child artifacts' points within the project. If the parent artifact has children in other projects across the TeamForge, you can include those points in the total as well. The icon, , indicates that the artifact's points include its foreign child artifacts' points.

 **Note:** The Tracker Admin needs to set the tracker to include foreign children in points calculations.

## What is a task dependency?

Task dependencies show what must be done before a given task and what will be done after the task is complete. Project managers can create dependencies between any two or more tasks in a project.

Task dependencies allow you to establish relationships between tasks. This can help to frame a task in the context of a larger project plan.


There are two types of task dependencies:

- *Predecessors* are tasks on which another task depends.


For example, before you can complete the task "Finalize project plan," you might need another project member to complete the predecessor task "Hire new staff."


- *Successors* are tasks that depend on your task.

In this example, "Finalize project plan" is a successor task to "Hire new staff."

 **Note:** Task dependencies and tracker artifact dependencies are different things.

- For tracker artifacts, an artifact with dependencies (a "parent" artifact) can't be considered closed unless all of its dependent artifacts ("children") are closed.
- For tasks, a dependency means one task can't start until another task is completed.

 **Important:** A task cannot be dependent on itself. You cannot make a task both a predecessor and a successor to another single task. If you attempt to do so, you will receive a circular dependency error.

 **Note:** When a dependency is added to or removed from a tracker artifact, a notification mail is sent to users monitoring the artifact.

An option is provided at site level and user level to make sure whether the notification mail has to be sent or not. For more information on this, see [Configure your site's settings](#) on page 365.

## What is a tracker?

A tracker is a collection of records that follow the development of a unit of work from conception through to completion. You can create a tracker to manage almost any kind of work that your project calls for.

In each TeamForge project, you can create any number of trackers. Each tracker tracks a different type of data.

For each tracker, you can define values for status, category, and other default fields. You can create your own fields to capture additional data that is specific to your project or organization. You can also create tracker workflows to specify the criteria necessary for users to change tracker status values.

Individual tracker entries are referred to as tracker artifacts. The role-based access control system enables you to control which project members are allowed to view, create, and edit tracker artifacts. Within a project or across the projects, a tracker can be cloned along with the workflow.

Summary information about the number and status of project tracker artifacts is provided on each project's **Tracker Summary** page. The Project Dashboard also provides an at-a-glance overview of the status of each project member's projects, including information on the number and status of the tasks and tracker artifacts in each project, and project overrun and underrun statistics.

## What is tracker work flow?

To help users handle their tracker items effectively, you can set up some work flow rules. Workflow rules require a user to do something to a tracker before they can reassign it or update its status.

Administrators can define these kinds of workflow rules:

**Status transitions that a user can make based on an artifact's current status.**

For example, if an artifact is **Open**, you can specify that it can be changed to **Pending** or **Cannot Reproduce**, but not to **Closed**.

**Status transitions that a user can make based on his or her role.**

For example, if an artifact is **Pending**, you can specify that only users with the role **QA Engineer** can change it to **Closed**.

**Field values that a user must provide when making a specific status change.**

For example, if an artifact is **Closed**, you can specify that a user must enter a comment in the **Comments** field before changing the artifact's status to **Open**. You can also require an attachment.

You can create a workflow for each combination of tracker status values in a tracker. A tracker can be cloned within a project or across the projects along with the workflow. If a user role is not existing in the destination project, a new role is created with the same name. The permissions associated with the role are not copied from the source project.

## What is a user group?


To manage multiple users at once, create a group to represent them.

You can create a group to facilitate managing many users who share one or more characteristics.

For example, giving all the users in the accounting division access to all financial projects might be laborious if you assigned the permissions one at a time. Instead, create a group and assign it access to the financial projects category.

A user group can have any number of roles. When a role is assigned to a group, every member of that group has that role.

If a project is a subproject of another project, it may inherit user groups and their associated roles from the parent project.


 **Note:** A user who has a role in a project by virtue of group membership is not necessarily a member of that project. Becoming a member of a project is a separate process.

Group permissions are cumulative. This means that each member of the user group has all the access permissions allowed by all of the assigned roles, plus any permissions that may have been assigned by other methods, such as application permissions or individually assigned roles.

## Velocity and average velocity

Your team's velocity is the amount of work the team has shown it has completed in an iteration. You can use this information to help estimate how much work can be completed in future iterations.







In the planning folder summary page, you can see your team's velocity for the specific Iteration planning folder.

 **Note:** Velocity is displayed only on the Iteration planning folder summary page.

Velocity is calculated based on the sum of points of closed artifacts in an iteration.

Note that velocity calculation takes into account the following:

- Points of the parent artifacts, that is, only closed artifacts which are at the top level of the planning folder. For example, in the following illustration, only the points of the highlighted artifacts and not their child artifacts are taken into account for velocity calculation.

<input type="checkbox"/>	2 HIGH	▼  artf195171
<input type="checkbox"/>	2 HIGH	 artf247043
<input type="checkbox"/>	2 HIGH	▼  artf242252
<input type="checkbox"/>	2 HIGH	 artf246675
<input type="checkbox"/>	2 HIGH	▼  artf245601
<input type="checkbox"/>	2 HIGH	 artf246677

- Autosummed points of these closed parent artifacts do count as well.

Velocity only makes sense as a relative measure. There is no specific velocity that is good or bad or standard, because no two teams take on work of exactly the same scale or complexity. However, if your team's velocity is increasing from sprint to sprint, you can surmise that you are on the right track.

### Average velocity





The average velocity is calculated for completed and ongoing iterations and is displayed only on the Release planning folder summary page. Average Velocity calculation is as follows:

Average velocity = Sum of iteration velocity of all iterations within the specific Release/Total number of iterations

Note the following for velocity and average velocity calculation:

- Nested iteration planning folders (child iteration planning folders or subfolders) are not included for velocity or average velocity calculation for the parent planning folder. For example, in the following illustration:

Average velocity for Release 1 = Sum of velocity of Iteration 1 and Iteration 2 / 2

▼  Release 1
▼  Iteration 1
 Iteration 1.1
 Iteration 2

In the above calculation, Iteration 1.1 (being the subfolder) is not included. Similarly, velocity calculation for Iteration 1 does not take into account its subfolder Iteration 1.1.

## What is an activity table component?

An activity table is a special kind of text component that gives project members quick access to a focused set of artifacts, action items and work products.

An activity table includes organized links to up-to-the-minute information about work that is going on in the project. You can add, remove or change these links with the HTML editor that comes with every text component.

You can add static links to particular artifacts, or you can link to a search query. When you link to a query, you enable project users to get all the results of the query with a single click.

The activity table template is just a suggestion for how you might want to organize your information. You can change any of its properties to support the way your project members work.

## What is a documents component?

A documents component allows team members to work with documents directly from a project page.


A documents component is like a window into your Documents tool. Instead of making users go to the documents tool and search for a document on their own, you can create a documents component on your project page that is linked to a folder in your documents tool. Each page can have its own list of relevant documents.

Users still go the documents tool, but now you guide them to what they need there.

## What is a global project role?

A global project role is a ready-to-use role available in all projects. Only site administrators or restricted site administrators can create and manage a ready-to-use role.

As a project administrator, you can use global project roles provided by the site administrators instead of creating and managing roles tailored to your projects.

 **Note:** You can use the ready-to-use roles to set up your team faster and with little fuss. However, you may not be able to edit the ready-to-use roles.

Before you create a role in your project, it is a good idea to check all the available ready-to-use roles. You are likely to get ones that grant the desired set of permissions.

Global project roles serve a different purpose from that of a site-wide role. The site-wide roles enable site administrators to create restricted site administrators for providing assistance in site management. Besides that, site-wide roles can be used to grant tool/application access across the site to a user.

## What is a project page?

A CollabNet project page is a place where users can see and add information about the project, such as messages from the project manager, open issues or documents you want people to read.


You can build your own project pages to design, manage and track your project's lifecycle. When you create a project page, users automatically see it in their navigation panes.

To post information or provide functionality on your project page, you add a page component of the appropriate type for the information or functionality you are working with.

For example, to let people know about how the project is coming along, add a Project Statistics component to the page. To let project members upload documents for other users to read, you add a documents component.

If the page is a part of a parent page, it appears as a node in the tree in the left navigation area.

You can add a page directly to the top of the project, or as a subpage.

 **Tip:** When you have your project pages defined and your process honed, you can export your project pages in the form of a project template, so other projects in your organization can reuse the resources you created. See [Create a project template](#) on page 303.

## What is a project page component?

A project page is a collection of simple portlet-like components that enable you to add custom HTML content, reports, project tracker queries, and much more to the project home page so your team quickly access what is important.

## What is an Available upon Request role?

The Available upon Request role is a role which a project member in TeamForge can submit a request for. You must be a site administrator or project administrator to create an Available upon Request role.

While creating global or project roles, you can select the **Available upon Request** option to allow the project member to see the role in **Request Additional Roles** section on the **Project Home** page.

When a project member requests a role, the request is submitted to the project administrator for approval. The project member receives an email notification when the request is either approved or denied.



**Note:** Based on the project administrator's discretion, a few role requests may be granted immediately on request, while the other role requests may need approval.

## What is a news component?

A news component is a way to get developing information to project members via a project page.

You can use a news component to maintain a journal or blog about your project, to announce milestones met, or to share information about a rapidly changing situation.

There is a single set of news posts for your project. If you put a news component on more than one project page, the same new posts will show on all of them.

## What is a project page component for?

Project page components add functionality to a project page. You can use components to communicate with users or project members, or to invite them to contribute information or resources to the project.

- Promote a standard lifecycle for all of your projects by using project pages to lay out information the way that works best for your business.
- Increase developers' productivity by putting the information most relevant to your project right on the homepage.
- Improve project coordination by combining related documents, graphs, artifacts, and content into single, digestible web pages.
- Maintain information security. The project manager controls who has access to each of the pages and components.
- Keep project members informed with live data and content that can be quickly updated as often as you like.

Project pages put the project manager in the driver's seat. You define how you want your project to look and what information you wish to share with your members and customers. Using a simple portlet-like layout, project pages let you add custom HTML content, reports, project tracker queries, and much more to the project home page so your team can quickly access what is important. All your pages are organized and displayed in a simple navigation menu on the left side of the home page.

In addition to customizing the project home page itself, you can also create and integrate additional web pages on the fly. Whether you need a page to display daily status reports or want to publish an HTML blog, a few clicks is all you need. No need to use any fancy tools or even know HTML.

For example, on a project page titled **QA**, you might have:

- A trend graph showing the open bugs in your project over a period of time.
- A Tracker query list showing you what hot bugs are currently in the system.
- A block of custom HTML content that the QA lead updates daily giving users his or her take on how the quality of the product is doing that day.

## What is a Tracker Search Results component?

A Tracker Search Results component provides quick access to Tracker information from a project page.

Some tracker searches are used so frequently that you want people to be able to see their results without exploring. For example:

- People working on your project may need to know at the beginning of every work day which artifacts have been updated overnight by a remote team.
- A cross-project leadership team may need a daily count of started and completed work items to support forecasting and planning.

Such heavily-used information might be a good candidate for a tracker search component on a prominent project page, such as the project home page. When you set up this component based on a popular tracker search that you devised and saved, project members and users can save time by viewing the information at a glance, without having to go find the information themselves.



**Note:** If your site administrator has set up Project Tracker for your project, you can also use a project page query component to search your Project Tracker data.



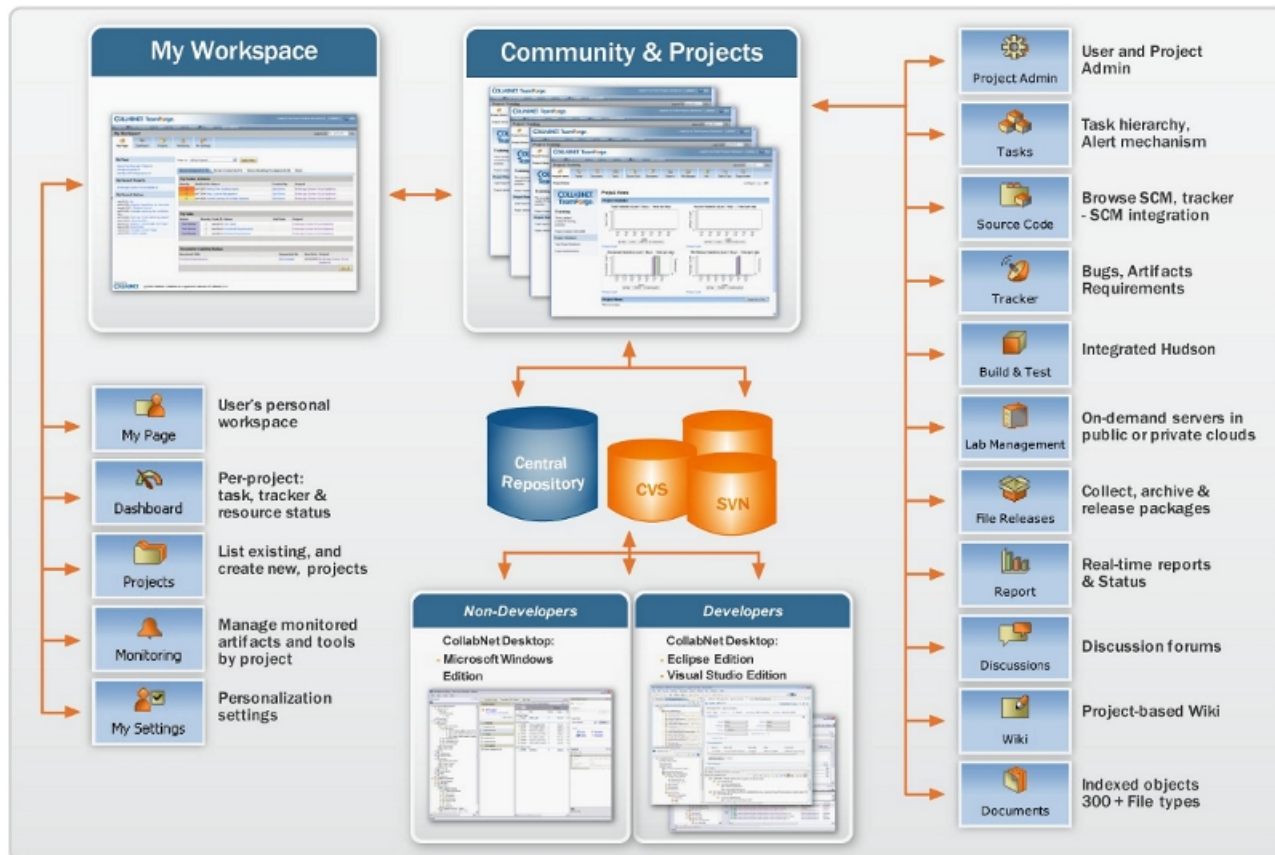
## What is a project statistics component?

A *project statistics component* gives users a graphical view of recent statistics for tasks, trackers, documents, and file releases from the project page.

For more information about the project statistics component see [Create a project page component](#) on page 136.

## What is TeamForge?

TeamForge integrates software configuration management (SCM), continuous build and test, issue tracking, project management, lab management, and collaboration tools into a single Web-based application lifecycle management platform that empowers distributed teams to deliver great software.



1. [My Workspace](#) on page 13
2. [Join a TeamForge site](#) on page 12
3. [Manage source code](#) on page 84
4. [My Workspace](#) on page 13
5. [The Project Dashboard page](#) on page 387
6. [What is a project?](#) on page 993
7. [Watch for changes](#) on page 129
8. [Join a TeamForge site](#)
9. [CollabNet Desktop - Microsoft Windows Edition](#)
10. [CollabNet Desktop - Eclipse and Visual Studio Editions](#)
11. [Help for project managers](#)
12. [Work on tasks](#) on page 66
13. [Manage source code](#) on page 84
14. [Plan and track anything](#) on page 32
15. [Build and test your software](#) on page 125
16. [Provision hardware](#) on page 117

17. [Release products](#) on page 125
18. [Reporting in TeamForge](#) on page 202
19. [Participate in discussions and mailing lists](#) on page 74
20. [Contribute to the project wiki](#) on page 82
21. [Document your work](#) on page 152

## What is a text component?

A text component is a self-contained editor with which you can write anything you want on your project page.

You don't have to know anything about HTML.

Text components include everything you would expect in a fully functional HTML editor. You can add tables, upload images to embed into the page, create hyperlinks, and assign font colors and styles. Advanced users can work directly with the HTML source.

## What is a wiki component?

A wiki component lets you link to an existing wiki page from a project page.

## What is the complete wiki syntax that TeamForge supports?

Since TeamForge makes use of the JSPWiki rendering engine internally, you may reference the JSPWiki syntax documentation at the link below for a complete list.

<http://www.ecyrd.com/JSPWiki/wiki/TextFormattingRules>

Please note however that JSPWiki has progressed beyond the release available in TeamForge currently, so there may be some discrepancy between what JSPWiki supports and what TeamForge supports. TeamForge also does not currently support the JSPWiki plugin framework, so any formatting plugins referenced on the URL will not be available for TeamForge.

## What is Team?

Teams helps you create logical groups of cross-functional members comprising architects, developers, testers and so on.

A team's view of backlogs complements the planning folder view. While planning folders represent backlogs (release, iteration, etc), a team's list view represents the team's view of the backlog (across releases, iterations, etc).

In an agile environment, project activities are broken into smaller units of work items and are assigned to various project teams. The Teams feature allows you to create these physical teams in TeamForge so that you get the specific team's view of backlogs. This knowledge of the team's view of work items helps you facilitate effective communication and execute project activities in a more structured and organized fashion. For example, you can easily view and filter artifacts which are taking longer than estimated within a team, analyze the scenario, and identify the impediments. Once identified, you can communicate them to the relevant team(s), re-assign it to other appropriate team(s) or team member(s), and resolve them quickly.

## What is a team owner?

Any project member can be designated as a team owner.

The team owner does not have the permission to create or delete teams. They can view their team information and edit them; while editing their team details, they can add or remove members, or designate another member as the team owner.

## General FAQs

---

Common questions on TeamForge usage.

## What's the best way to search the CollabNet User Information Center?

Searching this site with the Eclipse search tool is a little different from searching the Web with a service such as Google.

Here are some guidelines for getting the most out of a search on this site.


### Setting the search scope

To quickly locate topics on a particular subject in the documentation, enter a query in the **Search** field. You can narrow the scope of your search by selecting only the sections you are interested in. To set the search scope:

1. Click **Search scope** next to the **Search** field.
2. Select the **Search only the following topics** option.
3. Click **New**.
4. Type a name in the **List name** text box.
5. Select one or more guides (select the check boxes) from the **Topics to search** list. For example, you can select the check boxes corresponding to the CollabNet TeamForge 7.0 user and administration guides.
6. Click **OK**.

Once the search scope is set, you can enter your search query and click **GO**. Your search is restricted only to the guides you selected.

### Search query guidelines

- Use AND to require the term on each side of the AND operator be present in the topic.
  -  **Note:** When you don't put anything between terms, the search tool assumes you mean AND. Topics that contain every term in the query are listed in the search. For example, if you enter database service, topics that contain both the term database and the term service display. Topics that contain only the term service or only the term database do not display.
- Use OR before optional terms. For example, if you enter database OR service, topics that contain either the term database or the term service display.
- Use NOT before a term that you want to exclude from the search results. For example, database NOT service displays topics that contain the term database, but do not contain the term service.
- Use ? to match any single character. For example, plu? displays topics that contain plug.
- Use \* to match any set of characters, including an empty string. For example, plu\* displays topics that contain plug or plugin.
- Use double quotation marks to enclose a term that is to be treated as a phrase. For example, "edit menu" displays topics with this entire term, not topics with only the term edit or the term menu.
- Case is ignored. For example, database service displays database service, Database Service, and DATABASE SERVICE.
- Punctuation acts as a term delimiter. For example, web.xml displays topics that contain web.xml, web, and xml. To display only topics that contain web.xml, enclose the term in double quotes.
- In a search query, if you enter create, topics that contain create, creates, creating, and creation are displayed. To only see the term create, enclose the term in double quotes.
- The following English words are ignored in search queries: a, and, are, as, at, be, but, by, in, into, is, it, no, not, of, on, or, s, such, t, that, the, their, then, there, these, they, to, was, will, with.

## What is a Publishing repository? How does it work?

Publishing repository, like the branding repository, is one of the default repositories that's created automatically when a TeamForge project is created and is intended to contain publicly-consumable files.

The Publishing repo has a `www` directory. Files in the `www` directory are checked out to a working directory and served by the Apache server with no user authentication checks. In other words, files stored in the `www` directory do not go through TeamForge's RBAC checks and are publicly accessible even if the user is not logged in (accessible via a direct link to the file). By design, the files stored in the `www` directory are meant to be public on both "public" and "private" projects no matter whatsoever. However, files stored in no other directories but `www` are publicly accessible.

## Additional security Enhancements added to Publishing Repository

Site administrators can now toggle access to Publishing Repositories and restrict access based on defined RBAC. See [DISABLE\\_REMOTE\\_PUBLISHING](#) on page 740 for more information.

## Is the CollabNet TeamForge SOAP API backward-compatible?

The CollabNet TeamForge Enterprise Edition 5.x and CollabNet TeamForge Enterprise 6.x APIs (including Service Pack updates) are fully compatible with CollabNet TeamForge.

Applications developed using these earlier APIs will continue to function after you upgrade to CollabNet TeamForge. A release of CollabNet TeamForge may be accompanied by updates to the API. These changes are always backward compatible with earlier versions of the API. However, the calls from the different API versions are not interchangeable.

- For updates and patch releases, existing methods are incremented, not overwritten. For example, if an update is made to the `createDocument` method, it is reflected in a new method `createDocument1`. You do not have to update your existing applications to reflect a new API with any Service Pack or Hot Fix release.
- For a major CollabNet TeamForge release, such as this 17.4 release, updates are merged into a new API version. At this point, if you want to use the new API calls, you must update your existing applications.

## What can I do with the SOAP API?

You can use the TeamForge SOAP API to do almost anything a user can do in the TeamForge web user interface.

TeamForge exposes a subset of the APIs defined by the application server as web services, through the SOAP protocol.

A SOAP proxy server and a SOAP API layer, both running on Apache Axis, expose a set of web services representing each TeamForge application. The SOAP server serves the following functions:

- Provides web services by accepting SOAP requests from the clients.
- Performs SOAP client authentication.
- Implements TeamForge role-based access control (RBAC) and caching service.
- Accesses the application server via RMI stubs.

You can get the API here: <http://www.collab.net/community/teamforge>

The SOAP API can provide more types of functionality than the Web UI alone. For example, it can be time-consuming for a user with Tracker Admin permissions to copy a tracker workflow from one tracker to another in the Web UI. However, the SOAP method that provides a workflow copy function makes this task easy.

### Access permissions

For access permissions (authorization, as distinct from authentication), the SOAP API follows the same rules imposed by the TeamForge role-based access control (RBAC) system. A user using a SOAP API-based client has no more access to data than they have in the Web UI.

The CollabNet TeamForge architecture allows you to quickly and easily develop integration points between CollabNet TeamForge and applications you develop.

### User-centric services

All services and APIs are user#centric, meaning that all integrated applications must establish an individual connection to the SOAP server for each user. This differs from programming directly with an application server where one connection can be established for any number of users.

Activities that can be performed using the SOAP interface are by definition user#based, such as retrieving a list of a user's projects, tasks, or assigned tracker artifacts. These activities therefore require an individual connection for each user.

Requiring individual connections also ensures that role-based access control is checked for each action performed by each user. To ensure that security is enforced, RBAC checks are performed on each SOAP API call and cannot be disabled at the client level.

### Consistent interaction

While each CollabNet TeamForge service has its own SOAP interface, interaction with each service is designed to be as consistent as possible. The calls for each service are similar, although the data format and specific call parameters may be different.

For example, the following calls are consistent across all services:

- list
- get
- set
- delete
- create


The call parameters, however, are different. For example:

- When working with the CollabNet service, you might call `getProjectList(string sessionId)`.
- When working with the TaskApp service, you might call `getTaskList(string sessionId, string taskFolderID)`.

## How should my API client store user passwords?

Any reputable method of storing passwords will work, as long as your site is protected by SSL.

All client tools rely on the TeamForge SOAP API for authentication, and therefore use whatever authentication method TeamForge is using.

 **Important:** Because SOAP is simply XML transmitted over HTTP, all values are sent in clear text. For that reason it is very important that your TeamForge site be SSL-enabled and protected by server-side SSL certificates. This will ensure that any usernames or passwords sent from a client tool will be encrypted.

Many standalone client tools are able to cache a copy of the user's credentials to make it easier for them to access the site. The CollabNet Eclipse Desktop stores passwords in the encrypted Java keystore, and the CollabNet Windows clients use the Windows keystore.

CollabNet's Subversion clients and other Subversion clients, such as Tortoise and Subclipse, are also able to store user credentials. While CollabNet has no control over how third-party tools store such credentials, it our experience that the mainstream tools all use an appropriate keystore for secure storage of user credentials. CollabNet recommends that customers independently verify the storage methods of those tools and set a policy appropriate with their own security guidelines.

Subversion users on Linux systems have the option to use the Gnome keyring to securely store user credentials. CollabNet recommends that customers set their own policy for how their users should use the Linux Subversion client.

## How does an application interact with TeamForge SOAP services?

TeamForge exposes a subset of the APIs defined by the application server as web services, through the SOAP protocol.

### SOAP APIs

A SOAP proxy server and a SOAP API layer, both running on Apache Axis, expose a set of web services representing each TeamForge application.

The SOAP server provides the following functions:

- Provides web services by accepting SOAP requests from the clients.
- Performs SOAP client authentication.

- Implements TeamForge role#based access control (RBAC) and caching services.
- Accesses the application server via RMI stubs.

While each TeamForge service has its own SOAP interface, interaction with all the services is designed to be as consistent as possible. The calls for each service are similar, although the data format and specific call parameters may be different.

For example, the following calls are consistent across all services:

- list
- get
- set
- delete
- create

The call parameters, however, are different. For example:

- When working with the CollabNet service, you might call `getProjectList(string sessionID)`.
- When working with the TaskApp service, you might call `getTaskList(string sessionID, string taskFolderID)`.

### User-centric services

All services and APIs are user#centric, meaning that all integrated applications must establish an individual connection to the SOAP server for each user. This differs from programming directly with an application server where one connection can be established for any number of users.

Activities that can be performed using the SOAP interface are by definition user#based, such as retrieving a list of a user's projects, tasks, or assigned tracker artifacts. These activities therefore require an individual connection for each user.

Requiring individual connections also ensures that role#based access control is checked for each action performed by each user. To ensure that security is enforced, RBAC checks are performed on each SOAP API call and cannot be disabled at the client level.

## How do I require approval for new user accounts?

You can configure the system so that new users can create their own accounts, but the accounts are not activated until a site admin approves them.

To enable this mode of operation, add the following line to `/opt/collabnet/teamforge/sourceforge_home/etc/sourceforge_configuration.properties`:

```
sf.approveNewUserAccounts=true
```

Once this line has been added to the file, restart TeamForge for it to take effect.

Please note that site admins can still create accounts for new users and they will not be held for approval. Also note that the user will receive an email from TeamForge telling them to confirm their password by clicking on the given link, and the link will not work. The password is properly set on account approval.

## How does TeamForge automatically sum up effort estimates?

As you refine your agile project plan, you break down user stories into smaller stories, and eventually into tasks. TeamForge can watch the changing effort estimates all the way down the hierarchy, and give you a running total for each parent artifact and all its children.

### Example

Imagine that you originally created a user story to describe the need for a "Shopping Cart" feature in your e-commerce application. You gave this story a rough effort estimate of 20 units, and entered that figure in the **Effort** field.

Upon further discussion, your project team recommends breaking the story down into three tasks:

- Database (5 units)
- Back End (15 units)
- UI (8 units)

You create artifacts for those three tasks, identifying them as children of the user story artifact, and enter their respective effort estimates in the **Effort** field in each task artifact.

Back in the "Shopping Cart" story artifact, you select **Autosum effort** to indicate that effort number for this artifact should be rolled up from its children.


Now, rather than showing the manually entered number (20 units), your user story artifact shows the figure derived from adding the effort estimates for all three child artifacts (28 units).

### Double-counting?

Effort numbers are never "double-counted" when auto-summing is on. For example, if the sum for planning folder A would include effort for child C and parent P, and P has autosumming turned on, C's effort number is only counted once.

However, when auto-summing is off, the effort numbers from parent P and child C can both contribute to totals. So if a planning folder A includes parent P (autosum=off, effort=3) and child C (effort=5), then the pair contributes a total of 8 to A's total.

An artifact's effort number is calculated from the effort numbers of its immediate children artifacts only. This is true whether or not those children's numbers are themselves automatically derived.

 **Note:** When a tracker is disabled, artifacts from that tracker do not contribute to the effort totals calculated for any planning folder they are in.

### Permissions

If you are a long-time user of TeamForge tracker features, be aware that autosumming can lead to some situations you haven't seen before. For example, picture these scenarios:

- Artifact A is assigned to Sanil. He therefore has permission to edit that artifact, but he does not have permission to edit the artifact's parent, artifact B. Artifact B is assigned to Sergey, who has edit permission for both artifacts. Sergey turns on autosumming for artifact B. When Sanil updates his effort estimate for artifact A, the content of the **Effort** field in artifact B is updated as a result, even though only Sergey is explicitly authorized to change values in artifact B.
- Connie creates artifacts X and Y, and declares both of them children of artifact Z. Then she assigns artifact Z to Thiru. Thiru now has edit permission for artifact Z. But he cannot edit artifact X or artifact Y, because he does not have access to the separate tracker where they live. When Thiru turns on autosumming for artifact Z, that artifact's **Effort** field includes the sum of artifacts X and Y, even though Thiru is not explicitly authorized to see any data from those artifacts. (In fact, he may not even know that artifacts X and Y exist.)

## How does TeamForge use Velocity templates?

Velocity is the templating language that CollabNet TeamForge uses to render areas of the site with dynamic information.

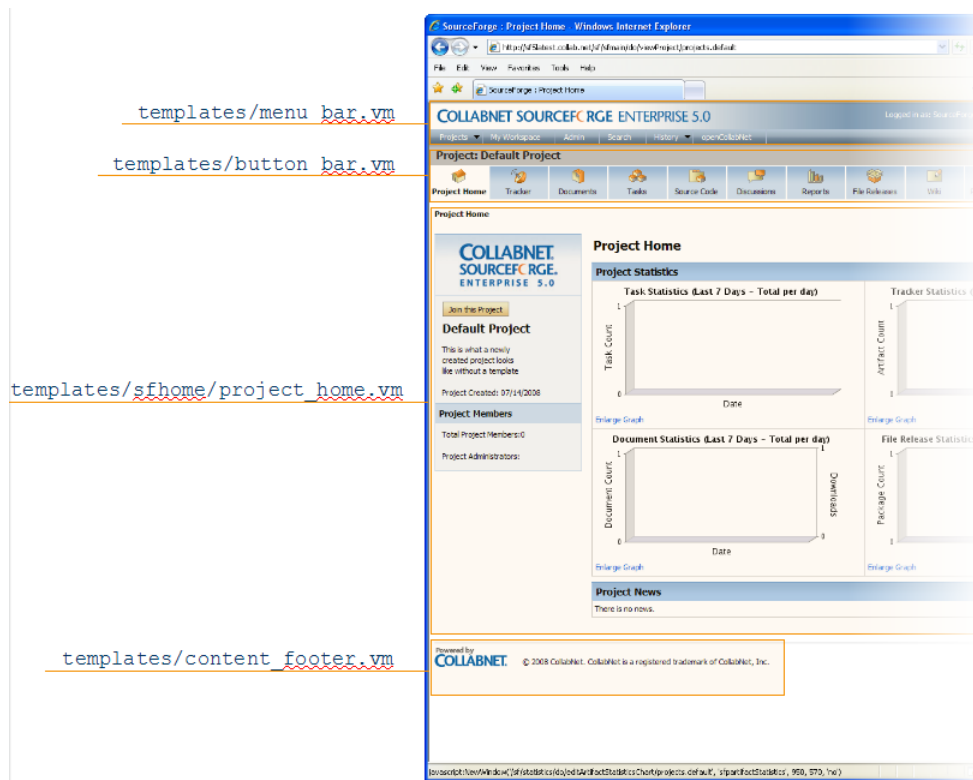
You can override the instructions contained in any of these Velocity templates by placing a file of the same name in the equivalent path in the `branding` repository in the `look` project on your site.

Velocity templates are located in the `templates` directory in the branding repository.

Velocity file	Description
<code>menu_bar.vm</code>	Controls the rendering of the top bar across all pages in the system. Displays a small login form, the site logo

<b>Velocity file</b>	<b>Description</b>
	and current user information as well as the search and projects drop down menus.
<code>blank_menu_bar.vm</code>	Contains only the top logo, without the menu that appears below it.
<code>body_header.vm</code>	Rendered immediately after the opening body tag. If a site requires everything to be contained in some other container, this template can be used.
<code>body_footer.vm</code>	Rendered immediately before the closing body tag. If a site requires everything to be contained in some other container, this template can be used.
<code>button_bar.vm</code>	Controls the rendering of the bar beneath the menu bar, which contains the 'Quick Jump" link as well as the buttons that appear on any project page (the one containing the applications). Site admin pages, user settings pages (e.g. my workspace, dashboard) and project pages use different sets of buttons that are passed into this template for rendering.
<code>content_header.vm</code>	Rendered after the button bar; wraps the actual contents of the page being viewed.
<code>content_footer.vm</code>	Rendered before the body footer; wraps the actual contents of the page being viewed. Contains the Copyright notice.
<code>sfmain/home.vm</code>	Velocity template that generates the site home page.
<code>sfmain/project_home.vm</code>	Velocity template that generates the default project home page.






## How do URLs in CollabNet TeamForge 17.4 compare with URLs in CollabNet Enterprise Edition?


The URLs in CollabNet TeamForge 17.4 and CollabNet Enterprise Edition are not exactly the same, but they can be made to work together.

Any bookmarks you have saved in CollabNet Enterprise Edition projects can now be accessed from within CollabNet TeamForge 17.4.

CollabNet Enterprise Edition URLs that specify item identifiers, such as artifact ID, document ID or discussion post ID, are specified are routed to the exact item page in CollabNet TeamForge 17.4. However, CollabNet Enterprise Edition URLs that do not have exact item identifiers may be routed to the tool's main page or project home page. As an example, link to CollabNet Enterprise Edition's specific post ID is redirected to the exact post ID page whereas, the link to a discussion forum may be re-directed to the **Discussions>Forum Summary** page.

 **Note:** If any of your bookmarked CollabNet Enterprise Edition URLs are not routed to the CollabNet TeamForge 17.4 URLs, contact your system administrator.

Only the URLs listed here will work the same way in TeamForge 17.4 as in CollabNet Enterprise Edition. You may need to change some other URLs manually and communicate the changes to your site members. For example, there may be a form submission to enter Project Tracker artifacts which may look like: `/servlets/tracking/action/CreateArtifact?type=defect`. The equivalent URL in TeamForge 17.4 would be: `/tracker/tracking/linkid/prp11042/action/CreateArtifact?type=defect`. You should get the corresponding URLs from the CollabNet TeamForge site and update your HTML or project pages accordingly.

 **Note:** If you are making the transition via a staging site, do this in production, after the upgrade is complete, because changes you make on the stage site are not preserved.

### Issue Tracker

CollabNet Enterprise Edition URLs, where issue tracker IDs or PCNs are specified, are re-directed to the exact artifact or tracker appropriately. Any other Issue Tracker tool-specific URLs are re-directed to the CollabNet TeamForge 17.4 Tracker application's home page.

- After migration, enter a CollabNet Enterprise Edition URL of an Issue Tracker issue in the browser address bar like: `https://project.domain/issues/show_bug.cgi?id=xxx`. This gets re-directed to a URL like: `https://host/sf/go/artfxxx`.

## Discussions

CollabNet Enterprise Edition URLs, where forum IDs, topic IDs or post IDs are specified, are re-directed to the exact forum, topic or post appropriately. Any other discussions tool-specific URLs are re-directed to the CollabNet TeamForge 17.4 Discussions application's home page.

- Bookmarked forum URLs of the form `https://project.domain/ds/viewForumSummary.do?dsForumId=xxxx` are correctly redirected to the corresponding forum in the CollabNet TeamForge's Discussions tool like `https://host/sf/discussion/do/listTopics/projects.project/discussion.forumId`.
- Bookmarked message url's of the form `https://project.domain/ds/viewMessage.do?dsForumId=xxxx?dsMessageId=xxxx` are correctly redirected to the corresponding topic or post page depending on whether the message represents a post or a topic, like: `https://host/sf/discussion/do/listPosts/projects.project/discussion.forumid.topicid...`
- Any other url's of the form `https://project.domain/ds/....` are redirected to the Forum summary page as: `https://host/sf/go/projects.project/discussion..`

## Documents

CollabNet Enterprise Edition URLs, where document IDs are specified, are re-directed to the appropriate exact document. Any other document tool-specific URLs are re-directed to the CollabNet TeamForge 17.4 Documents application's home page.

- CollabNet Enterprise Edition URL: `https://project.domain/servlets/ProjectDocumentReservationAdd?documentID=xxx` is re-directed to `https://host/sf/go/docxxx`.

## CVS

- For repository browsing purposes, CVS URLs are appropriately redirected to CollabNet TeamForge 17.4 ViewVC URLs. `https://project.domain/source/browse/project/dd.txt?r1=x&r2=y` is re-directed to `https://host/integration/viewvc/viewvc.cgi/dd.txt?root=project&system=exsyxxxx&r1=x&r2=y`
- For purposes of checking out or committing with a client, see the **Source Code** tab of the CollabNet TeamForge 17.4 project for the new repository path. CVS working copies must be recreated (checked-out afresh).

## SVN

Subversion URLs for specific revisions of a file, log messages or difference between two revisions are appropriately redirected to CollabNet TeamForge 17.4's ViewVC URLs.

- CollabNet Enterprise Edition URL: `https://project.domain/source/browse/project/trunk/test.txt?view=markup&pathrev=x` is re-directed to `https://host/integration/viewvc/viewvc.cgi/trunk/test.txt?root=project&system=exsyxxxx&view=markup&pathrev=x`
- CollabNet Enterprise Edition URL: `https://project.domain/source/browse/project/trunk/test.txt?view=log&pathrev=x` is re-directed to `https://host/integration/viewvc/viewvc.cgi/trunk/test.txt?root=project&system=exsyxxxx&view=log&pathrev=x`
- CollabNet Enterprise Edition URL: `https://project.domain/source/browse/project/trunk/test.txt?r1=x&r2=y&pathrev=x` is re-directed to `https://host/integration/viewvc/viewvc.cgi/trunk/test.txt?root=project&system=exsyxxxx&r1=x&r2=y&pathrev=x`

## PD/Task

Any URL of CollabNet Enterprise Edition's Dashboard/Task pages are redirected to the CollabNet TeamForge 17.4 Tasks application's home page.

- CollabNet Enterprise Edition URL: [https://project.domain/servlets/activity/\\*.do](https://project.domain/servlets/activity/*.do) is re-directed to <https://host/sf/taskmgr/do/viewSummary/projects.project/taskmgr>.

### Wiki

All CollabNet Enterprise Edition Wiki URLs are redirected to the CollabNet TeamForge 17.4 Wiki application's home page.

- CollabNet Enterprise Edition URL: <https://project.domain/wiki/...> is re-directed to <https://host/sf/wiki/do/viewPage/<projectname>/wiki/HomePage/sf/wiki/...>

### Project Metrics

CollabNet Enterprise Edition URLs of Report tool are redirected to the CollabNet TeamForge 17.4 Reporting application's pages.

- CollabNet Enterprise Edition URL: [https://project.domain/servlets/Report\\*...](https://project.domain/servlets/Report*...) is re-directed to <https://host/sf/reporting/do/listReports/<projectname>/reporting>.

### Announcements

All CollabNet Enterprise Edition News URLs are redirected to CollabNet TeamForge 17.4's Announcements.

- CollabNet Enterprise Edition URL: <https://project.domain/servlets/ProjectNewsAdd...> is re-directed to <https://host/sf/sfmain/do/myPage?selectedTab=news>.

### Help

CollabNet Enterprise Edition Help page URLs are redirected to appropriate help pages on "help.collab.net".

- CollabNet Enterprise Edition URL: <https://project.domain/nonav/servlets/HelpTOC> is re-directed to <http://help.collab.net/index.jsp>.

## Moving to a CollabNet TeamForge project from a CollabNet Enterprise Edition project

If you have used the tools in CollabNet Enterprise Edition, you'll find the equivalent tools in TeamForge familiar.

Most of the tools and their functionality are similar between CEE and CTF 17.4. However, within each tool, some work-around might have been done to bring over your CEE project information.

This table shows how tools in CEE map to CTF 17.4 tools:


CEE Tool	CTF 17.4 Tool
Projects	Projects
Category (Project)	Category (Project Grouping)
Discussion Services	Discussion
Docman	Document
Issue Tracker	Tracker
PCE	PCE
Wiki	Wiki
Project Tracker	Project Tracker (Integrated Application)
Framework	Roles and Permissions
SCM	SCM
Remote Publishing	Remote Publishing
Project Dashboard	Project Dashboard

<b>CEE Tool</b>	<b>CTF 17.4 Tool</b>
Announcements	News

## Do I have to use the administrator-given password always?

No, you don't have to use the administrator-given password beyond your first login into CollabNet TeamForge.

The site administrator may have provided you with a user name and password after creating your user account in CollabNet TeamForge. When you login using those credentials, you might be asked to change your password for security reasons. At this time, you can set a password of your choice.

 **Tip:** You will not be asked to change your password if you had created your own account, or if a password-request had been raised for you or if an administrator had reset your password.

## What software configuration management tools are available in CollabNet TeamForge ?

In CollabNet TeamForge, users can browse the contents of a project's source code repositories and view detailed information about code commits, changed files, and associations with other CollabNet TeamForge items.

CollabNet TeamForge is not intended to replace your SCM tool. Code must be checked in using Subversion, CVS or GIT.

## How do I determine which version of TeamForge is currently running?

To determine the version of TeamForge you are running, click on the question mark icon in the upper-right corner of the user interface and choose About TeamForge from the drop-down menu.

You can also find the version number inside the following files:

```
* $SOURCEFORGE_HOME/etc/version
* $SOURCEFORGE_HOME/etc/integration-version
```

## Why can't I reply when someone comments on my review?

By default, TeamForge doesn't add the review initiator to a review, which would be needed to facilitate this.


To work around this, simply add yourself as an optional reviewer when creating the review. Alternatively, some customers have chosen to create a forum in TeamForge, and then include the forum's posting address in the review notes.

## How does TeamForge support dynamic planning?

TeamForge helps you maximize your team's effectiveness by keeping you in close touch with the multiple moving targets facing your project.

In a development project, each piece of the picture constantly changes in response to changes in the other parts. As you continue to iterate through the process, a feedback loop like this takes shape. (Click a node for more detail.)

The new Dynamic Planning features in TeamForge 17.4 give you a more open and extensible platform that integrates and centralizes the software development tools necessary in modern application life cycles.

 **Note:** The Dynamic Planning features helps you effectively utilize Agile--like processes, but you can use any process model you like.

### Tracker summary screen

TeamForge 17.4 incorporates a new Tracker Summary screen display to accommodate planning folders, treeviews, and multiple tracker viewing. The Tracker Summary section is available at the top of the screen with summaries of open and closed artifacts as well as a summary of open artifacts by priority. The Planning Folders section is located at the bottom of the Tracker Summary screen. This section includes summaries of open and closed artifacts, a summary of open artifacts by priority, and a summary of Effort for each planning folder.


## Tree view


TeamForge 17.4 incorporates an expandable and collapsible tree view of the planning folder hierarchical structure to display parent/child relationships of artifacts.

The tree view allows for the viewing of artifacts in a hierarchical structure and displays parent/child relationships across multiple trackers.

## How do you measure "effort?"

Effort is a uniform span of time for measuring work on your product. Defining the unit of effort is an essential part of planning your work.

TeamForge 17.4 shows you the effort that has been estimated, and actually expended, for each artifact on the tracker summary screen. Parent artifacts can automatically add up these effort figures from their child artifacts' effort figures. The calculator icon indicates that the artifact's effort is a sum of its child artifacts' effort within the project. When the effort from children artifacts in projects across the TeamForge (foreign children) is included in calculations, the icon,  appears.

 **Note:** When a tracker is disabled, artifacts from that tracker do not contribute to the effort totals calculated for any planning folder they are in.

The unit of effort should be something that makes sense in your environment. TeamForge does not require any particular unit. Your unit of effort might be hours per person, days, weeks, or something else. (If you are using a scrum-based project methodology, you may have opted to measure effort in relative terms, using points (story points), in which case you can leave the **Effort** fields blank.)

For example, some teams use the "ideal hour" as their standard unit of effort. To define an ideal hour, consider all the activities in a standard work day that must be done but don't directly contribute to development: installing and configuring tools, eating lunch, responding to email and instant messages, providing customer support, etc. For every hour of direct development work, how much time goes into these activities, on the average? If the answer is about a half hour, then your ideal hour is 1.5 clock hours.


Now consider a task that you judge to represent about four hours of direct development work. The value you'll enter in the **Effort** field is 6, because for each hour of direct development you'll need an extra half hour to make that development work possible.

In an environment with a lot of overhead -- for example, a group that relies on a very complex tool set -- your ideal hour might equal two clock hours, or perhaps much more. This is not in itself a problem: the point is not to suppress needed activities, but to plan realistically, in order to reduce the need for routine scheduling adjustments and to forecast more reliably.

(However, when units of effort seem radically out of line with reality, this may be a clue that something in the environment could be better optimized.)

For a quick view of the effort values you are working with, check the **Planning Folder Summary** page.

- The estimated, remaining and actual values are listed for the planning folder as a whole.
- The **Est** column shows the original estimated effort value for each artifact in the folder.
- The **Rem** column shows the value remaining for each artifact after the latest update to the artifact.
- The **Act** column shows the actual effort expended for each artifact, calculated after the work is done.

 **Tip:** To make your tracking easier, consider having TeamForge automatically generate a running total of estimated effort for a whole hierarchy of user stories or tasks.

## Why is TeamForge Explorer forgetting my settings and not logging me in?

Certain conditions can cause TeamForge Explorer to stop functioning properly and to start behaving oddly. This can manifest itself as the app not logging you into TeamForge, or forgetting your settings, or other general strange behavior.

To correct this behavior, you can delete your stored profile from within TeamForge Explorer.

If deleting the profile doesn't correct the behavior, try deleting the profile again, then uninstall and reinstall TeamForge Explorer.

If that still doesn't fix it, delete the profile, uninstall TeamForge Explorer, and delete the registry key `HKEY_CURRENT_USERSoftwareVA ?SoftwareSourceForge ?ExplorerWindowsMainWindow` (and all subkeys). Then reboot, then reinstall TeamForge Explorer.

## What is the file size limit on uploads or attachments in TeamForge?

TeamForge has a 2GB file size limit on uploads or attachments.

If you attempt to upload a file larger than 2GB using a Windows-based browser, you will not be able to upload the file and you may not receive an error message.

## How can I import tasks into TeamForge Task Manager?

Tasks can be imported into the Task Manager by using the CollabNet TeamForge Task Management Plugin for Microsoft Project.

## Can I associate objects of different projects?

Yes, you can associate an object (for example from document to subversion commit) in one project to an object in another project if you have access or are a Site Admin.

You must have permission to view the object that you're associating with before you can associate with that object, unless you are a Site Admin.

## How can I find the number of files in a repository without checking it out?

To list the files in a repository use the command: `svn ls -R file:///svnroot/REPONAME | wc -l`  
(Requires local access to server)


## What happens when log files get too big?

Log files can grow very large over time. To maintain reasonable log file sizes, you can rotate logs on a schedule.

When you rotate logs automatically, live logs are archived every day at 00:00.

Archived logs are stored in compressed form in a directory alongside the live log. For example, if live logs are stored at `<LOG_DIR>/{apps, apache, ...}`, then compressed log archives are stored at `<LOG_ARCHIVE_DIR>/{apps, apache, ...}`.

The directory structure of the log directory is preserved in the log archive directory.

 **Note:** Empty log files are not compressed.

## How can I make the project pages invisible for any user who is authorized to see the project?

This can be customized by a Project Administrator.

The Project Administrator can set this using these steps:

1. Click the link 'On' next to Configure, on the top right-hand side of the project home page.
2. Scroll to the page you want to hide and click the Edit Properties icon. This can be found on the right side of the title bar for the page you want to edit, next to Edit or Create buttons. The tooltip displays if you hover over the tool.
3. Choose 'Hidden - Visible only to project administrators in configure mode'.
4. Click Save.

## How do I make the monitoring messages be sent from Forge Administrator?

You can change the default behavior for site options by changing the value from "false" to "true" in this statement: #  
MONITORING\_EMAIL\_FROM\_ADMINISTRATOR=false

If the site option MONITORING\_EMAIL\_FROM\_ADMINISTRATOR=true, then "From:" field is the Forge Administrator, else it is from the user who made the change that initialized the monitoring email.

## How can I access the native install version of TeamForge?

You will need to contact your account manager.

## How do I enable/disable path-based permissions via SOAP?

The Path Based Permissions (PBP) are handled via the \*roleList, \*Cluster methods of rbacAppSoap. To enable PBP, use the "scm\_fgp" (fine grained permissions) argument to addCluster.

See the below psudeocode.

```
from com.collabnet.ce.soap60.webservices import *
    from com.collabnet.ce.soap60.webservices.ClientSoapStubFactory
import getSoapStub
    from com.collabnet.ce.soap60.types import *

hostname = "http://server/"
username = "admin"
password = "admin"
project = "proj1007"
roleName = "tracker"

sfSoap = getSoapStub(cemain.ICollabNetSoap, hostname)
sfSession = sfSoap.login(username,password)
rbacAppSoap = getSoapStub(rbac.IRbacAppSoap, hostname)

roles =
rbacAppSoap.getRoleList(sfSession,project).getDataRows().tolist()
roleId = None

for row in roles:
    if row.getDescription() == roleName :
        roleId = row.getId()
        print "found tracker role, %s" % roleId

if roleId == None:
    raise Exception("Cant find role ")

clusters = rbacAppSoap.listClusters(sfSession,
roleId).getDataRows().tolist()
for row in clusters:
    print row.getFolderId(), row.getOperationClusterName()
    if row.getOperationClusterName() == "scm_commit":
        print "found target!"
        rbacAppSoap.removeCluster(sfSession, roleId,
row.getOperationClusterName(), row.getFolderId())
        rbacAppSoap.addCluster(sfSession, roleId, "scm_fgp",
row.getFolderId())

sfSoap.logout(username,sfSession)
```

## How do I enable post-commit logging?

You do this by editing the post-commit.py file.

Edit the `/opt/collabnet/teamforge/runtime/sourceforge_home/integration/post-commit.py` file.

Search for `log.setLogging(False)` and modify the value from `False` to `True`.

## How do I put a notice to my users on the project home page?

To do this you will need to create a news item within the project. News items are posted and displayed on the project home page. News items are also displayed on the TeamForge home page.

To post a news item:

1. Navigate to the home page of the project in which you want to post the news item.
  - From within the project, click **Project Home**.
  - From anywhere in TeamForge, choose the project from the list of projects available under **My Workspace > My Projects**.
2. In the Project News section, click **Add**. The **Create News Post** page is displayed.
3. Enter a title for the news item.
4. Enter the text for the news item in the message field (up to 4000 characters including spaces), then click **Add**.

The news item is submitted. It is displayed on the project home page and the TeamForge home page immediately.

## What is a private IP address and what are the private IP ranges?

Any IP address that falls specified ranges is a private IP address and is non-routable on the Internet.

These addresses are reserved for use only within private/corporate network and cannot be seen outside the private networks. These private addresses are translated at the company's firewall into an external (public) IP address, which will be some address that does 'not' fall within these ranges.

10.0.0.0/8=10.0.0.0 – 10.255.255.255

192.168.0.0/16=192.168.0.0 – 192.168.255.255

172.16.0.0/12=172.16.0.0 – 172.31.255.255

An address is Private if it starts with:

10 or

192.168 or

172.[16-31]

It is possible for anyone to see their external (public) IP by going to any one of a number of sites that provide this information as a free service. One example that's easy to remember is <http://whatismyip.com>.

## What is the suggested log configuration for a production system?

To troubleshoot installation issues, the default log4j configuration is set to `DEBUG`. This can cause the log files to become quite large. Once your system is successfully installed and in use, you should drop the log levels down to `INFO`.

See [Change the logging level on your site](#) on page 720 for how to do this.

If you still have a problem with over-large log files, you may want to set up log rotation. See [Rotate TeamForge log files](#).

## What is the difference between a stagger and normal header in query result heading settings?

The normal and the stagger header preference setting allows you to set a standard header for your query view or a staggered header.

For more information about queries in Issue Tracker see: [Query database of issues](#)



## Can the query result be listed without the issue id?

No. Query results can never be listed without the issue id because issue id is the only field which identifies each issue uniquely.

## How do I remove a news item?

You can delete any news item that you no longer want displayed on the project home page. Deleting a news item from a project also deletes it from the TeamForge home page.

To delete a news item:

1. Navigate to the project home page.
  - From within the project, click **Project Home** in the project navigation bar. From anywhere in TeamForge, choose the project from the **Projects** menu in the navigation bar.
2. In the **Project News** section, click **Delete** next to the news item that you want to delete. When prompted, confirm that you want to delete the news item, and it will be deleted.

## How do I remove the OCN link from my TeamForge instance?

Remove the OCN link from your TeamForge instance by commenting out the header.ocn file in menu\_bar.vm.

Remove the OCN link from your SourceForge instance by editing the menu\_bar.vm file in the following directory:

```

$SFHOME/runtime/jboss/server/default/deploy/saturn.ear/sf.war/WEB-INF/classes/templates/menu_bar.vm file

```

Open the file and comment out or remove the following lines:

```

<li class="yuimenuBarItem">
  <a class="yuimenuBarItemLabel" id="openCollabNet" href="/sf/sfmain/do/ocn">${MESSAGE.get("common", "header.ocn")}</a>
</li>

```

## How do I send email from a specific sender address instead of the member address?

To send mail from a specific sender address change the `MONITORING_EMAIL_FROM_ADMINISTRATOR` token to true in the `site-options.conf` file.

TeamForge uses the member's email address as the sender address when it delivers the mail. To reconfigure this setting in the `site-options.conf` file, set the `MONITORING_EMAIL_FROM_ADMINISTRATOR` token to true, as shown in the following code sample:

```

MONITORING_EMAIL_FROM_ADMINISTRATOR=true

```


## Where can I download the TeamForge DTD?

You can download the TeamForge Document Type Definition (DTD) from the link below.

## How do I remove the build and test link from TeamForge pages?

The site administrator can remove these links by checking out and editing the branding repository from the look project.

To remove the build and test links from your TeamForge navigation panel, check out the branding repository from look project and reconfigure the links as shown in the following code sample:

 **Note:** You must be logged on as administrator to perform this task.

```

[branding_stage]$ cd branding
[branding]$ mkdir -p il8n/com/vasoftware/sf/il8n/apps/sfmain
[branding]$ echo "configure_build_and_test.systemUrl.default=" >
il8n/com/vasoftware/sf/il8n/apps/sfmain/application.properties
[branding]$ svn add il8n
A il8n
A il8n/com
A il8n/com/vasoftware
A il8n/com/vasoftware/sf
A il8n/com/vasoftware/sf/il8n
A il8n/com/vasoftware/sf/il8n/apps
A il8n/com/vasoftware/sf/il8n/apps/sfmain
A il8n/com/vasoftware/sf/il8n/apps/sfmain/application.properties
[branding]$ svn commit
" Old: "enter paragraph information and use html tags for bullet
points."


```

## Can I query the CollabNet TeamForge database?


Yes, you can query the database if you are a site administrator or have been given access to **System Tools** by another site administrator. You can raise a database query by using the **Admin > System Tools > Ad Hoc Database Query** page.

In the **Ad Hoc Database Query** page, select a data store, Operational Datastore or Datamart, type the "select" query and click **Run Query**. The query is executed against the CollabNet TeamForge database and the results are displayed.


For security reasons, the `sfuser` and `password_history` tables of the operational datastore are restricted for ad hoc querying from the **Admin > System Tools > Ad Hoc Database Query** page. Alternatively, use the following views, `sfuser_view` and `password_history_view`, for retrieving all other data but passwords.

 **Note:** The results of your query may be limited or your query session may time-out as per the settings in the `site-options.conf` file.

You can submit read-only queries of the format: `SELECT [FROM] [WHERE] [GROUP BY] [HAVING] [ORDER BY]`.

 **Tip:** You can use the following special keywords while drafting the query:

- "\d" or "show tables" - To list all the tables.
- "\d <table name>" - To view the description of a specific table.
- "select \* from <table name>" - To view all the contents of a specific table.

 **Note:** This feature is not supported by Oracle at the backend.

## Advantages of using the Apache TIKA parser library for indexing

Starting TeamForge 7.0, the underlying parser library for indexing has been changed from Stellent to Apache TIKA.

The Apache TIKA parser library has the following advantages over the Stellent parser library:

Issue	Stellent	Apache TIKA
<b>Stale process issue</b>	Parsing of corrupt or unrecognized files by the Stellent parser libraries often result in stale processes that consume swap space and add to the load on the system, which may at times lead to site outage. To manage such processes, you may choose to create and deploy stale process monitors and the stale processes, when detected, must be removed manually to prevent site outage.	Parsing of unrecognized or corrupt files by Apache TIKA libraries is robust and needs no manual intervention as there are no stale process issues.

Issue	Stellent	Apache TIKA
<b>Search queue processing speed</b>	It takes five minutes to timeout when the Stellent parser library encounters a corrupt or unrecognized file that it knows not how to parse. If there are more such corrupt or unrecognized files, more time is wasted by the indexer waiting for a response (or a timeout) from the Stellent parser, which in turn adversely impacts the search queue processing speed.	The Apache TIKA parser library is capable of determining whether a file it encounters can be parsed or not. As no time is wasted by the indexer waiting for a response (or a timeout) from the parser, the search queue processing speed is better with the Apache TIKA.
<b>Multiple processes Vs Single JVM</b>	For parsing files, the Stellent parser library spawns one subprocess per file. Meaning, the number of subprocesses is equal to the number of files to be parsed and it is possible that we may end up with the stale process issue as discussed earlier. As a result, if the Stellent processes consume more resources, other processes and applications are left with scarce resources.	The Apache TIKA, being a Java-based parser library, works within the JVM and makes the external resource pool available exclusively for other processes and applications. As the search JVM, where the Apache TIKA library lives, can also be separated starting TeamForge 7.0, it can be managed better.

## How do I connect to the Teamforge Postgres database?

You can use the *psql-wrapper* script to connect to the TeamForge application database.

### Usage

Run this script as below:

```
sudo [RUNTIME_DIR]/scripts/psql-wrapper
```

## How do I capture the output of "top" command?

Top is the realtime monitor of the running processes in a Linux system. To log the top running processes, use the following command: `top -b -n 1`.

-b = Batch mode operation - Starts top in 'Batch mode', which could be useful for sending output from top to other programs or to a file.

-n = Number of iterations limit as: -n number specifies the maximum number of iterations, or frames, top should produce before ending.

## How do I use the TeamForge updater to manage backups of old versions of TeamForge?

You can safely delete the items in your `<sourceforge_base_install_directory>/sourceforge_home/backups` as long as you are comfortable with your version of TeamForge, and have no desire to go back. This directory may also be safely omitted from your backup plan.

## How do I resolve timeouts when calling web services?

This is due to the requested operation taking longer then your client SOAP stack is configured to wait before throwing a timeout. You will need to reference your client documentation to see how to update the timeout properties of the connection.

For AXIS in java, you can do this via the `sun.net.client.defaultReadTimeout` property.

```
System.setProperty("sun.net.client.defaultReadTimeout", "600000"); //10
minute timeout, in ms
```

## What resources can be searched on a TeamForge site?

The resources that show up in search results depend on the context in which you are working.

Here is a summary.

Tool	Searchable resources
Projects	Project id, title/description, created by, project status
Project categories	Title, description
Discussions	<ul style="list-style-type: none"> <li>• Web UI: author, subject, content, attachment.</li> <li>• Email: we index the original email after we process it.</li> <li>• Topics: author, title and description.</li> <li>• Forums: title, description, author.</li> </ul>
Documents	<ul style="list-style-type: none"> <li>• Document folders: title, description</li> <li>• Documents; version comment, title, status, description, the attachment itself (all versions), authors</li> </ul>
SCM	Commit message, title, author
Tracker	<ul style="list-style-type: none"> <li>• Tracker title, description, author</li> <li>• Artifacts: title, group, category, customer, status, description, authors, tracker, all text flex fields, single-select fields, multi-select fields</li> <li>• Artifact attachments: the attachment itself and comments</li> </ul>
News	Body, title, author
File Releases,	<ul style="list-style-type: none"> <li>• Packages; title, description, author</li> <li>• Releases: title, description, author, maturity, status</li> <li>• Files: description, filename</li> </ul>
Tasks	Title, description, authors, planned
Users	Username, full name, email, status, details
Pages	<ul style="list-style-type: none"> <li>• HTML components: title, content</li> <li>• Subpages: page title, component title</li> </ul>
Wiki	Content of wiki page, using wiki syntax

## Locales and Languages in TeamForge

TeamForge supports English locale only. No support for Chinese Korean and Japanese (CJK) locales with TeamForge 17.1 and later.

## Customization

---

Questions on how to customize the TeamForge site.

Concept definition.

## What elements of a site can I customize?

You can customize the site home page and the default home page of every project on the site. You can also customize the menu bars, headers and footers of any page.

For more information about branding, see: [Customize anything on your site](#) on page 358.


## How does TeamForge use stylesheets?

The look and feel of much of TeamForge is controlled by cascading style sheets (CSS).

All default CSS styles can be customized to alter the look of the application. You can customize fonts (color, size, font face, etc.), links, backgrounds, headings, tables, tabs, and anything else that CSS can control.

The default TeamForge CSS file is `css/styles_new.css`.


New CSS files can be added to the `css/` directory and reference them via `templates/body_header.vm`.

 **Tip:** If you override an existing CSS file, it will be used instead of the default CSS file. So you must be sure to include all the default styles in your customized file. A best practice is to add any new or overridden styles to the bottom of the CSS file so that they can be easily identified.

## Can I substitute my own images for the default TeamForge images?

Every image in TeamForge can be edited or replaced by a new image file.

Images are stored in an `images` directory in the `branding` repository.

 **Note:** The default TeamForge image files are included in both the Quick Start and Advanced branding zip file.

### Examples

- To replace the masthead graphic, replace the image file `images/masthead/logo.gif`
- To replace the folder graphic, replace the image file `images/my/all_projects.gif`

## Can I use my own custom JavaScript on my site?

The JavaScript scripts used in TeamForge can be customized.

Existing JavaScript scripts are located in the `js/` directory in the `branding` repository.

New scripts can be checked into the `js/` directory then referenced from Velocity templates.

## Can I customize the web interface?

You can customize the way the web interface looks and functions through the use of Velocity templates.

Please note, however, that while Velocity as a technology is supported by CollabNet, Inc., the customizations themselves are not supportable. Any future upgrade of TeamForge may, in fact, break your customization.

Furthermore, making these types of changes to your installation is considered a customization and will impede our ability to support you. Should you experience issues and open a ticket with Technical Support, you may be asked to remove these customization to debug your issue.

## Integrated applications

---

Questions about integrated applications in TeamForge.


Concept definition.

## How is an integrated application described?

An integrated application is described using two XML files - a deployment configuration file and an application configuration file - that provide information to TeamForge about the configuration options exposed by the application.

In TeamForge version 6.1.1 and later, you have the ability to configure some integrated application settings using the user interface. You can also export these settings in XML format and make changes. To edit configuration settings, you would upload the XML file containing the updates.

### Integrated application settings

 **Note:** Some of the tags are internationalized so that the application will display languages based on the browser locale. See [Internationalize your integrated application](#) on page 690 for more information.

#### <name>

This is the title of the integrated application. When the integrated application is added to a project, the button that appears on the project pages has this name. This name must be unique -- you cannot use it for any other integrated application on the same TeamForge server.

This tag is used in both deployment and application configuration files.

#### <adminurl>

When an application has an administration screen for configuring its parameters, this field contains that URL. It is optional.

This tag is used in the deployment configuration file.

#### <baseurl>

This is the URL to which a user will be directed on clicking the integrated application button in a project.

This tag is used in the deployment configuration file.

#### <endpoint>

This is the SOAP endpoint for the integrated application. The endpoint contains the various methods exposed by the integrated application that are called during the lifecycle of TeamForge.

This tag is used in the deployment configuration file.

#### <gourl>

This indicates which URL must be used when an object id for an integrated application is specified (either via `Jump_to_id` or on the URL as `/sf/go/<objectid>`). This URL can support a couple of dynamic parameters.

- %o -- The object id entered by the user will be dynamically replaced here.
- %p -- The project id for the object entered will be dynamically replaced here.

For example, if the Go URL is `http://go.tourl.com/tracking?id=%o` and the object ID entered is XYZ123, then the URL will be replaced and redirected to `http://go.tourl.com/tracking?id=XYZ123`.

This tag is used in the deployment configuration file.

#### <category>

This is used to display the integrated application under a separate **Category** tab. It is an optional parameter that is manually inserted based on the requirement. If the category is not defined, the integrated application appears under the **Generic** tab.

This tag is used in the deployment configuration file.

**<isbaseapp>**

This validates whether the integrated application is a base application configured for the specified category. It is an optional parameter that is manually inserted based on the requirement.

This tag is used in the deployment configuration file.

**<is-search-supported>**

This validates whether the integrated application needs to be displayed as an object category in the drop-down list for search options in TeamForge such as **Jump to ID** and **Advanced Search**. It is an optional parameter that is manually inserted based on the requirement.

This tag is used in the application configuration file.

**<config-parameters>**

There can be any number of configuration parameters for an integrated application and they are displayed when associating the application to a project. These parameters are filled in by the project administrator and are available in the integrated application SOAP interface as configuration parameters. The integrated application gets a chance to validate these parameters and indicate back to TeamForge whether this project association is successful by passing in a "TRUE". It can return a "FALSE" if it doesn't want this project association to succeed. Each configuration parameter is placed inside the "param" tag, which can contain multiple elements to describe the parameter.

**<title>**

The internationalized title that appears for a project administrator to fill in while associating the integrated application to a project.

**<name>**

The Java variable under which the value for this parameter will be available on the integrated application.

**<description>**

The internationalized description that appears when a project administrator fills in or enters a configuration parameter.

**<default value>**

The default value for the parameter that will appear in the user interface during the association of an integrated application to a project.

**<display type>**

This is the type of display control used for the configuration parameter. We support "TEXT" for text fields, "CHECKBOX" for checkbox type controls, "RADIO" for radio buttons, and "SELECT" for select dropdowns. This field can also take an attribute that says what the value type for the field should be -- whether it should be an "Integer", "String" and so on. So if the field is expecting numbers, then entering "foo" as a value will throw a validation failure.

**<option>**

If the display type is "RADIO" or "SELECT", then these fields contain the individual options available for the display controls. This will contain a "name" attribute that will be sent to the integrated application when that option is selected from the UI. The value of this option should be an internationalized field as it is the value visible to the user.

**<editable>**

This specifies whether the configuration parameter should be editable once the integrated application is associated to a project. These configuration parameters are available when you add or edit an integrated project. If a parameter should not be "edited" post association, setting this to "false" will make it non-editable.

This tag is used in the application configuration file.

#### **<description>**

This is an internationalized string for the integrated application's description. It contains information for TeamForge project and site administrators to know what the application does.

This tag is used in the application configuration file.

#### **<id-pattern>**

When trying to link to an integrated application id, this regular expression gets used for mapping. By default (if no value is provided), it looks for alphanumerical characters; in case you need specific characters to be matched (for example, JIRA, which has hyphens in ids), this value is used.

This tag is used in the application configuration file.

#### **<page-component>**

These settings are used for Project Content Editors. The integrated application content can become part of the standard Page Component data that appears in project home pages. The settings indicate the type of information that will be available from the integration application.

##### **<input-type>**

This is the input type control for an integrated app Page Component. We only support 2 types now. Either "select" so that the inputs can be shown from a "SELECT" dropdown and the users will be able to pick a value from there. Else, it can be a "text" where a simple "text" field will be entered for taking the user input.

##### **<result-format>**

This is the format in which the output of Page Component is returned. This can be a "list" which indicates that it will be a Table like output. The integrated app will send the results in an XML format and the Integrated app framework converts this into a list of records. The other option is "html", where the output from the Integrated application is just displayed on the screen.

##### **<page-component-description>**

The description that will appear when you add an Integrated application Page Component (Link to the page where " Add component" is available)

##### **<page-component-title>**

The title that will appear when adding an Integrated application Page Component (Link to the page where "PCE Add component" is available)

This tag is used in the application configuration file.

#### **<permissions>**

This is a collection of permissions that are exposed by the integrated application. There could be any number of such permissions. These permissions will appear as a part of the project's roles (existing ones, as well as ones newly added) and can be assigned along with other tool permissions. You can map one of these permissions with a "dapMappedTo" attribute -- this indicates the permission to be used when a user logs in without authentication (for example, for public projects). Typically, this is the permission to read data so that it doesn't need a login name; it varies from one application to another.

This tag is used in the application configuration file.



**<prefix>**

If the "require-per-project-prefix" attribute is false, the value of this tag is used for identifying the integrated application in Go URLs, associations, and linkifications. If the "require-per-project-prefix" attribute is true, the value is used only for the "Host" project. Each project must fill in its value as part of adding the integrated application. Click here for [steps to add integrated applications to a project](#).

This tag is used in the application configuration file. The prefix can contain alpha-numeric characters and cannot be more than six characters in length. For more information about prefix, see [How does an integrated application interact with other TeamForge tools?](#)

**<require-per-project-prefix>**

An integrated application can indicate to TeamForge whether the object ids that it generates are uniquely identifiable across the entire application (if yes, the value for the attribute is "false") or whether they need to be project-specific (in this case, the value for the attribute is "true"). If an integrated application needs per-project prefix, you must enter the prefix value when the integrated application is added to a project.

This tag is used in the application configuration file.

**<require-scm-integration>**

This indicates whether SCM commits need to be validated. Some applications might have business rules which indicate that a commit can be made only if certain conditions are met. If the integrated application has any such rules, the value for the attribute should be "true". There are also a couple of methods to be implemented in the SOAP endpoint.

This tag is used in the application configuration file.

**<require-page-component>**

Some integrated applications choose not to expose details as Page Components. For those that don't, set this tag to "false" and for those that do, set it to "true". If the value is "true", you must provide the "page-component-details" tags as well.

This tag is used in the application configuration file.

**<servicetype>**

TeamForge 6.1.0 and earlier releases supported only SOAP as the mechanism to talk from TeamForge to the integrated application. TeamForge 6.1.1 and later support REST calls. The servicetype tag indicates whether the protocol used for communication is REST or SOAP.

This tag is used in the deployment configuration file.

For examples of how these tags are used in the integration of the Pebble blogging application, see [pebble-dep.xml](#) and [pebble-app.xml](#).

**How does an integrated application interact with other TeamForge tools?**

When you integrate an external application into your TeamForge site, the application can take full advantage of object IDs, links and Go URLs.

To look at how this works, we'll use the Pebble application as an example. Pebble is a blogging tool that you can quickly integrate with TeamForge.

**Object IDs**

Integrated application object IDs are of the form "prefix\_objectId". Object IDs uniquely identify a TeamForge object so that you can access and use it in different contexts. For example, to get to artifact `artf1234` quickly, you just enter `artf1234` in the Jump To ID box. In the Pebble tutorial application, the date of a blog post, in YYYYMMDD format, is used as the object ID.

A prefix is an alphanumeric string attached to the beginning of an object ID that TeamForge uses to manage object IDs from different tools. For example, in the Pebble app, `<prefix>_20100601` gets you a page showing all the blog posts in the project that were published on June 1, 2010.

In an object ID such as "prefix\_objectId", the "prefix" is case-insensitive, whereas the "objectId" is case-sensitive. For example, the two object IDs, "PT\_SC1" and "pt\_SC1", refer to the same object in TeamForge. Whereas, the two object IDs, "PT\_SC1" and "PT\_sc1", refer to two different objects in TeamForge. Here, PT and pt are case-insensitive and the SC1 and sc1 are case-sensitive.

The prefix can either be the one specified when an integrated application is added to a project by project administrator, or the one in the XML Application configuration file depending on the "require-per-project-prefix" setting. The "require-per-project-prefix" setting can be true or false. If it is false, each project integration would not need to provide a project prefix; so the one provided in the XML application configuration file takes effect. If the "require-per-project-prefix" setting is true, a prefix needs to be provided by the user during every project association.

The amount of information the prefix carries depends on the kind of application you are integrating into your TeamForge site.

- With applications that use object IDs, such as Project Tracker and JIRA, you can identify the project that the object belongs to from its object ID.
- For applications that don't have uniquely identified objects, or don't have the notion of "project," such as MoinMoin or Review Board, you can choose a prefix that's specific to the project where the integrated tool is used.

### Setting up multiple prefixes for integrated applications

At times, you may want to use more than one prefix for an integrated application such as the TeamForge EventQ. It is possible to have multiple prefixes set up for integrated applications. You must have the prefixes, separated by commas, included in the XML application configuration file and upload the file to your TeamForge site. For more information about uploading the application configuration file, see [Edit an integrated application](#).

Consider the following while setting up multiple prefixes for an integrated application:

- Prefixes, once set up using the XML application configuration file, cannot be modified.
- A prefix can be up to six alpha-numeric characters in length. However, the combined length of all the prefixes cannot exceed 128 alpha-numeric characters.
- The "require-per-project-prefix" must be set to false in the application configuration file. In case it is set to true, an error message appears when you upload the application configuration file.
- Do not use existing prefixes. You cannot upload an application configuration file consisting of one or more prefixes already in use in TeamForge.

### Go URLs

Go URLs allow a user to get to a particular object ID with a short, handy URL. To use this for Pebble, construct a URL like this: `https://mysite.com/sf/go/<prefix>_<date in format YYYYMMDD>`.

For example, if the Pebble tool in your project has the prefix PA, and you want to send someone all the blog posts published on app June 1, 2010, send them this link: `https://mysite.com/sf/go/PA_20100601`.

### Associations

The object ID can be used to associate objects with other TeamForge objects. For example, if you want to associate a document with the blogs published on June 1, 2010, go to the document's **Associations** tab and add an association to PA\_20100601 as the object ID.

### Automatic links

When you type text of the format `<prefix> _<date in YYYYMMDD>` in any TeamForge text field, the text is converted to a link. When you click the link you see the blog posts for that date, if any.

## How can I import tasks into TeamForge Task Manager?

Tasks can be imported into the Task Manager by using the CollabNet TeamForge Task Management Plugin for Microsoft Project.

## SCM commits in TeamForge with EventQ integration

Consider the following while integrating TeamForge and EventQ.

### Uploading the event handler JAR file

After adding *EventQ* to your site, it is mandatory to upload the custom event handler JAR file to get the pre and post commit notifications. [Click here](#) for steps to add a custom event handler.

### Implications of TeamForge-EventQ integration on SCM commits and associations

*Pre and post commit* notifications will be sent from TeamForge to EventQ only if the "*Association Required on Commit*" is enabled for the repository and if the "*require-scm-integration*" is set to "true" for the EventQ application.

Post TeamForge-EventQ integration, SCM commits will fail in TeamForge if all the following conditions are true and if the EventQ application is down or if there are errors while processing the commit request (on the EventQ's side):

- The **Association Required on Commit** is enabled for a repository.
- The **require-scm-integration** is set to "true".
- The commit message contains the integrated application's ID. In other words, the EventQ application's ID in this case.

SCM commits can also fail if the custom event handler exists in the TeamForge even after the removal of the EventQ application from the TeamForge site. It is highly recommended to remove the custom event handler once you remove the EventQ application.

## What is in the integrated application's configuration file?

Besides the basic information regarding the application, you can specify several important parameters in the configuration file.

The xml file can contain the following information about the integrated application:

- name - Name of the integrated application.
- description - Description of the integrated application.
- baseurl - Location where the integrated application resource locator is expected to go, when invoked.
- permissions - Set of permissions associated with the integrated application.
- prefix - The prefix identified for the integrated application.
- gouri - The Go URL of the integrated application.
- endpoint - The URL that contains the webservice descriptor language(WSDL).
- adminurl - The URL that provides integrated application administrator privileges through TeamForge site administration.
- Category - This is used to display the integrated application under a separate **Category** tab.
- isbaseapp - To validate whether the integrated application is a base application configured for the specified category.
- is-search-supported - To list the integrated application as an object category under the TeamForge search options, **Jump to ID** and **Advanced Search**.
- require\_prefix\_project - Confirmation if the integrated application requires a per-project prefix.
- require\_scm\_integration - The indication for the source code integration between TeamForge and the integrated application.
- Configuration Parameters - The parameters that the integrated application needs for initializing itself.
- editable - Confirmation if the value specified in the configuration file is editable using the **Edit Integrated Application** page.

- `servicetype` - The `servicetype` tag indicates whether the protocol used for communication is REST or SOAP. The protocol could either be REST or SOAP.

You can specify permissions for integrated applications and map the same to all the roles, as required.

- 👉 **Note:** The associations between CollabNet TeamForge and an integrated application can be created only from CollabNet TeamForge to the integrated application and not vice-versa.
- 👉 **Note:** To associate an object in an integrated application from within CollabNet TeamForge, use the [`<prefix_objectid>`] format. Successful associations appear hyperlinked.
- 👉 **Tip:** Each integrated application displays its prefix on moving the mouse over the application name in the tool bar.

## Lab Management

---

Use this background knowledge to help you support TeamForge Lab Management projects.

Concept definition.

### Can I get more systems for my project?

If your project needs more systems allocated to it, you must contact your Domain Administrator to increase your limit and allocate new systems to your project.

Your Domain Administrator will usually be a known support contact in your organization.

If additional systems are available in other projects, your Domain Administrator will be able to re-allocate those resources to you. If there are no free systems available, and none that can be reassigned, the Domain Administrator will be able to request additional systems from CollabNet, and then assign them to your project.

### What happens when I move a machine between projects?

When you move a machine from one project to another, the machine is not automatically rebuilt. It still runs the same profile and all the same software it was running in the original project.

There can be reasons for wanting to move a machine between projects without rebuilding it. For example, due to a reorganization, the same users and hosts may be providing the same service, but under a different department or project.

#### **Moving a machine does not reboot the machine or stop any running processes.**

If the old project was running processes on the machine that you do not want, you will have to stop them manually.

#### **Moving a machine changes the access rights on the machine, but does not end existing login sessions.**

Once the machine has changed projects, after a few minutes, the authentication will have changed so that users from the old project can no longer log in and users from the new project can now log in.

Existing user login sessions, however, are not terminated. Those users will not have `root` (UNIX) or Administrator (Windows) access any more.

#### **Moving a machine does not reset the contents of the `localadm` group.**

The `localadm` group is used to maintain a list of users on the system who are local administrators. This group is not emptied when a machine moves projects. It could contain users which you do not want as administrators on the host in the new project.

### What is an audit log?

Every action performed by the user in the TeamForge Lab Management system is recorded in the Audit Log.

For example, when a host is rebuilt using a profile, these are some of the details captured in the Audit log:

- The old profile.
- The new profile.
- How long it took to complete the rebuild process.

## What is the correct procedure for modifying a hosted Lab Manager profile?

All profile modifications must be done through the Lab Management UI, under Administration > Manage Profiles. Lab Manager profiles should not be directly modified and changes should not be committed to subversion.

To modify your profile, follow these steps:

1. In the browser, login to <https://mgr.cubit.domain.com/>
2. Click on Administration
3. In the left pane, click on Manage Profiles
4. Click on the profile (your\_profile\_name)
5. Click on the Packages tab and choose your options


## What is port forwarding?

Use port forwarding to let TeamForge Lab Management hosts connect to machines on other networks.

While TeamForge Lab Management is a secure and isolated environment, occasionally there are valid reasons to let other traffic in and out of TeamForge Lab Management.

Port forwarding consists of routing traffic from one network port on one host to another (same or different) network port on another host.


Take great care when exposing network services running on TeamForge Lab Management hosts to the outside world. Acquaint yourself with your organization's security policies, or develop them if you don't already have them, and make sure all services that are exposed comply with these policies. Even better, work with CollabNet to establish access controls around the access to TeamForge Lab Management, so only authorized hosts from within your enterprise -- and trusted partners -- can access TeamForge Lab Management.

 **Important:** CollabNet strongly recommends that any service that can yield a shell account on a TeamForge Lab Management node -- such as SSH -- not ever be port forwarded to outside the TeamForge Lab Management environment using this interface.

For CollabNet Hosted customers: CollabNet performs regular scans of its network, and if we see a dangerous or vulnerable network service configured, we may take any steps necessary to protect the overall security of the TeamForge Lab Management customer environment, including disabling the offending network service and the associated port forwarding.

## What is involved in administrating profiles?

If you are a project admin in TeamForge Lab Management with profiles assigned to your project, or if you are a TeamForge Lab Management domain admin, it is important that you have an understanding of how to administer profiles.


 **Important:** It is essential that you first have a solid understanding of how to retrieve and interpret profile details. Making changes to profiles without understanding the effects of what you are doing can cause disruption to your own project, and if the profile you are administering is a public profile, possibly many other projects that use this profile as well.

Profile definitions are versioned, which means the history of changes to profiles is preserved. This allows users to build with various versions of a profile, go back to earlier ones if a new one doesn't work for them, and upgrade to a newer version when available. As an administrator, you have the power to set descriptions on profiles and individual versions, block the use of certain versions, and so on.

Within the Profile Library, you can reach the Profile Admin page by selecting a profile in the Profile Library, and clicking on the **Admin** tab.

Properties can be version-independent or version-dependent.


- Version-independent properties are common across all profile versions -- they do not vary, regardless of the profile's version. These are basically the settings that are configured when the profile is first being added to Lab Management with profiles assigned to your project, or if you are a TeamForge Lab Management. They are displayed at the top of the Profile Admin screen.
- Version-dependent properties are associated with a particular version of a profile.

 **Note:** These properties, when set, will apply to the version in question, as well as to any subsequent version. Setting the property on a later version will overwrite the property on earlier versions.

These properties are displayed below the version-independent properties on the Profile Admin screen. Most of the properties that you can edit are version-dependent properties.

Only one version's properties are displayed at any given time for editing. To view multiple versions at the same time, use the Profile Details page.

### Version-independent properties

 **Note:** Some version-independent properties, set at profile creation time, cannot be changed.


There are three version-independent properties that you can change:

<b>Summary</b>	This is a brief summary of the profile. You can enter a more detailed description for each version in the properties section.
<b>Project</b>	The project that the profile belongs to. If you select "no project" using the "-" option, the profile will belong to the domain. The <b>Project</b> selection box is only available to domain admins.
<b>Is Public?</b>	This setting governs whether the profile is usable by all profiles, or just to members of this project.

### Version-dependent properties

There are many version-dependent properties that you can change for each profile. Some of these properties include:

- Those that specify hardware requirements needed to build this profile. This will make sure a profile is assigned to a host that can actually successfully be rebuilt with the profile.
- An option to specify whether the profile version can be used or not. Useful for marking a profile version "bad" so it cannot be inadvertently selected by users.
- The logo to associate with the profile version. This logo will be displayed in many locations throughout the system to easily identify profiles. Use the link to the Profile Logo Gallery to add and maintain logos.

 **Note:** When working with these properties, be sure you are operating on the correct version!

<b>Description</b>	A description for this version of the profile. Try to put some informative text here, so potential users of this profile will have some guidance as to what this version contains, or how it is different from other versions.
<b>Tag</b>	A symbolic name for a profile version. Tags can be used to make versions easier to remember, and can be moved around between profile versions, similar to a "tag" in Subversion or CVS. Valid characters for Tags are: letters A-Z, numbers 0-9, and underscores ; although a profile cannot be all numeric, and must contain at least one non-numeric character. The Tag HEAD is reserved, and

<b>Can new systems be built with this version of the profile? (also known as "Buildable").</b>	always refers to the latest version -- buildable or not -- of the profile.
	Tags are commonly used by project admins and other project leaders to instruct their users about the proper versions of profiles to use.
	For any number of reasons, you may wish to restrict profiles so that one or more versions of that profile are not buildable. For example, you may wish to force your users to always use the latest version of your profile: this would be easily accomplished by making all the profile versions not buildable except the most recent. You can change the buildability of a version at any time.
<b>Icon</b>	You can choose from any of the available icons for your profile, although a profile icon is strictly optional. Icons are not private to your project, and are shared among the whole domain, so do not upload anything too secret (or naughty!).
<b>Buildable CPU Types</b>	The types of CPU that can be used to build the profile, for example, "Xeon" or "UltraSparc IV". Setting this property is strictly optional, even if the profile has CPU type restrictions.
<b>Buildable CPU Archs</b>	The CPU architectures that can be used to build the profile, for example, "x86_64" or "sun4v". Setting this property is strictly optional, even if the profile has CPU architecture restrictions.
<b>Buildable Number of CPUs</b>	The number of CPU's required to run the profile. Setting this property is strictly optional, even if the profile has restrictions around the number of CPU's it can use.
<b>Buildable Hardware Models</b>	Specific hardware models which are required to run the profile, for example, "PDP-11". Setting this property is strictly optional, even if the profile has hardware model restrictions.
<b>Size (in GB)</b>	The minimum hard disk size, in gigabytes (GB), required to install and run the profile.

## Who controls which profiles can be used in a project?

As a project administrator, you can control which operating system profiles the users in your project can build hosts with.

TeamForge Lab Management's Profile Library gives ownership of profiles to individual projects, or to the entire domain. Profiles can either be public (available to all projects on the site) or private (available only to the project which owns the profile).

However, that fact that a profile is public, or belongs to the project, does not necessarily mean that it can be used to rebuild hosts. The profile must be specifically allowed for use in the project by a project administrator before it can be used to rebuild hosts.

Sometimes it is desirable not to restrict an entire profile, but only one or more versions of that profile. Individual versions of a profile can be prohibited by the project administrator of the project which owns the profile using the Profile Admin screen.

If the profile does not belong to your project, you cannot restrict the use of that profile at the individual version level: that decision is made at the discretion of the owners of the profile.

If you are the project administrator of the project which owns a profile, you can change the profile's public/private status at any time using the Profile Admin screen.

When you prohibit a profile from your project, this has no effect on hosts which are already built inside your project. They will continue to function normally running the profile they were already running. The owner of the system, however, will be unable to rebuild the system with its current profile, or any other profile which is not allowed in the project.

As a project administrator, you may wish to force your users to rebuild their systems once you have prohibited the use of a profile. On the Profile Summary page you can find a listing of all systems that are running each profile. This helps you track down systems running a profile that you have prohibited, and with your project administrator privileges, you can rebuild those systems with a profile of your choice.

## How are my project systems being utilized?


Regardless of the number of hosts in a project, in practice it is common to find shortages of free hosts. At the same time, there are almost always hosts which are under-utilized or even completely idle, which could be re-allocated or consolidated.

Finding and reallocating these hosts allows more efficient use of your project infrastructure.

### Project-Level Analytics

The **Analytics** tab in your project shows the following metrics. These metrics are the base available across all operating systems that TeamForge Lab Management supports.

- Load Average
- Processes
- CPU
- Logged-in Users

 **Note:** More metrics, including host-specific and operating-system-specific metrics, are available for individual hosts when you click on the host name in the results table.

To see a metric, click the name of the metric. You can see these time ranges:

<b>Daily</b>	Approximately the past 24 hours
<b>Weekly</b>	Approximately the past 7 days
<b>Monthly</b>	Approximately the last 30 days
<b>Yearly</b>	Approximately the last 365 days

If a given host says "no value" in any of its columns, this means that TeamForge Lab Management has been unable to collect this data over the requested time interval. If the machine is up, and not collecting data, contact a TeamForge Lab Management administrator to investigate why data collection is not working properly. If the machine is currently down, or was down during the requested time range, you will not be able to get any performance data for those times. There is no way to retroactively "catch up" data if collection is not working properly.

The data used to build each graph and chart presented to you can be exported in CSV (Comma-Separated Value) format, suitable for opening in any spreadsheet application. This allows you to build your own visualizations of the data to complement the ones TeamForge Lab Management creates.

### Beating The System: Dealing With the Possibility of Users Generating "Fake" Load To Make Machines Seem Busier

Since we publish the metrics for determining how busy TeamForge Lab Management thinks machines are, it is possible for irresponsible users to generate automated jobs which simulate a busy machine, even if the machine is really not being used for anything.



We encourage all users and administrators of the project to make sure people in your projects understand that this type of behavior is not acceptable and may lead to loss of privileges or other actions against them. Presumably, if someone has a good reason for wanting to keep a machine, it is good for their project, and your organization if they do so.

As administrators, we strongly recommend you take the time to understand and listen to your users' concerns about how machines are allocated. You may just need more machines in your project, or you may be overzealous about de-allocating machines once they drop below a certain usage threshold. De-allocating machines that seem to be not busy may seem efficient, but if those machines took their users a long time to set up, that might be counterproductive. Perhaps you can reach a middle ground and use virtual machines, or a smaller virtual or physical machine, for that user.

Finally, we always recommend talking to your users and trying to understand how they are really using their machines. The statistics that TeamForge Lab Management gathers are a starting point for managing your software development and testing infrastructure, but not the final word.

## What is host URL mapping?

Host URL Mapping in TeamForge Lab Management allows you to access web services running inside the TeamForge Lab Management environment from anywhere, using a simple and consistent URL, with optional SSL encryption services added on.

Host URL Mapping provides three major benefits:

- Maintains a consistent URL even if the service is moved to a different TeamForge Lab Management node or different base URL.
- Provides external access to resources that would otherwise be only accessible inside the TeamForge Lab Management environment. Because Host URL Mapping uses standard HTTP/HTTPS ports, you will not be blocked by firewalls that sometimes prevent port forwarding from working properly.
- Allows you to transparently add SSL encryption to web services running inside your TeamForge Lab Management environment.

If your application mixes absolute and relative URL's, host URL mapping dynamically rewrites the absolute portions of your URLs to help ensure that your applications display properly when they are mapped. While not perfect, this feature has been tested with a number of web applications and found to be effective.



**Caution:** Take the utmost care when exposing web services running on TeamForge Lab Management hosts to the outside world. Acquaint yourself with your organization's security policies, or develop them if you don't already have one, and make sure that all services that are exposed comply with these policies. Even better, work with CollabNet to establish access controls around the access to TeamForge Lab Management, so only authorized hosts from within your enterprise -- and trusted partners -- can access TeamForge Lab Management.

CollabNet strongly recommends that any web service that you expose be password-protected, or otherwise require authentication to access.

CollabNet Hosted customers: CollabNet performs regular scans of its network, and if we see a dangerous or vulnerable web service configured, we may take any steps necessary to protect the overall security of the TeamForge Lab Management customer environment, including disabling the offending service and the associated URL mapping.



**Note:** While you can convert non-SSL URL's into SSL using Host URL Mapping, you cannot map SSL URL's. You can use *Port Forwarding* to expose SSL URL's outside of the TeamForge Lab Management environment, however.

## Why doesn't URL mapping work for me?

My page doesn't look right. Images and graphics are wrong, or functionality doesn't work. The same page looks and works fine when viewed through Port Forwarding.

Host URL Mapper attempts to "clean up" HTML pages which pass through it in order to clean up absolute links. But it is not hard to construct an application that will slip through TeamForge Lab Management's filters and still not properly render all of its referenced objects inside of its pages. We make all reasonable attempts to clean up HTML,

but not all applications can be properly rendered using Host URL Mapping. This is especially true for applications which make heavy use of Javascript, ActiveX, Java applets, and other types of rich client-side web programming.

To test this, use either a direct connection to the host (if available) or a port forwarded connection to the host to see if this behavior is present on the original version of the page. If the original page does not have this behavior, the first step is to verify the **Dynamic Rewriting** level is set to **More Aggressive**. If that does not work, please file a support request with CollabNet to evaluate the page.

## How does host URL mapping compare with port forwarding?

URL mapping is good if you don't want your connection blocked, while port forwarding is good if you need non-HTTP services.

- Port forwarding is a facility for making any TCP or UDP network service available outside the TeamForge Lab Management environment. Host URL mapping only allows you to expose HTTP-based services.
- A major limitation of port forwarding is that many organizations' firewall security policies prohibit outgoing connections to arbitrary high ports. Because URL mappings all use standard HTTP/HTTPS ports 80 and 443, they will never be blocked. If you can access TeamForge Lab Management itself, you will be able to access any application configured with host URL mapping.
- Host URL mapping can automatically add SSL encryption to your web services that are not running SSL encryption. With port forwarding, if you want SSL encryption, you must set it up on each host.
- Using host URL mapping, you can expose only a part of a server's URL space. With port forwarding, the entire URL space is visible.
- Port forwarding does not need to rewrite links inside the HTML, so more web applications will work under port forwarding.

## How are virtual guests different from physical machines?

Virtual guests work like physical hosts, with some important differences.

### Creating a new virtual guest

The process of creating a new virtual guest is very similar to the process of creating a new physical host, with the following differences:

- You must select a virtual host -- a physical machine -- to run the virtual guest on. You cannot run a virtual guest inside of another virtual guest.
- The virtual guest's project cannot be set independently of the virtual host, and the virtual guest will always be in the same project as its virtual host.
- You are constrained, with hard limits, by the RAM and hard disk available on the virtual host. TeamForge Lab Management requires that each host have a minimum of 512MB of free RAM and 10GB of disk free.
- Even if you are within the hard limits for RAM and hard disk space usage, you are still sharing other resources -- notably disk I/O bandwidth, network bandwidth, and CPU cycles -- with the host machine and any other guests already on the virtual host. Before creating a new virtual guest on a virtual host, we recommend carefully examining the virtual host's system performance to make sure it can handle the additional load. Of course, if you do find out later on that performance of your virtual guest is not as good as you would like, you can always migrate the virtual guest to another virtual host.
- You are not as constrained on your selection of MAC address with virtual guests ; you can choose any available address in the allowed range. Valid values for MAC addresses for virtual guests in TeamForge Lab Management are 00:50:56:01:00:01 to 00:50:56:3F:FF:FF. Using anything outside of this range will either result in the host not being reachable on the network, or the host coming in conflict with another MAC address on the network.
- You cannot specify the architecture or chip type for the virtual guest, since those properties are inherited from the parent.
- You do not need to specify a Lights Out Management IP address for the virtual guest, since the IP address used to manage the guest is always the IP address of the virtual host.

## Disk size

While disk space is allocated to virtual guests at the time of virtual guest creation, it is not actually occupied on the host until it is needed by the guest. At the same time, TeamForge Lab Management does not keep an up-to-date count of exactly how much disk space is in use on the virtual host.

In practice, this means:

- You will likely have more disk space on your host than your virtual guests would indicate. But it is also possible to have *less* space than your virtual guests would indicate. For example, let's say you have a 100GB disk on your host, with two virtual guests, each with a 40GB disk allocated. But if you're only using 5GB of that 40GB in each guest, the remaining 70GB is unallocated. But on the flip side, let's say you allocated those two virtual guests at 10GB each, but they were using a total of 90GB on your local disk to store files. TeamForge Lab Management would let you make this allocation, but your virtual machines would crash when you tried to put more than a combined 10GB on them both.
- This translates into freedom in your management of disk space on the virtual hosts: you are free to temporarily "borrow" disk space from virtual guests that is not being currently used and use it for your virtual host.
- Along with this freedom comes the responsibility of being vigilant about maintaining sufficient space on your virtual hosts to always have enough disk space for any guests on the system. Monitor your usage carefully using Analytics page for the host.

## Modifying hardware parameters of existing virtual guests


Some virtual guest hardware parameters can be modified after the virtual guest has been created.

### Changing disk size

While changing the size of the disk is supported in TeamForge Lab Management, the change will not be reflected until you re-image your guest. And, in certain cases, it is possible for a change in the disk size of your guest to cause the guest to become unreachable, especially if you reduce the disk size. Therefore, we recommend that you change disk size on virtual guest only if you are going to immediately rebuild the guest.


### Changing RAM size

Changing RAM size is a very low-risk operation. In order for the change to take effect, the virtual guest must be rebooted after the change is made in TeamForge Lab Management. Reducing the RAM available to your virtual guest can make the system run much slower.

 **Note:** In between changing RAM size of the virtual guest and rebooting the virtual guest, you must wait approximately 5 minutes to insure your change is propagated.

### Changing the number of CPUs

Virtual guests can support either one or two processors. Two CPUs for virtual guests are supported only if the virtual host has at least two CPUs. Like changing a virtual guest's RAM allocation, changing the number of CPUs in a virtual guest is a low-risk operation, and the virtual guest must be rebooted after the change has been made in TeamForge Lab Management.

 **Note:** In between changing number of CPU's for the virtual guest and rebooting the virtual guest, you must wait approximately 5 minutes to insure your change is propagated.

## What is involved in migrating a virtual guest?

When you migrate a virtual guest to a different host, there are a few things to keep in mind.

- The user performing the migration must be a TeamForge Lab Management Domain Admin.
- The current virtual host and guest, and the desired new virtual host, must all be in the same project. In other words, you cannot change both the project allocation and the parentage of a host in a single step.
- The target virtual host must have enough RAM, disk, and CPU to accommodate the virtual guest you wish to migrate to it. The rules governing the RAM, disk and CPU needed are the same as if you were creating a new virtual guest on the host.

TeamForge Lab Management calculates these values for you and presents you with a drop-down list of hosts that fulfill all these criteria. But if you are wondering why a host is not a valid potential new virtual host, see the host's current hardware parameters and allocations on the **Host|Admin** page for the virtual host.

- If your TeamForge Lab Management installation is composed of more than one subnet, both the virtual guest being moved and the target virtual host must be on the same subnet. If you are not sure if your installation has more than one subnet, check with your local TeamForge Lab Management support contact.
- Only one virtual guest at a time can be migrated to a given virtual host, although there is no limit on the number of virtual guests that can be moved from a given virtual host. The reason for this restriction is that moving a virtual guest to a host is a very I/O-intensive operation that will consume all available disk I/O on the destination host. Allowing more than one simultaneous move will not be any faster, and increases the risk of over system instability for the virtual host.
- If more than one virtual guest migration to a given virtual host is started at the same time, they will complete serially. The exact order of the migrations may not correspond exactly to the order in which they were initiated and cannot be accurately predicted.
- Occasionally, after the migration has completed successfully, the virtual guest will require an extra reboot to start up properly.

## What's a good way to read log information?

To read log entries, download the log as a CSV file.

You can also filter the logs using the **Filter** option at the top of all audit log screens.

On any log page, you can click **Download all log entries in csv format** to download the corresponding CSV file.

- 👉 **Note:** The CSV file is formatted in Microsoft Excel format. If you import it into OpenOffice, set **Text Delimiter** to " (double quotes).

## Email

---

Questions about email settings in TeamForge.

Concept definition.

## Can I customize my site's email notifications?

The default text for CollabNet TeamForge email notifications can be customized.

Email templates are located in the `templates/mail` directory in the branding repository.

Some email templates that are commonly customized are:

- `user_welcome.vm`
- `account_request_rejection.vm`
- `project_approve_pending.vm`
- `user_forgot_password.vm`

Email template formatting follows the standard Velocity conventions:


```
user_welcome.vm
```

```
##subject
Welcome to CollabNet TeamForge 5.0!
##subject
##body
The Welcome message goes here...
##body
Email subject line
Email message text
```

## Can I specify an alternate email address?

Yes, of course you can specify one or more alternate email address. CollabNet TeamForge supports user profiles with one primary email address and up to three secondary email addresses.

The email address specified while creation of a user account is considered the primary email address. The alternate email addresses are optional and can be specified while updating the user profile.

 **Tip:** To add your secondary email addresses, use the **Admin > Users** page to update your profile.

## How can I check if port 25 is open?

If you know the mail server is up and running, check whether you can talk over port 25 to your mail server. This can be done using a one-line command: telnet <appserver name> 25 Substitute the <appserver name> with your own server.

Once you type this into your DOS window and hit return, you should see some sort of response from your mail server, as shown below:

```
Trying 208.75.196.84... Connected to cu190.cubit.sp.collab.net (208.75.196.84). Escape character is '^]. 220
cu190.cubit.sp.collab.net SMTP Server (JAMES SMTP Server 2.2.0) ready Mon, 27 Jul 2009 06:38:20 -0700 (PDT)
```

## When a discussion forum is set up, do all members receive a notification mail?

Yes. A mail informing users about the creation of discussion forum is sent to all the members of the project who have the **Discussions** (check box) selected as a monitored application.

Users can enable **Discussions** as a monitored application. To do that:

1. Go to **My Workspace > My Page > Monitoring** and select a project from the **Edit Monitoring Subscriptions and Preferences** pane.
2. Select the **Monitored Applications** tab.
3. Select the **Discussions** check box and click **Save**.

For more information about discussion forums, see [Administer forums and mailing lists](#) on page 77.

## Does the project owner receive a mail when a discussion forum is started in his project?

Yes. Whenever a discussion forum is started the project owner gets an email.

## I am unable to edit a specific artifact via email, but I can via the web UI. Why is this?

There may be work flow rules applied to the tracker that require specific fields to be set. As you can only define the artifact title and description via email, the artifact creation fails. If you wish to create artifacts in this tracker via email, the tracker admin will need to disable these work flow rules.

In addition, you may not be able to edit artifacts belonging to a tracker via email, if the tracker has two mandatory flex fields with parent-child relationship, but no default value relationship between the default values of the parent and child flex fields. In such cases, it is recommended to always establish the default value relationship between flex fields that have a parent-child relationship.

## How do I set up a local alias via James?

In situations where you need to obtain a SSL certificate for your domain, and your SSL certificate provider only permits you to use addresses related to your TeamForge domain, it may be necessary to generate an email alias from within TeamForge. Since there is currently no way to do this through the UI, you'll have to do it from the James administrative interface.

First, you'll need to connect to the James administrative interface on your system. If you've followed our best practices guide in our knowledgebase, you'll know that you should have port 4555 firewalled to everyone but localhost. SSH to your TeamForge server, and then issue the following command:

```
telnet localhost 4555
```

This will bring up the Remote Administration Tool:

```
[root@appl root]# telnet localhost 4555
Trying 127.0.0.1... Connected to localhost (127.0.0.1).
Escape character is '^]'. JAMES Remote Administration Tool 2.2.0
Please enter your login and password
Login id: admin Password: (text is echoed locally)
Welcome admin. HELP for a list of commands
```

First, we'll need to add a new user:

```
adduser <username> <password>
```

Then, we'll need to set the forwarding address of that user

```
setforwarding <username> <email address where you want email to go>
```

finally, we'll exit the James administrative interface.


```
quit
```

Your changes should be in place.

## Does TeamForge support using /etc/aliases for local mail delivery?

No, TeamForge uses the James SMTP server, which does not use the /etc/aliases file.

To enable local mail aliases, you will need to configure user mapping in the XMLVirtualUserTable in the /opt/collabnet/teamforge/runtime/james/apps/james/SAR-INF/config.xml file.

 **Note:** Please note that while the James SMTP server is used as part of TeamForge, customizations such as these cannot be supported by CollabNet.

## How can I stay informed about events on my CollabNet TeamForge site?

To keep up with changes, you can be notified by an automatic email when an item you are interested in is updated.

Monitoring items lets you stay up to date automatically on all changes without having to log into CollabNet TeamForge and check the status of each item.

You can monitor individual items, folders, and entire applications in each project for which you are a member. You can configure the frequency of email notifications to suit your personal preferences, and suspend monitoring messages entirely if you are out of the office or for any reason do not want to receive messages.

Items that can be monitored include:

- Entire applications, such as all tasks or all documents.
- Folders, such as task folders or document folders.
- All items in a tracker, forum, or package.

- Individual items, such as a task, a document, a tracker artifact, or a release.

You can tell by looking at each item whether you are currently monitoring it and how to begin or end monitoring it.

- For each item, list of items, or folder, you see a **Monitor** button or menu option.
- If you are currently monitoring an item, the **Monitor** option is replaced by a **Stop Monitoring** option, and the monitoring icon is displayed.

Whenever you create an item or edit an item, you automatically begin monitoring that item.

You do not receive monitoring notifications for your own changes to a monitored item.

If you have the appropriate permissions, you can see who is monitoring an item.

You can add other users to the monitored item. After a user is added to a monitored item, the user can continue monitoring the item, configure monitoring preferences for the item, or stop monitoring the item.

## How can I reduce the number of emails from monitoring a folder?

Monitoring a folder might spam the mailboxes with emails that are generated for every small change.

Email notification preferences under the **Monitoring** sub-tab under My Workspace can be modified to Send Daily Digest Email, which provides a summary of the email alerts.

1. Go to **Projects > My Workspace**.
2. Select **Monitoring** from the **My Page** menu.
3. In **Edit Monitoring Subscriptions and Preferences**, click All Projects and click on **Email notification preference**.
4. For each project, the monitoring folders under various components (Tracker/File Releases/Tasks/Discussions/Wiki) can be customized to Daily Digest Email.

## Can TeamForge accept email for more than one domain?

You can configure James to accept email for more than one domain by adding the additional domains to the `<servername>` section in the config.xml file.

Add the domains to the `<servername>` section of this file: `/opt/collabnet/teamforge/james/james-2.2.0/apps/james/SAR-INF/config.xml`.

Around line 53, you should see the following:

```
<servername autodetect=""true"" autodetectIP=""true"">
  <servername>localhost</servername>
</servername>
```

You can add other host names for James to accept mail for by adding more `<servername>` blocks. The comments in the config.xml file explain this further. Please keep in mind that these changes may be overwritten by future TeamForge upgrades.

## How do I configure TeamForge to send mail on a specific network adapter in a multi-NIC configuration?

When a host has multiple NICs, James will try to do the right thing when sending mail. In some network setups, this is not correct, and manual configuration is needed.

James requires multiple changes to fully configure how it interacts with the network. Open the config.xml file, located in `$SF_HOME/apps/james/james-2.2.0/apps/james/SAR-INF/` for version 5.1.

Locate the '`<mailet match="All" class="RemoteDelivery">`' section. add a subnode '`<bind>$addr</bind>`' where \$addr is the ip address that James should be sending mail from.

Near that area, there is a `<servername...></servername>` section. Confirm/change the two autodetect options (autodetect, autodetectIP) to false. Next, add the fully qualified host name, and the ip address that will be used, to their own `<servername>` entry.

After the changes are complete, save the config.xml and restart the application.

## Reports

---

Questions about reports in TeamForge.

Concept definition.

### Why do I get performance issues when retrieving flex field information from Datamart?

Flex fields are stored in XML format in the flex fields table. Due to the complexity involved in parsing the XML data and retrieving the relevant flex field information, you can see some performance issues. To overcome this, the XML to non-XML conversion based solution or feature has been implemented to store the flex fields in non-XML format. This approach has redefined the flex field storage and retrieval mechanism and thus has improved the performance of reporting queries while retrieving flex field information from Datamart.

#### How to enable or disable this feature in Datamart?

To enable or disable this feature respectively, update the field “attribute\_value” to either true or false for the record where entity\_name is 'flex\_field\_xml\_to\_nonxml\_etl' and attribute\_name is 'ALLOW\_FLEXFIELD\_XML\_TO\_NONXML' in ETL Attributes table.

```
update etl_attributes set attribute_value='true' where entity_name =
'flex_field_xml_to_nonxml_etl'
and attribute_name = 'ALLOW_FLEXFIELD_XML_TO_NONXML';
```

#### Does a new job have to be executed to get the benefits of this feature?

No, the existing tracker initial job and tracker incremental job will do the required data processing.

#### What would happen if a parsing error is encountered during XML to non-XML conversion?

In such cases where auto correction is not possible for incorrect XML records, the incorrect records are to be corrected manually and to be updated with the status "closed". Once the next ETL process picks up this record and if it is valid, then the job removes the entry from this table, updates the status from "closed" to "processed" and populates the parsed data into new flex field bridge table. If the manually corrected xml is still invalid, then the job updates the status of that record from “closed” to “open”, so that the user can work on correcting the entry.

#### What are the changes related to custom reports development on flex fields?

The custom report developers need to query the new flex field bridge tables to get the benefits from the XML to non-XML conversion feature implementation.

#### How to handle the ETL job failure due to OutOfMemoryError: GC overhead limit exceeded?

For a permanent fix, see [Why am I getting an email specifying that the ETL job has failed?](#) on page 956

To fix this issue during runtime, perform these steps:

1. Stop the ETL service.

```
kill -9 <etl process id>
```

2. Open the file /opt/collabnet/teamforge/runtime/conf/set-env.sh.
3. Increase the JVM heap size in ETL\_JAVA\_OPTS to Xms512 and Xmx2048m.

```
ETL_JVM_OPTS="-Xms512m -Xmx2048m -server -XX:+HeapDumpOnOutOfMemoryError -
XX:HeapDumpPath=/tmp -verbose:gc -XX:+PrintGCTimeStamps -XX:
+PrintGCDetails -
Dsun.rmi.dgc.client.gcInterval=600000 -
Dsun.rmi.dgc.server.gcInterval=600000 -
Djava.awt.headless=true -
Dsourceforge.home=/opt/collabnet/teamforge/runtime/sourceforge_home -
```



```
Dsourceforge.logdir=/opt/collabnet/teamforge/log/etl -
Dapp.data=/opt/collabnet/teamforge/var -
Dapp.runtime=/opt/collabnet/teamforge/runtime -
Dapp.distribution=/opt/collabnet/teamforge/dist -
Dapp.log=/opt/collabnet/teamforge/log"
```

4. Start the ETL service.
5. Run this command to make sure that you see the updated memory setting in the command output.

```
ps -ef | grep etl | grep -v jboss
```

6. Run the ETL job again.

### How to redo the entire XML to non-XML conversion if required for some reason?

**Workaround 1:** Bootstrap the reporting service and execute the tracker initial job. As XML to non-XML conversion is part of initial job as well, this would redo the conversion process.

**Workaround 2:** Truncate the relevant tables and reset the previously processed artifact flex field key to 0. This would reinitiate the XML to non-XML conversion while executing the tracker incremental job.

Relevant SQL process:

```
truncate table stage_flex_fields_bridge;
truncate table text_flex_fields_bridge;
truncate table user_flex_fields_bridge;
truncate table ms_flex_fields_bridge;
truncate table ss_flex_fields_bridge;
truncate table date_flex_fields_bridge;
truncate table bad_xml_records;

update etl_attributes set attribute_value = 0 where attribute_name =
'PREV_RUN_ARTIFACT_FLEX_FIELDS_KEY' ;

commit;
```

## What kind of objects can I create reports on?

Reports can report on data from trackers, tasks, planning folders, and repositories.

Task reports can report on data in a single project or across multiple projects. Tracker reports can report on data in a single tracker or across multiple trackers in a project or multiple projects. Task and Tracker reports are in tabular format and grouped under Table reports. You can also create reports on data in planning folders and on data in source code repositories.

## How do I connect to the Datamart?

You can use the *psql-reporting-wrapper* script to connect to the datamart.

### Usage

Run the following command:

```
sudo [RUNTIME_DIR]/scripts/psql-reporting-wrapper
```

## How do I change the time to run the ETL jobs?

The *ETL\_JOB\_TRIGGER\_TIME* can be modified to specify a different time.

By default, the ETL job runs at 2:30 AM (local time) everyday. It is recommended to run this once a day to avoid any performance degradation of the Teamforge site. See [ETL\\_JOB\\_TRIGGER\\_TIME](#) on page 748 for more information.

## How can I check the status of ETL?

The `[RUNTIME_DIR]/scripts/collabnet status etl` displays the status of the ETL process.

You can get additional information about the various ETL jobs that are configured using the command `[RUNTIME_DIR]/scripts/etl-client.py -a`


## What can I learn from a burndown chart?

Burndown is the estimated amount of work that remains to be done in an iteration or a release, compared with the work originally estimated.

At any time, you can click on a planning folder (release, iteration or standard) and view the burndown chart which shows the work to be completed in terms of story points within the given time frame. This helps you see how the team is progressing towards "done". In the burndown chart of a standard planning folder ('Folder') you can see how the sum of the estimated and remaining effort and points (story points) for all contained and descendant features is changing within a given time frame by choosing the options available in the View Mode.


In both 'Release' and 'Folder' burndown charts, you can use the trend line shown in your burndown chart to do some useful things:

- Project the time when all the work for this iteration will be completed.
- Predict the amount of work your team can expect to complete during any given iteration.



## Release burndown chart

A release burndown chart shows an iteration-wise breakup of the work to be completed, that is, it shows how much work (in terms of points) is remaining at the start of each iteration.



The information you derive from your burndown chart will help revise the product scope, make accurate planning decisions, and refine implementation details.

In general, a burndown chart trends downward until it reaches zero. In practice, some events can reverse the downward trend of your burndown chart. For example, development work frequently uncovers a greater scope for a user story than was initially estimated. As a result, you'll revise your estimate of remaining work on that story, and your burndown figure may be pushed upwards.

The burndown chart uses whatever effort units you are using in the planning folder.

- 👉 **Note:** The story told in your burndown chart is only as reliable as the underlying data. The owners of individual task artifacts can help keep the burndown chart accurate by regularly updating their remaining effort and points (story points).

## What can I learn from a capacity chart?

The capacity chart shows the project team's judgment of how much work can be done with the resources available and the time period represented in the planning folder.

- When you compare average capacity with individual users' assignments, you can see which team members are relatively over- or under-assigned. You can also see how much of the estimated and remaining effort is unassigned.
- Comparing average capacity with the relative priorities of artifacts is one way to gauge your team's ability to deliver the work that the product owner has defined.
- When you view average capacity against the count of open and closed artifacts, you have another way to assess your team's likelihood of meeting its sprint commitment.

The data behind the capacity chart is updated in real time, so that your team can respond quickly to changes in effort estimates or relative priorities.

Capacity is expressed as a numeric value, using the same scale as you are using for estimated, remaining, actual effort accounts.

You can see the capacity of a planning folder in the planning folder summary view.

- 👉 **Tip:** A planning folder's capacity must be equal to or greater than the total estimated effort (or total remaining effort) for the artifacts in the planning folder. To avoid confusion, you may want to wait for estimates to emerge before setting the planning folder's capacity.

## How do I export reports?

You can export data of a Table report.

Clicking the **Export** button that is placed at the top right corner of the **View Report** page exports tracker or task report data to a .csv, .xml, .xlsx or tab-delimited file.

## What can I learn from an Open by Priority chart?

Dynamic planning is most effective when you know the relative priorities of work items you are tracking. The **Open by Priority** chart helps you check that your team is working on the optimal mix of artifacts.

Use the **Open by Priority** chart for a quick overview of the type of work your project is working with. A glance at this chart can suggest some broad generalizations, which you'll then want to test by examining your planning folder contents more closely.

- When the darker bars on the left side of the chart are high, you are probably working on the high-impact issues that come up earlier in a work cycle. Don't be surprised if the trend line in your **Burndown** chart is not yet sloping steeply downward to the right.
- When the lighter bars on the right side of the chart are high, you are likely to be working on cleanup issues and refinements. You may be approaching the end of the work cycle you defined in this planning folder. At this stage, the Closed slice of the **Open vs. Closed** chart is likely to be equal to or larger than the Open proportion.

- 👉 **Note:** The **Open by Priority** chart won't tell you how much work is involved in the artifacts that it measures. For that, check the **Burndown** chart.

## What can I learn from an Open vs. Closed chart?

Sometimes you just want a raw count of the number of work items you are looking at.

Use the **Open vs. Closed** chart for a quick overview of your team's progress in purely numerical terms. A glance at this chart can suggest some broad generalizations, which you'll then want to test by examining your other charts and your planning folder contents more closely.

The **Open vs. Closed** chart is good for some basic summary information. You may want to share this chart with executive sponsors who need to know how things are going in quantitative terms, but don't need a lot of detail about the types of work being addressed.

## How does TeamForge deliver activity reports?

The data in your reports comes from a special database that extracts live site data from the production database at intervals you specify.

You can specify the time at which the reporting data is refreshed from the production database. By default, the extraction takes place daily at 2:30 a.m. in the TeamForge application server's time zone. See [Schedule data extraction for reporting](#) on page 712.

The reporting database can be deployed on a separate machine to help channel load away from the application server. Historical data is available even if the application server no longer stores it.

## Where does the reporting data come from?

An ETL application extracts data from the live production PostgreSQL or Oracle database where the TeamForge site stores most of its critical data. (Information about reporting configurations is also stored in the production database.) Some data is also gathered from the file system.

## How is the production data converted into reporting data?


TeamForge extracts a snapshot of the production data, transforms it into a format that supports reporting requirements, and loads it into the datamart, which is optimized for fast retrieval. The Extract-Transform-Load (ETL) application is a Tomcat JVM running as a TeamForge service under the TeamForge integration server architecture.

## Where is the reporting data kept?

After the ETL app collects and processes the live site data, it is stored in a separate database called the datamart. If the TeamForge site uses a PostgreSQL database, then the datamart is also a PostgreSQL database; likewise for Oracle. The datamart uses a Star Schema-based design for tables.

## How are the reports shown in the TeamForge user interface?

The reports are rendered in the TeamForge UI using Adobe Flex.

 **Note:** When a site is upgraded, there will be a delay before reporting data is available to users, until the scheduled ETL run has occurred. Performing a manual ETL run immediately after an upgrade is not advisable, since it could consume a lot of system resources leading to performance problems.

## Why do I get duplicate records in a tracker report?

Yes, this is a known limitation.

For a tracker report, when you select more than one value from a multi-select, user-defined field as filter and if an artifact is associated with all of the selected values, then that artifact's record is duplicated for each of the selected values.

For example, assume that you have a 'Select User' multi-select, user-defined field with values 'User 1', 'User 2' and 'User 3' in a tracker report. All these three values are associated with 'artifact 1001'. Select all three values as filter and generate the tracker report. You will see 'artifact 1001' record being duplicated, that is, you will see three individual 'artifact 1001' records created for each of the three users.

## Subversion and CVS


---

Questions about Subversion and CVS.

Concept definition.

### What's the difference between a 'managed' and 'unmanaged' CVS server?

You can use TeamForge to control a CVS server, or you can run the CVS server stand-alone, or 'unmanaged'.

 **Note:** Only CVS servers can run unmanaged. Subversion servers are managed by TeamForge.

#### Managed server

A *managed server* gets its instructions from TeamForge to which it is integrated with. Use TeamForge for SCM tasks such as:

- Creating repositories
- Adding and removing users
- Enabling or disabling user access to the repositories

Both new SCM servers and those with existing repositories can be managed by TeamForge.

#### Unmanaged server

On an *unmanaged CVS server*, you use the CVS's own utilities to handle repository management, user management, and access permissions.

Add an unmanaged server if you know that all repository and user creation and management must be done manually. For example, set up an unmanaged CVS server when:

- A corporate IT policy prohibits applications such as TeamForge from making any changes to the users or repositories on the corporation's CVS servers.
- A CVS server has existing repositories that cannot be changed to adhere to TeamForge standards.

## Subversion replication in TeamForge

Here's some information on how replication could be useful in your TeamForge site and what to consider when you plan to set up a replica.

### Why replicate?

Typically, you would deploy a replica for one these reasons:

**Projects in remote locations with lower bandwidth or higher latency want the performance of a "local" server.**

Your company has a number of developers clustered together at a remote location. When you install a replica inside the LAN of these developers, they can greatly improve their Subversion performance and keep a lot of their traffic off the WAN. In this scenario, you would probably want a replica in each such location. Keep in mind that a replica can only be a proxy for one master -- so if your company has more than one Subversion master server, you may need more than one replica at each location.

**You want to reduce the load on the master server.**

For example, continuous integration tools can place a lot of load on the server and moving that load to a separate server can increase the response time for other users. In this scenario you probably only need to add one replica; you'd add it as close to the master as possible so that synchronization is quick. Of course the previous

point can be a factor here. If the continuous integration server is at remote location, then you would want to put a replica near the continuous integration server.

### Rules for using a replica

To create a replica in TeamForge, you start with Subversion Edge. A Subversion Edge wizard lets you convert the server into a replica of a Subversion server in TeamForge.

When you configure the replica server, you provide the TeamForge username and password to use for the replica. These are the credentials the replica uses when it replicates Subversion content. This user must be added to projects and given permissions to the repositories being replicated. Those permissions also control what parts of the repository will be replicated. So if you have folders that should not ever live on remote servers, you can set up path-based permissions and that content will never be replicated to the server.

If you forget to set up permissions, the replication will fail. However, there's no real harm done, and once you fix the permissions, you can do it again.

The replica user can be a normal user account -- it does not have to be an Admin account. If the replica is set up and maintained by an Operations team, they might want to just use an Admin account so that project teams do not have to worry about adding the user to the project or setting up permissions.

Permissions for end users accessing those repositories will follow the normal TeamForge rules.

## How do I move an existing CVS repository into TeamForge?

You can import and manage an existing CVS repository with TeamForge.

Do the following to move an existing CVS repository:

1. Stop CVS access to the old repo.

```
chmod -x /usr/bin/cvs
```

2. Tar the old repo.

```
cd /cvsroot/old_repo
tar zcvf /tmp/old_repo.tar.gz
cd..
mv old_repo /tmp
```

3. Restore CVS access.

```
chmod +x /usr/bin/cvs
```

4. Transfer the repo to the TeamForge CVS server.

5. Create the new repo from within TeamForge.

```
Browse to your project
Click the Source Code button
Create your new repo
```

6. Untar the old repo.

```
cd /cvsroot/new_repo
tar zxvf /tmp/old_repo.tar.gz
```

7. To synchronize permissions, perform the following steps.

- Login as a TeamForge site admin
- Click the Admin link
- Click the Integrations button
- Check the CVS integration you want

- Click the Synchronize Permissions button
8. Verify the new repo.
  9. Remove the old repo.

```
/bin/rm -r /tmp/old_repo
```

## How do I move an existing SVN repository into TeamForge?

If you have an existing SVN repo, you can manage it with TeamForge.

To move an existing the SVN repo, perform the following steps:

1. Stop SVN access to the old repo.
2. Dump the old repo.

```
svnadmin dump /svnroot/old_repo > /tmp/old_repo.dmp o mv /svnroot/  
old_repo /tmp
```

3. Restore SVN access.
4. Transfer the repo to the TeamForge SVN server.
5. Create the new repo from within TeamForge.
6. Browse to your project and click the **Source Code** button, then create your new repo.
7. Load the old repo.

```
cat /tmp/old_repo.dmp|svnadmin load /svnroot new_repo
```

8. To synchronize permissions, perform the following steps:

- Login as an TeamForge site admin.
- Click **Admin**.
- Click **Integrations**.
- Select the SVN integration you want.
- Click the **Synchronize Permissions** button.

9. Verify the new repo
10. Remove the old repo

```
/bin/rm -r /tmp/old_rep
```

## How do I enable Neon debugging in my Subversion client?

This is easily done in a Linux environment by editing the existing servers text file. On Windows it is not easily done and not recommended.

On a Linux system, edit the `~/ .subversion/servers` file by adding the line `"neon-debug-mask = 130"` (without quotes) to the [global] section of the file, making sure that you also un-comment the [global] line as well. Once Neon debugging is enabled, you should see much more output from each svn command.

Although Neon debugging is possible on Windows, it involves steps that are too complex for most end users to undertake, including compiling Windows binaries for the specific platform and manually handling any errors that arise during that process.

## How do I create a local backup of a remote SVN repository?

You can back up your remote Subversion repository to your local system by keeping a mirror copy of the repository locally using `svnsync`.

## Trackers

---

Questions about Trackers in TeamForge.

Concept definition.

### Can I assign an artifact or task to a group of users?

Unfortunately this is not possible with the current implementation of groups in TeamForge.

As a workaround, you can create a new user account and set its email address to that of a mailing list that all group members are subscribed to. Then the submitting user can assign the artifact to this user account, and the mailing list will be notified.

### How do I create an artifact through email?

To create a tracker artifact using email, send an email message to <tracker id>@<TeamForge server>.

You can find the tracker id on the **Artifact List View** page. Fields are mapped as follows:

- To: Tracker email address
- Subject: Artifact title
- Body: Artifact description
- Attachments: Attachments

### I don't need the priority and assigned-to fields, how do I get rid of them?

These fields are used internally by the system and cannot be deleted or disabled.

### Why do the open tracker counts differ from when I filter on 'Open'?

The 'Open' category on the summary screen is based on the meta status 'Open'. It includes multiple statuses; Open, Fixed, and any other statuses that are defined as equivalent to Open by the tracker admin for your project.

For example, suppose the Tracker Summary block shows 21 Open artifacts. However, when filtering on the 'Open' status, only five artifacts are displayed.

If you filter by 'All Open' you will see all 21 issues with statuses defined as Open. If you filter specifically by 'Open' you will see only the artifacts that specifically have the status 'Open'.

### How do I edit the predefined data in the various fields in the issue tracker?

Use **Configuration Options** to edit the required data in the fields.

For more information about customizing your installation see: [Customizing your TeamForge installation](#)

### Can I limit the number of characters for fields in a tracker?

No, it is not possible in the current version of TeamForge.

Currently, while creating the tracker custom flex field of input type 'Text Entry', you cannot control the maximum number of characters to be allowed. There is a 'Field Width' that takes values to control the size of the text field, but not to limit number of characters entered.

### Is it possible to move an artifact from one tracker to another?

Yes. It is possible.

Simply use the cut button to remove the artifact, then paste it to another tracker. You can also change the tracker type while editing the artifact using the **Edit Artifact** page. The destination tracker need not be in the same project, but if the tracker definitions differ data could be lost.



## How can I change the display order of the fields in my tracker?

Simply drag and drop the fields within your browser.

Horizontal rules may also be added in this manner.


## How do I create a parent or a child artifact?

A tracker artifact can be a parent or a child to another artifact.

- A child artifact can have only one parent.
- A parent artifact can have any number of children.
- A parent artifact cannot be closed if a child is open.

To create a new parent or child artifact of an existing artifact:

1. Click on the **Dependencies** tab of the artifact.
2. Click on **Create New Parent** or **Create New Child**. Fill in the form as you would for submitting an unrelated artifact.
3. Click **Next**.
4. Enter a comment that describes the relationship, if appropriate.
5. Click **Finish**.

 **Note:** If **Choose Parent** or **Create New Parent** is not visible, the artifact already has a parent artifact.

## What does the status of a planning folder mean?

The status of a planning folder communicates where it currently stands in the development process. You can set the status of a planning folder to values you define yourself.

The status of a planning folder can help project members make sense of what can be a complex development process. For example, consider these scenarios:

### Drop everything!

Just when you have prepared a planning folder and are ready to kick off a sprint, the product owner gets word of a more pressing set of features and asks you to postpone this work and get right to the newly revealed priority. You are worried that your project's momentum might lead project members to keep picking up tasks in this planning folder instead of switching to the new work.

Create a status called 'On hold', and specify it as an 'Inactive' status. When you give the planning folder your new 'On hold' status, project members whose view shows only active folders would not see this one in their planning folder list.

### It's a wrap

Your sprint has just concluded. You don't want any further development work to happen in this planning folder, but you do want team members and others to be able to peruse it for insight into what went right and what went wrong.

Create a status called 'Retrospective', and specify it as an 'Active' status. When you assign this status to your planning folder, participants will be able to confirm that the sprint in question is finished and their comments are invited.

### Status tips

- It is a good idea to choose a default status, such as 'Under construction', so that project members are not confused when new planning folders appear.
- You can organize your planning folder statuses in a way that makes sense for your team, such as alphabetically or chronologically.
- Delete statuses that are no longer being used, so you do not clutter up the selection.

- When you delete a status that is in use by one or more planning folders, the status of those planning folders is changed to 'None'. Planning folders that were created in an earlier TeamForge version also get a 'None' status at first.


## How is a project template structured?

When you create a template that includes project content, each tool brings in its own kind of structure, depending on the type of content it manages.

### Folders

Some tools, such as Documents, Tasks, and File Releases, can be thought of as folders that contain individual items, subfolders, or both.

For example, the Documents tool has a Root Folder, inside which you can create a tree of subfolders to organize your documents according to your project's needs. Every document lives inside one of these folders.

 **Note:** This works the same way regardless of the name of the root folder. For example, in the Tasks tool the root folder is called the Tasks Summary, and in the Discussions tool the root folder is called the Forum Summary.

In a project that has been in use for any significant time, users have probably created some number of documents, which they have shared by adding them to a folder in the Documents tool. When you create a project template from that project, you can choose to include those documents in the project template or not.

- If you choose to include the documents in the template, those documents will appear in any new project that is created using that template.
- If you do not include the documents in the template, new projects based on that template will include the Documents tool, including the Root Folder and its sub-folders, but none of the documents that existed in the original project.

### Not folders

The Reports and Wiki tools are not organized in folders. When you include the content from these tools in a template, new projects created from that template include a flat collection of all the Wiki pages or reports in the original project, with all their text and data.

## How do I create new values in the 'Reported in Release' field?

SFTF will populate that field automatically from the packages that are defined in the file release tool.

## What fields can I use in a tracker?

You can use fields provided by TeamForge, and you can create your own custom fields. Most trackers use a mix of TeamForge-provided fields and user-defined fields.

### TeamForge-provided fields

Some fields are provided by TeamForge for all trackers.

**System-defined fields:** Some TeamForge-provided fields are always present in every tracker. These are identified as "system-defined" fields in the field list.

- The **Artifact ID**, **Title**, and **Description** fields are system-defined fields. They are always present and always come first when you view an artifact.
- Some system-defined fields can be reordered. For example, the **Assigned To** field is a system-defined field, so it is always present, but you can move it around in your tracker's display.

**Configurable Fields:** Some fields are available for any tracker, but can be modified in different ways. These are called "configurable" fields.

- Any configurable field can be a required field. When a field is required, a user cannot submit an artifact without providing a value for that field.
- Some configurable fields can be disabled. When a field is disabled, TeamForge does not store data for it or include it in calculations. For example, the **Reported in Release** field is meant for tracking bugs, so you may want to disable it for a tracker that's for user stories. But the **Status** field is important for any tracker, so it can't be disabled.
- For some configurable fields, you can specify whether users see the field when they submit an artifact.
- Some configurable fields need values for users to choose from. You have to provide those choices.
- You can set the width and height of configurable text fields.

### User-defined fields


When none of the TeamForge-provided fields captures exactly the information you need to track, you should create a field tailored to your own needs. Such fields are identified in the field list as "user-defined" fields. See [Create your own tracker fields](#) on page 54 for details.

## What happens when a changed value makes a dependent field invalid?

When a user changes a field value on which the values of other fields depend, some dependent values may need to be corrected.

This can happen if:

- A parent value has been changed in such a way that its child values are inconsistent.
- A text value has been changed in such a way that it no longer matches the regular expression that describes the acceptable values.

 **Note:** Linking fields in dependent relationships does not modify existing data. However, when users later modify fields that are linked, they do have to adhere to the relationships among the fields.

### Changing values

When a user changes the value of a field, and the value in another field becomes invalid as a result, the latter field (we call it the 'dependent' field) is reset to its default value the next time someone updates that artifact.

For example, picture a tracker called 'Interior decoration' in which, if the **Color scheme** field is set to Warm, the **Carpet color** field can be set only to Red, Orange or Yellow. Red is the default value.

User 1 sets **Color scheme** to Warm and **Carpet color** to Yellow.

User 2, after consulting the client, comes in and changes the color scheme to Cool. Now the allowed values of **Carpet color** are Blue, Gray and White, with Blue as the default. But the actual value is still Yellow, which is now invalid. It's up to the user who is updating the artifact to choose a new color from the new set of valid colors. If the user doesn't make a choice, TeamForge automatically changes the value of **Carpet color** to Blue, because that's the default value when **Color scheme** is set to Cool.

### Changing dependency rules

When you restructure a hierarchy of dependent values in a way that leaves invalid child fields in multiple artifacts, each of those artifacts must be fixed the first time a user edits it.

For example, in the artifact `room1`, User 1 has set the **Color scheme** field to Warm and the **Carpet color** field to Red. Suddenly User 2, the project manager, realizes that red is a terrible color for carpeting. User 2 changes the tracker settings so that **Carpet color** can be set only to Orange, Yellow, or Maroon.

The next time User 1 comes back to the `room1` artifact, TeamForge warns that the field value is in an inconsistent state. If the user tries to edit the specific field that is inconsistent, changes to the artifact cannot be saved until the user selects one of the new allowed values for that field.

## Changing text validation rules

If you change the regular expression that governs what content is allowed in a text field, each affected field must be fixed the first time a user edits it.

For example, imagine that your 'Interior decoration' tracker includes a text field in which users can record the telephone number of a carpet installer. You set up a validation rule like `[ ( ) \d \d \d ( ) ] \s \d \d \d [-] \d \d \d \d` to ensure that only standard U.S.-format phone numbers can be entered.

The tracker has been in use for a week when you realize that email would be a better way to communicate with carpet installers. You change the validation rule to something like `\b [A-Z0-9 . _ % + -] + @ [A-Z0-9 . -] + \ . [A-Z] { 2 , 4 } \b`, to ensure that users will enter an email address.

The next time a user tries to update an artifact with a telephone number value in that field, TeamForge warns that the field value is in an inconsistent state. If the user tries to edit the specific field that is inconsistent, changes to the artifact cannot be saved until the user enters a value for that field that matches the new validation rule.

## Deleting a parent field

If you delete a field that contains values that another field's values depend on, the dependent field becomes a standard single-select field on its own. No correction is needed.

## Changing 'requiredness'

- When a field has a parent field that is required, the child field's default value is set to None. If that required parent field is deselected, the child field no longer has to be required, but Required remains the default.
- If you require a specific field value before an artifact can be placed in a given status, that field's children are subject to the same requirements. See [Create a tracker workflow](#) on page 55 for more about controlling what status an artifact can be in.

## Discussions

---

Questions about Discussions in TeamForge.

Concept definition.

### Is the discussion forum creator subscribed by default?

Yes. The user creating the discussion forum is subscribed by default.

For more information about discussion forums, see: [Communicate with project members](#)

### How can I monitor a forum or mailing list?

There are two methods for monitoring a forum or mailing list in TeamForge: at the Discussion Forum level and at the My Workspace level.

At the Discussion Forum level:

- Check the box beside the forum in question.
- Select the **Monitor** button.
- Select **Monitor Selected**.

At the My Workspace level:

- Select **Monitoring**.
- Select the project name.
- Select the **Monitored Applications** tab.
- Select **Discussions**.
- Click **Save**.



**Note:** You must monitor a forum to receive email from the mailing list associated with that forum.

## How do I create a new forum topic?

A forum can have any number of topics. A topic starts a message thread to which other users can reply. Create a new forum topic to begin discussion of a new subject.

To create a forum topic:

1. Click **Discussions** from the **Project Home** menu.
2. On the **Forum Summary** page, click the title of the forum in which you want to create a topic. The **Topic Summary** page is displayed.
3. Click **Create**. The **Create a Topic** page is displayed.
4. Enter a description of the topic in the **Subject** field.
5. Enter the message text in the **Message** field. After the topic is created, other users can reply to this message.
6. (Optional) If you want to add an attachment to the topic, click **Browse** to select the file.
7. Click **Save**. The forum topic is created.

## How do I find the email address for a forum?

The email address for the forum is displayed on the **Topic Summary** page. If a mailing list is not enabled, you will not see an email address on the **Topic Summary** page.

To find the email address for a forum:

1. Click **Discussions** from the Project Home menu.
2. On the **Forum Summary** page, click the title of the forum in which you want to create a forum topic. The **Topic Summary** page is displayed. The email address is listed in the Mailing List field.

## How do I remove a user from a discussion?

Unfortunately there is not currently a way for a TeamForge admin to remove a user from monitoring.

You will need to have the user login and remove themselves from monitoring via the **Monitoring** button on the **My Workspace** page.

## How do I post a reply to a forum topic?

To post a reply to a forum message, follow the steps below.

1. Click **Discussions** from the **Project Home** menu.
2. On the **Forum Summary** page, navigate to the forum topic in which you want to post a message.
3. In the section containing the message you want to respond to, do one of the following:
  - Click **Quote** to quote the original message in your response message. The **Post a Reply** page is displayed with the original message quoted.
  - Click **Reply** to omit the original message from your response message. The **Post a Reply** page is displayed without the original message.
4. Enter the message text.
5. (Optional) If you want to add an attachment to the message, click **Browse** to select the file.
6. Click **Save**.

The forum message is now posted. Other project members can reply to it using the same process.

## I can see the message I posted to a discussion in the web UI, but I didn't receive any of the responses through email. Why is this?

To receive form posts through email, you must monitor the forum.

You can do this by placing a check next to the forum you wish to monitor and choosing 'monitor selected' from the **Monitor** drop-down list at the bottom of the Forum list. For more information on this, see [Monitor many items](#).

## Why are some of the discussions threaded?

Posts that are sent the mailing list address of the discussion will create new topics, and hence will not be threaded.

Posts that are made in response to another post will be threaded beneath the original post.

## What happens when I post to a moderated discussion forum?

A message to a moderated discussion forum is held until a moderator acts on it. (Except if you are a trusted user. These messages do not require moderation.)

- The moderators of the discussion receive an email notification that you have posted a message. The email notification contains the URL path to moderate the post.
- A moderator can either approve or reject your message.
- If the moderator accepts the message, the message status changes to "Accepted" and the message is posted to all the users monitoring the forum.
- If the moderator accepts the message and trusts the sender, the message status changes to "Approved and Trusted" and the message is posted to all the users monitoring the forum. All subsequent posts from you are automatically approved.
- If the moderator rejects the message, he/she can include his/her comments or reasons in the moderation rejection email, and the message is removed from the archive.
- In the **All Topics** tab, an hourglass icon indicates which topics contain posts that await approval.


## Who can be a moderator?

A discussion forum moderator is selected by the forum administrator.

- You can be a moderator by being a project member with forum post permissions.
- A moderator can moderate a discussion forum using emails or the web UI or one of the CollabNet desktops.

If the mailing list isn't enabled for a moderated discussion forum, moderation can be done only through the web UI.

- A moderator can trust the sender and approve a message or reject the message.

 **Note:** To enable or disable moderation or add or remove moderators, you must be either the forum administrator or project administrator.

## Can I move a discussion to another project?

No. Discussions cannot be moved between projects.

For more information see: [Discussions and mailing lists](#)

## Do project owners get automatically subscribed to the discussion forum started by another member?

A project owner does not get subscribed to the discussion forum started by another user.

To make sure the project owner is included, add the user as a member when creating the forum.

## Can I subscribe to a TeamForge discussion forum's mailing list through an email?

Sure, you can subscribe to a discussion forum's mailing list to keep up with changes, through an email. When you subscribe or monitor a discussion forum, you are notified by an automatic email whenever there is any update to the discussion.

**To subscribe to a mailing list:** You can subscribe to the discussion forum's mailing list by sending an email to `<mailing list name>-project name-subscribe@<domain>`.

You must have the 'Discussion: view' permission to be made a subscriber. Check with your forum administrator, if you do not have the permission.

As with other items that you may be monitoring, the discussion forum to which you subscribe through an email is also added to the **My Workspace > Monitoring > Monitored Objects** list. Your user name is added to the 'Users Monitoring' list of the discussion forum. Both these entries are removed when you unsubscribe from the mailing list.

To subscribe in digest type, send an email to `<mailing list name>-project name-digest-subscribe@<domain>`

**To unsubscribe from a mailing list:** You can unsubscribe to the discussion forum's mailing list by sending an email to `<mailing list name>-project name-unsubscribe@<domain>`.

## Why would I want to make a discussion forum moderated?

Automated processes can only do so much to protect forum quality. Moderated posting can be thought of as a temporary transfer of control from automated processes to a human decision maker.

Software can instantly analyze an email to determine whether the sender is allowed to post messages directly to the list or if the message should be sent for moderation. However, the software doesn't recognize the sender as a person who participates in the organization, nor does it have the sophistication needed to determine whether the message contents are on-topic for the mailing list.

Advantages of moderated posting:

- Moderators help keep content on-topic by rejecting off-topic messages and providing helpful suggestions to users whose posts are rejected.
- Moderators help maintain a positive environment for list users by rejecting messages containing harsh or abusive language.
- Moderators can allow deserving nonmembers to post to lists that are ordinarily closed to nonmembers.
- Moderators prevent spam.

### Example

Ashish makes the development forum in his project a moderated one as he wants to make sure that all the messages posted to the discussion come to him for approval before they're included in the forum. When a message arrives, he reads the message, and if it's appropriate for the discussion, he accepts it; if not, he rejects it.

Over time, Ashish finds that the traffic in his discussion has increased and he is no longer able to moderate all the posts by himself. So he adds a couple of other senior developers in his project as moderators, who can share the responsibility of moderating the forum.

After a while, Ashish realises that he doesn't have to reject any messages posted to the forum as everyone seems to understand the purpose of the forum and uses appropriate language in emails. So he removes the restriction and make the forum an unmoderated one. Now Ashish and the other moderators no longer receive emails for approval when a user posts a message to the discussion. Messages are directly included in the forum and delivered to the forum subscribers.

## Documents

---

Questions about Documents in TeamForge.

Concept definition.

## How does TeamForge support documents?

A key element in a successful product development project is ensuring that stakeholders from all functional groups are involved in an iterative document review and approval process, particularly for critical documents, such as Product Requirements Documents.

It is also critical that any changes to key documents are recorded and the new information quickly communicated to everyone concerned.

The TeamForge Document Manager helps you manage your documents throughout their life cycle and ensure the appropriate level of involvement of other project members. In addition to storing documents for reference purposes, the Document Manager provides a document review work flow process to allow you to actively engage other users in the review and approval of a document.

### Simple notification process

If you would like to keep other TeamForge users informed when the status of a document, such as when a document is posted or updated, or there is a change in the contents or status of an existing document, but do not require a formal review and approval process, the following steps provide a best practice for this level of document management:

- Post the document to the Document Manager in the desired location. If you would like to prevent other users from editing, moving, or deleting the document, use the 'Lock Document' feature. (If you choose not to lock the document, any user with the document edit permission can edit and update the document.)
- Notify other users that the document is available and request that they begin monitoring it. All users monitoring an item receive email notification whenever the item is updated.

### Document work flow review and approval

In cases where more formal document review is desired, the TeamForge Document Manager provides an easy-to-use work flow process for managing a document's review and approval.

After submitting a document, you can start a document review. You are prompted to provide the following information:

- The names of all required and optional reviewers. (All reviewers must have permission to access the document.)
- The date by which the review must be completed.
- Email message text.

You will be notified each time a reviewer submits review comments.

When selected to review a document, users receive an email notification with the relevant details and a link to the document.

You can also choose to attach the document to the email notification. An item also appears in the **Documents Awaiting Review** section of each reviewer's **My Page**, including the due date.

A **Submit a Response** section is provided within the Document Manager where reviewers can enter their review comments. Reviewers should be instructed to include "I approve this document" or similar text in the **Submit a Response** section to indicate that they have approved the document.

All reviewers and the document submitter will be able to read the reviews of all other reviewers once they have been submitted.

## Why can't I edit a document when it opens in my browser?

There are chances that your browser swallowed the running application.

If you click on a document in TeamForge (for example, a Word doc), and it opens in your browser instead of launching a separate MS Word window, chances are you will not be able to edit this document. This is because your browser has either called one of the MS document viewers that do not have native edit capabilities, or the browser has 'swallowed' the running MS Word process and the document has written to your %TEMP% directory as a read-only file.

You can either choose to download the file to your desktop and then click on it to edit, or configure your browser to not swallow the application (this is described in various MS KB articles).

## Why are some uploaded documents missing icons when displayed in TeamForge?

TeamForge has a small internal mapping of which icon is associated with which mimetype(s). This mapping is necessarily small, as a complete mapping would be exceedingly large, and almost always out of date.



You can examine the mimetype that TeamForge is receiving by examining the document's attributes using our SOAP API. You can also override this mimetype info via the API by uploading a new version of the document (which can be the exact same document contents). There is sometimes a generic icon for documents from certain applications. TeamForge accepts and stores the mimetype that the user's browser sends on the HTTP upload of the original document. As an example, let's examine the MS Word document. Some browsers will send strings like `application/vnd.ms-word`; others will send `application/msword` and others will send something completely different.

## Can I link to documents outside of TeamForge?

TeamForge has the concept of a 'url doc' to support this very usage.

When creating a document in TeamForge, simply choose this type from the **Create** screen and then enter the URL that references the document in the existing external system. This will create a 'placeholder' document in TeamForge that can be associated to, reviewed, and so on as if it were a normal document, while maintaining the document's actual contents in the external application.

## Can I lock a document in TeamForge?

Absolutely, yes. You can specify a document as locked at any point in time (document create, document edit, etc).

You can do this through the UI or through the SOAP API. Simply check relevant boxes available in the **LOCK DOCUMENT** section on the **Document details** page.

By default, edit and download locks are disabled for a document. To lock and unlock your document, you must have the Document Edit permission in addition to the Document Create permission. To prevent others from editing or downloading a document, you can choose the edit lock or the download lock as required. Selecting the edit lock automatically enables the Document Administrator to edit the document and allows you to select the download lock option.

To edit or download a locked document, the user must have the Document Edit permission. The user who has locked the document can edit or download it at any point in time. A Document Administrator that has the option, **Allow Document Admin to edit** enabled in the **LOCK DOCUMENT** section or a Site Admin can also edit or download a locked document.

## Does TeamForge automatically resolve conflicts in documents made by multiple concurrent editors?

If you store your documents in CVS, Subversion, or some other SCM tool, then that tool will handle conflict resolution per its normal means.

However, if you are using the Document Manager component of TeamForge, then there is no automated means within the product to prevent another user from uploading a newer version of a given document that does not contain the changes you just uploaded. To prevent this, we recommend you use the edit lock and download lock features of TeamForge to prevent others from editing and downloading a document you are currently editing.

## What document types are supported in TeamForge?

TeamForge supports all known and unknown document types.

In fact, TeamForge makes absolutely no distinction on a file type when uploading a document. This means that you can even upload binary files (like a Zip archive) into the document manager. Note that while the TeamForge accepts any incoming data stream as a document, it does use the mimetype sent by the browser to determine if it has the proper icon to display for the document.

## Why doesn't an open review automatically close when a new version of the document is uploaded?

TeamForge does not automatically close a document review under any circumstances (new document version, review due date passes, and so on).

This was a conscious decision on our part in which it was decided that TeamForge cannot always be aware of the business rules or personnel availability at a customer. For example, TeamForge cannot know that the one person whose document review input is most needed is on vacation for the four days the document was under review. If TeamForge were to close the review, then it would disappear from the user's **My Page**. Additionally, TeamForge cannot know that a new version of a document supersedes the prior versions (or that it does not supersede it). You may be uploading a 'draft' or 'work-in-progress' of the new version as document review feedback is received and if TeamForge were to close the review at this point, you might have received feedback from less than one percent of the reviewers.

### Who can work with documents?

Any project member with the Document Edit permission can edit a document.

### Can I set permissions so that users can move documents but not delete them?

You cannot configure document management permissions so that a user can move documents but not delete them.

It is not possible to separate move and delete permissions, because a move is actually a copy/delete action. The document is not really moved, it is copied to the new location and then deleted from the original location.

### Why am I not able to access a folder in TeamForge documents?

You can access a folder in TeamForge documents only if you have permission for the folder or any of its parent folders.

For TeamForge documents, users having permissions to the parent folder can access all the subfolders. Users having permissions to a subfolder can only view the hierarchy of the parent folders; they cannot access the documents in those folders.

### Are role-based permissions allowed for subfolders in the TeamForge Documents?

Yes. It is allowed in TeamForge Documents.

The Project Administrator can give permission to access sub folders in the TeamForge Documents based on the user roles, using the Roles option.

### Can I set permissions for documents alone?

Yes, as a project administrator, you can set individual permissions only for documents.

The Document Admin permission allows a project member to create, edit and administer both documents and document folders. Though this does not include the delete permission, individual permissions to create, edit, and delete can be set separately for document folders and documents on the **Documents Permissions** page (**Project Admin > Permissions > Edit Role > Documents**). This is to restrict a project member from handling document folders or subfolders. For example, if you do not want a project member to create a document folder, but only create a document within a specific folder, you can set the 'Create/View' permission for documents alone. The available permissions are:

- Create/View
- Edit/View
- Delete/View
- View Only

## Wiki

---

Questions about Wiki in TeamForge.

Concept definition.

### How do I generate a wiki table of contents?

You can create a table of contents from any heading text that you have in your wiki page.

For versions 5.2 and earlier, generating a wiki table of contents requires the Wiki TOC plugin, available through CollabNet Professional Services.

To enable TOC for a Wiki page, place the following in your Wiki page at the spot where you want the Table of Contents to appear.

```
%%insert-toc
%%
```

The Table of Contents is generated automatically based on the heading markers in the wiki page, e.g. `!!!Heading`

## How do I delete an attachment from a wiki page?

You can delete an attachment from the wiki page.

Use the following steps to delete an attachment:

- Browse to the wiki page that contains the attachment and click **Edit**.
- Update the wiki page, even though you have not made any changes to it.
- Click **Show Details**.
- Click the **Attachments** tab.
- Select the attachment you want removed and click **Remove**.

## How can I detect orphan wiki pages?

No. You cannot detect orphan wiki pages.

Unfortunately at this time, TeamForge does not have the ability to locate or display orphaned wiki pages within the UI. This functionality is slated to be included in a future release. If you have an immediate need, please contact Technical Support and a suitable SQL query will be devised.

## How do I edit the wiki home page?

Navigate to the project home page and click the **Edit** button to edit the wiki home page.

Once you have enabled the wiki as the project home page option in Project Admin, you must return to the project home page by clicking the **Project Home** button. There will be an **Edit** button on the bottom right corner of the page under graphs. You may need to scroll your browser window to the right to see this button.

## How do I make the version comment required for wiki updates?

You can use Velocity to customize the pages of TeamForge and make the version comment required.

Create `/opt/collabnet/teamforge/sourceforge_home/templates/body_footer.vm` with the following contents:

```
<script>
var updateButton = document.getElementById('edit_wiki_page.update');
if ( updateButton ) {
updateButton.href = 'javascript:submitWikiPageUpdate();';
}
function submitWikiPageUpdate () {
if ( document.editPage.versionComment.value ) {
submitForm(document.editPage, 'submit');
return;
}
}
alert("Please include a detailed Version Comment for this change.");
}
</script>
```

Ensure that the file is owned by the sf-admin user:

```
chown -R sf-admin.sf-admin /opt/collabnet/teamforge/sourceforge_home/
templates
```

Ensure proper permissions:

```
chmod 0644 /opt/collabnet/teamforge/sourceforge_home/templates/  
body_footer.vm
```

TeamForge picks up the change immediately.

## I set my project to 'use wiki homepage' so why isn't my wiki showing up?

TeamForge currently uses two distinctly different wikis.

- If there is the wiki you have already edited, which is available by clicking on the **Wiki** button at the top of any project page, and there is the 'project home' wiki, which is what you enabled in the **Project Admin** screen.
- If you visit the project home page after setting this option, and if you have the proper RBAC permissions, there should be an **Edit** button in the bottom-right corner of the home page under the graphs. Use this button to edit the project home page wiki.

## Why would I attach things to a wiki?

There are specific scenarios that would require attachments to a wiki.

The most common example of when you would attach something to a wiki would be when you need an image in the wiki page and you are concerned that if the image is hosted remotely (a corporate web server, for example), it might be moved or removed. Additionally, you might wish to attach a file to a wiki page if the attachment is only truly important in the context of the wiki page and therefore is not important enough to be uploaded to the Document Manager of TeamForge.

# TeamForge 17.4 release notes

---

Released on: 8 May, 2017

## TeamForge 17.4 version history

- TeamForge 17.4 GA: 17.4.770-27 (05/08/2017)

## New features in TeamForge 17.4

---

TeamForge 17.4 has a lot of new features and enhancements. Here's a list of a few release-defining new features in TeamForge 17.4.

### TeamForge OAuth 2.0 authorization framework

In the traditional client-server authentication model, client applications must have access to the Resource Owner's credentials in order to request and access protected resources on the Resource Server. Sharing Resource Owner's credentials such as usernames and passwords to third party applications poses several problems, security being one of the major ones. Frameworks such as the OAuth 2.0 and OpenID Connect help in mitigating such security risks and hence TeamForge's authentication model has been completely rebuilt on top of OpenID Connect and OAuth 2.0 Authorization frameworks.

### Federated Identity Management

Federated Identity Management refers to the process of storing user credentials with an IdP (such as TeamForge, SAML and so on) and having the Service Provider (SP) rely on the IdP to validate user credentials when a user tries to access its resources.

With the new TeamForge Identity Management built on OpenID Connect (OIDC) and OAuth 2.0 authorization frameworks, TeamForge can now act as an ID Provider (IdP). As an IdP, TeamForge can authorize a third-party client application to obtain limited access to its services either on behalf of a Resource Owner (user) or on behalf of the client application itself.

TeamForge supports federated identity. By default, federated login is disabled in TeamForge. Site Administrators must enable federated login and select an IdP.

Both TeamForge REST and SOAP APIs use OAuth 2.0 access tokens for authentication. Clients can obtain access tokens from the token endpoint which is located at `/oauth/auth/token`.

For more information, see [Identity Management: Set up TeamForge OAuth 2.0 Authorization Framework](#) on page 314.

## Integrations

### TeamForge-JIRA integration

This release brings a new feature for JIRA site administrators to disable active JIRA-TeamForge project mappings in bulk (**Administration > Add-ons > TeamForge Associations**). The intended use case revolves around staging JIRA upgrades. When a JIRA server is mirrored into a staging/testing environment, any changes to issues in projects that are mapped to TeamForge will trigger events back to TeamForge, thereby polluting the production TeamForge event data store. Use this feature to disable all production JIRA-TeamForge mappings in stage environments and then test the TeamForge Associations add-on by creating a new mapping between a staging JIRA server and TeamForge server.

Applications Projects Issues **Add-ons** User management System Configuration Manager

ATLASSIAN MARKETPLACE  
Find new add-ons  
Manage add-ons  
Purchased add-ons

SOURCE CONTROL  
TFS4JIRA Configuration

TEAMFORGE ASSOCIATIONS  
**Manage**

### TeamForge Associations

Use this function to remove TeamForge project mappings for selected projects below. Existing sources, events and associations data will be preserved in TeamForge EventQ. The primary use case for this feature is in staging environments to prevent contamination of production data while testing.

Select the below TeamForge Mapping to remove from project:

- Select All
- 31Oct
- Jira7-Teamforge Integration
- JiraintA
- List
- New
- NewProject

Remove

For more information, see [Remove TeamForge project mappings](#) on page 641.

## Nexus integration

[CTF-Nexus-Integration-Plugin-2.1.1](#) is now available. To upgrade, see [Upgrade the TeamForge-Nexus integration plugin](#) on page 671.

## TeamCity integration

TeamForge EventQ now supports TeamCity versions 9.0 and 10.0.

## My Workspace

You can create one or more dashboards to categorize and view data in many different views. For instance, you may decide to create three dashboards: Dashboard A, B and C. With these three dashboards, you may choose to configure Dashboard A and B to view events from projects A and B respectively and configure Dashboard C to view just the code reviews. For more information, see [My Workspace](#) on page 13.

My Workspace My Page Jump to ID Ex: artf2345

Dashboard C

My Recent Projects

- RTFM
- TeamForge
- Helpdesk
- Customer Issues

My Recent Repositories

- output
- en

My Recent History

- cmmt370188: Commit by Meenakshi Kamaraj (kmeenakshi)
- cmmt370187: Commit by Meenakshi Kamaraj (kmeenakshi)
- cmmt370184: Commit by Meenakshi Kamaraj (kmeenakshi)
- cmmt370183: Commit by Meenakshi Kamaraj (kmeenakshi)
- cmmt370132: Commit by Meenakshi Kamaraj (kmeenakshi)

My Artifacts

ID	TITLE	STATUS	CREATED BY	PROJECT
artf289835	Doc task for artf236905 and artf289686	Open	Arun Panduran...	TeamForge
artf289235	Request to document tracker artifact limit in HCN	Open	Prem Kumar P...	TeamForge
artf288732	Docs: Add step to remove any add-ons prior to upgrade	Open	Richie Jarvis	TeamForge

New My Workspace Dashboard

Your My Workspace dashboard can be customized to suit your own needs. Click the customize icon (🔧) at the top of the page to get started.

The dashboard consists of three main things:

1. A name, which you can update via the text box at the top of the page,
2. A layout, which you can choose from a number of page layouts, which give you options in how to arrange your dashboard.
3. Widgets, which provide the functionality you want to expose on your dashboard. To add a widget just click the **Add widget** button. This will add the widget(s) to the top of your dashboard. A variety of widgets have been provided, including this one which is a Markdown widget that you can use for adding links and images to your dashboard. Most widgets have criteria or content that controls what information they display. These settings can be found by clicking the widget configuration icon (⚙️) on the widget.

Arrange widgets by clicking on the widget location icon (➕) and dragging the widget to the place in your layout that you want to see it.

When you are satisfied with your changes click the **Save** button.

There is a **Cancel** button if you do not want to keep your changes, and the **Use default** button will return the dashboard to the default settings.

Feel free to delete or edit the content of this widget once you understand the basic concepts.

## Configure Application tool

The **Configure Application** tool, added in TeamForge 17.4, makes TeamForge site level settings configurable via the user interface. This comes in handy for site administrators, who otherwise would have to work with the `site-options.conf` file and then recreate the TeamForge runtime for site level configuration changes. For more information, see [Configure your site's settings](#).

The Configure Application tool is not loaded with all the available site configuration settings available in TeamForge. For settings that are not available in the **Configure Application** tool, Site Administrators must go the `site-options.conf` route as always.

## Tags

The tagging functionality implemented for documents in TeamForge 17.1 has been extended to support tracker artifacts as well in TeamForge 17.4. Project members with CREATE/EDIT permissions can create new tags or add existing tags to tracker artifacts from the **Create Artifact** and **Edit Artifact** pages. You can add up to 10 tags to an artifact. You can add or remove tags to tracker artifacts when you search for artifacts within a single tracker or across trackers or during tracker mass update. You can also include "Tags" as one of the columns in the **List Artifacts** page.

For more information, see [Create a tracker artifact](#) on page 39 and [Edit a tracker artifact](#) pages.

The screenshot shows the 'Create Artifact' page in TeamForge. The page has a dark header with 'TeamForge' and 'Trackers' menus, and a search bar containing 'tracker22419'. The main content area is titled 'Create Artifact' and contains the following fields:

- TITLE\***: A text input field.
- DESCRIPTION\***: A large text area.
- TAGS**: A field with a search icon and an 'Add tag' button.
- GROUP**: A dropdown menu with 'None' selected.
- STATUS\***: A dropdown menu with 'Open' selected.
- CATEGORY**: A dropdown menu with 'None' selected.
- CUSTOMER**: A dropdown menu with 'None' selected.
- PRIORITY**: A dropdown menu with '4 - Low' selected.
- TEAM**: A dropdown menu with 'None' and a user icon.
- ATTACHMENT**: A section with an 'ATTACH ANOTHER FILE' button.
- TAGS**: A field with a search icon.
- CODE REVIEW**: A field with a search icon.

## Tracker workflow enhancements

TeamForge 17.4 introduces Advanced Transition settings, which you can use to make specific fields as 'required' and also to show or hide specific fields when a status transition happens. Advanced Transition settings are applied when you create or edit artifacts in a tracker and only when you edit artifacts in Planning, Task, and Kanban boards. For more information, see [Create a tracker workflow](#) on page 55.

**Advanced Transition**
✕

Transition Status To Under Consideration -> Completed

NAME	TYPE	REQUIRED	VISIBLE
Group	Single Select	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Status	Single Select	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Category	Single Select	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Customer	Single Select	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Priority	Single Select	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Team	Folder	<input type="checkbox"/>	<input checked="" type="checkbox"/>

Cancel
Save And Close
Save

## SCM/Git/Gerrit

### Repository Attachment Folder

You can now select a folder to upload code review and release notes file attachments. All the files attached to the repository are stored in the selected folder. For more information, see [Create a source code repository](#) on page 93.

### Highlights of TeamForge-Git integration: 17.4.6-2.13.8

- Based on vanilla Gerrit v2.13.8. See [Gerrit v2.13.8 release notes](#) for more information.
- Git installer enhancements to support TeamForge 17.4 installer enhancements.
- User interface performance enhancements to better support branches and Gerrit reviews that are substantially large in size.
- **GERRIT\_FORCE\_HISTORY\_PROTECTION** site-options token is no longer supported. Remove this token from the `site-options.conf` file while upgrading to TeamForge 17.4 or later. History Protection is enabled by default in TeamForge 17.4 and later. However, Site Administrators can disable History Protection at the site level if need be, after which Project Administrators can choose to have History Protection enabled or disabled for individual repositories.

### Highlights of TeamForge-Git integration: 17.4.11-2.13.9

- Based on 2.13.9 vanilla Gerrit version released on 2 August, 2017. See [v2.13.9](#) for more information.
- Fixed the issue in which the RabbitMQ connections and threads are kept permanently in memory even if the project is not active currently.

#### Cause of the issue:

This issue is caused by the default configuration of RabbitMQ client which opens the dedicated thread pool (that spawns up to 2 \* number of CPU threads by default) per connection. The issue manifests with large number of threads being spawned over time up to,

- (number of repositories + number of projects) for threads named **AMQP Connection AA.BB.CC.DD.YYYY**.
- (number of repositories + number of projects) \* (2 \* number of CPUs) for threads named **pool-XXX-thread-Y** (especially when the machine is heavily loaded and multiple changes are being reviewed and pushed to multiple projects).

#### Solution:



- **Shared thread pool used by all RabbitMQ connections:** It can spawn up to (factor \* number of CPUs) threads where factor can be modified by setting `teamforge.notificationMaxRabbitMQExecutorsFactor` in `gerrit.config` (default value is 2) file. Threads are removed from thread pool after being idle for more than `teamforge.notificationMaxRabbitMqThreadIdleTime` seconds in `gerrit.config` (by default 3\*60s) file. This addresses the issue with the **pool-XXX-thread-Y** threads explosion.
- **Connection Expiration:** Connection to get closed after being idle for more than `teamforge.notificationMaxRabbitMQConnectionIdleTime` seconds in `gerrit.config` (default is 5\*60s) file. This addresses the issue with the **AMQP Connection AA.BB.CC.DD:YYYY** threads explosion as they are not being quickly (marked as "not used" after `teamforge.notificationMaxRabbitMQConnectionIdleTime` idle time and guaranteed to be collected within (2 \* `teamforge.notificationMaxRabbitMQConnectionIdleTime` idle time)) collected.
- You can now search for the changes using Change-Id triplet (Repository~Ref~Change-Id). When the first 5 characters of the Repository name (Example: IdeaField) matches the Change-Id pattern, the changes searched for will be shown. This fix is sent back to open source Gerrit.

### Highlights of TeamForge-Git integration: 17.4.14-2.13.9

- For the vanilla based Gerrit release notes and the bug fixes, see Gerrit [v2.13.9-11-gf25da13683](#) released on November 3, 2017.
- HTTP webhooks functionality has been implemented to set up communication between TeamForge Git integration and the following products:
  - Lifecycle
  - Continuum
  - DLM 2.0
  - Any other tool that accepts HTTP notifications

TeamForge UI for webhooks (implemented in TeamForge 17.11) is not backported. However, you can use the following products to configure webhooks:

- CollabNet Desktop (latest development version)
- GitEye (latest development version)
- TeamForge WebHooks Configurator
- *Bug fixes*

Critical upstream (stable-2.13) Gerrit fix is now available for SSH clone/fetch failure when server takes more than 30 seconds to compute its operation. To handle this, a new parameter `sshd.waitTimeout` has been included in the `gerrit.config` file. By default, the value of this parameter is 30 seconds.

## EventQ

### EventQ Installation/Upgrade Changes

Starting from TeamForge17.4, you do not have to install EventQ separately on top of TeamForge. As long as you have your site's master configuration file (`site-options.conf`) specifying where your EventQ services (`eventq`, `mongodb`, `redis` and `rabbitmq`) are hosted, the TeamForge installer, by default, installs or upgrades EventQ services as part of TeamForge installation/provision.

In case you install or upgrade EventQ on a separate server, you must configure your TeamForge installation repository on the EventQ server and install EventQ using the `teamforge-eventq` meta RPM (`yum install teamforge-eventq`).

### Obsolete `site-options.conf` tokens

As a result of EventQ installer changes, the following EventQ tokens are no longer required/supported (see [Site options change log](#)).

- `ORC_HOSTNAME`
- `ORC_SSL_CA_CERT_FILE`
- `ORC_PORT`

- ORC\_PROXIED\_PATH
- ORC\_PROTOCOL
- ORCHESTRATE\_ENABLED
- POSTGRES\_INTERFACE\_IP
- POSTGRES\_INTERFACE

Remove these tokens from the `site-options.conf` file while upgrading to TeamForge17.4 or later.

For more information, see TeamForge17.4 install/upgrade instructions.

### Datamart performance enhancement for flex fields

A new XML to non-XML conversion based solution has been implemented to redefine the flex field data storage and retrieval mechanism. This new approach improves the performance of reporting queries used for retrieving flex field data from the Datamart. For more information, see [Why do I get performance issues when retrieving flex field information from Datamart?](#) on page 1040

## Install / Upgrade

### Changes to supported software versions

- Postgres 9.6.1
- James 2.3.2.1

### Default values for site-options.conf tokens

Default values have been assigned to the following `site-options.conf` tokens and are therefore removed from the default TeamForge 17.4 `site-otptions-default.conf` file.



**Warning:** If you are upgrading from TeamForge 17.1 (or earlier) to TeamForge 17.4 (or later) and if you have been using your own values for the following tokens, you must make sure you use the same values in your `site-options.conf` file post upgrade to TeamForge 17.4.

```
# CTF Core
SCM_ADMIN_PASSWORD=$auto$
IAF_DBPASS=$auto$
SCM_DEFAULT_SHARED_SECRET=$auto$
SOAP_ANONYMOUS_SHARED_SECRET=$auto$

# Database
DATABASE_USERNAME=teamforge
DATABASE_NAME=teamforge
DATABASE_PASSWORD=$auto$
DATABASE_READ_ONLY_USER=teamforge_reader
DATABASE_READ_ONLY_PASSWORD=$auto$

# Datamart
REPORTS_DATABASE_USERNAME=teamforge_datamart
REPORTS_DATABASE_NAME=teamforge_datamart
REPORTS_DATABASE_PASSWORD=$auto$
REPORTS_DATABASE_READ_ONLY_USER=teamforge_datamart_reader
REPORTS_DATABASE_READ_ONLY_PASSWORD=$auto$

# ETL
ETL_SOAP_SHARED_SECRET=$auto$

# Gerrit
GERRIT_DATABASE_PASSWORD=$auto$

# RabbitMQ
RABBITMQ_APP_ADMIN_USER=guest
RABBITMQ_APP_ADMIN_PASSWORD=$auto$
```

```

RABBITMQ_APP_CTF_USER=ctf
RABBITMQ_APP_CTF_PASSWORD=$auto$
RABBITMQ_APP_SERVICES_USER=eventq
RABBITMQ_APP_SERVICES_PASSWORD=$auto$
RABBITMQ_PERMISSION_PUBLISHER=permission_publisher
RABBITMQ_PERMISSION_PUBLISHER_PASSWORD=$auto$

# Mongo
MONGODB_APP_DATABASE_NAME=eventq
MONGODB_ADMIN_DATABASE_NAME=admin
MONGODB_APP_ADMIN_USER=admin
MONGODB_APP_ADMIN_PASSWORD=$auto$
MONGODB_APP_BACKUP_USER=backup
MONGODB_APP_BACKUP_PASSWORD=$auto$
MONGODB_APP_SERVICES_USER=eventq
MONGODB_APP_SERVICES_PASSWORD=$auto$

# HAProxy
HAPROXY_STATS_PASSWORD=$auto$

# Binary
IAF_DBNAME=iafdb
IAF_DBUSER=iafdbusr

# James Admin (for management interface)
JAMES_ADMIN_USER=admin
JAMES_ADMIN_PASSWORD=$auto$

```

## LISTEN\_IP

By default, services bind to the IP address corresponding to the `<host>:PUBLIC_FQDN` token. However, you can override this using the `<host>:<service>:LISTEN_IP` site-options.conf token. You can use this `<host>:<service>:LISTEN_IP` token to control which IPs the services bind to so that you can make sure that services are not overexposed than necessary.

For example:

```

mypostgresserver:LISTEN_IP = 1.2.3.4
localhost:mail:LISTEN_IP = 1.2.3.4

```

## More changes

- The `is_service_running.py` script is no longer available.
- TeamForge 17.4 installer expects the system locale to be `LANG=en_US.UTF-8`. TeamForge create runtime (`/opt/collabnet/teamforge/bin/teamforge provision`) fails otherwise.
- The Review Board application, which was used to be installed only on the TeamForge Application Server, can now be installed on a separate server.
- The default Review Board data directory has been changed from `/opt/collabnet/reaviewboard/data` to `/opt/collabnet/teamforge/var/reviewboard/data` in TeamForge 17.4. Make sure you restore your Review Board backup to the new data directory, which is `/opt/collabnet/teamforge/var/reviewboard/data`.
- The `publish/remotePublishLog` is now being generated in `/opt/collabnet/teamforge/log-archive/others/` as expected.
- As PostgreSQL upgrade is taken care of by the `teamforge` script during `teamforge` provision, `pgturant` utility is no longer required and therefore not supported any longer.
- CEE compatibility mode is no longer supported.
- Warning messages displayed during the installation or upgrade are now changed from "Yellow" to "Amber" color.
- For a list of obsolete and new site option tokens in TeamForge 17.4, see [Site options change log](#).

## snapshot.py

The `snapshot.py` script has been deprecated. The `teamforge` script subsumes the `snapshot.py` script's functions. For more information, see [TeamForge script](#) on page 785.

## Issues resolved in TeamForge 17.4

---

Here's a list of few noteworthy issues fixed in TeamForge 17.4.

- **Document Locked** and **Download Restricted** icons are now shown in the **List Documents** page.
- Document folder list in the **List Documents** page is displayed completely now.
- Once the session times out, the **List Documents** page redirects users to the TeamForge Login page as expected.
- You can now access documents and document folders in a "Public" project as an "anonymous" user.
- Fixed the issue with the footer bar not displaying properly in **List Documents** page when filters are applied.
- Fixed the issue in which the Flex fields become hidden after the user inputs invalid value in the fields and saves the document in the **Edit Document** page.
- While importing artifacts, a date format mismatch issue that prevented the import has been fixed.
- Fixed an issue that caused a TeamForge system error when a user with view only permission removed an association in code browser.
- Non-English content, if any, in the artifact description, is now wrapped properly to the next line.
- Fixed the bread crumbs display issue in Tags Management page.
- Fixed the user-defined date field's display issue in Tracker List page.
- Users with view only permission can now search documents using **Tags** as a filter.
- You can now export the list of Reported Tracker Artifacts, Fixed Tracker Artifacts or Planned Tracker Artifacts (with "team" data included) from the **File Releases** page without any issues.

### Bug fixes in TeamForge-Git integration: 17.4.6-2.13.8

- Fixed an issue to strictly enforce the "Association Required on Commit" policy. It is no longer possible to commit without the required association details.
- Fixed an iFrame issue. Clicking the **Browse** link in Gerrit no longer opens the TeamForge code browser on another iFrame.
- Fixed the Code Browser links in email to resolve a URL redirect issue in cases where the destination ref contains a "/" character.
- Fixed the incorrect Git repository creation time issue.

## Known issues in TeamForge 17.4

---

The following noteworthy issues, including any workarounds we may have, are known to exist in the TeamForge 17.4 release.

### Install / Upgrade

- During TeamForge provision, the Register SCM integration process fails on sites that use self-signed certificates. To work around this issue:
  1. Restart JBoss on the TeamForge Application server.
 

```
/opt/collabnet/teamforge/bin/teamforge restart -s jboss
```
  2. Reinitialize TeamForge on the SCM and Git server.
 

```
/opt/collabnet/teamforge/bin/teamforge reinitialize
```
- For successful creation of new review requests in Review Board, perform the following steps:
  1. Go to **Admin > INTEGRATIONS > SCM INTEGRATIONS**.
  2. Edit 'SVN' integration.
  3. Remove ':443' from REPOSITORY BASE URL and click **Save**.

### Display on Create check box

After marking a field as "Required" for a tracker/document, it is possible to clear the **Display on Create** check box of the field that you marked as **Required**.

### UserResourceCache is not configurable

While it is possible to configure the user resource cache size manually via the user interface, an issue prevents the custom user resource cache size from being applied.

### TeamForge System Error

When a user logs on after user session and OAuth token expiry, a TeamForge system error is being shown for the page that the user is trying to access right after the log on.

### Email notifications

Email notifications are not sent to the entire mailing list if it has an invalid email ID in it.

### Multiselect widget

If you have two multiselect drop-down lists, the first multiselect drop-down list fails to close once you select the second multiselect drop-down list.

### Required Fields

Even if you have set up a status transition rule with more than one "Required" field, the error message (see below) refers to only one of the fields as being non-compliant with the workflow transition rule.

```
Warning - the following artifacts were not updated:
<artifact id>: This change is not compliant with the existing workflow
transition rule.
Setting in "Category" field violates status transition constraints. Category
is required"
```

# CollabNet TeamForge 17.4 Update 1 release notes

---

Released on: August 14, 2017

## TeamForge 17.4 version history

- TeamForge 17.4 GA: 17.4.770-27 (05/08/2017)
- TeamForge 17.4 Update 1: 17.4.798-66 (08/14/2017)

## About TeamForge 17.4 Update 1

### Fixed issues

- The connection issue between TeamForge and EventQ (RabbitMQ) that blocked updates from being sent to the Activity Stream has been fixed.
- The Planning Folder Burndown chart shows the correct Velocity Trend line for the artifacts with P1 priority and "ActualEfforts" values.
- The *add\_eventq\_to\_projects.py* script used to add eventq to projects works fine.
- The users are now added to the list of monitoring users in the **Wiki** page as expected.
- A proper success message is now being displayed when you click the **Monitor** and **Stop Monitoring** buttons from the **View Artifact** page.
- Fixed an issue with permission checks when you add users to monitor an object. Both the TeamForge Web application and SOAP APIs now function the same way.
- Files downloaded from the FRS in Firefox/Safari Browsers are compressed properly.
- The RBAC issues found in the monitoring REST API leading to inappropriate user activities are fixed now.
- The SCM Viewer URL does not include a port number as expected.
- The *mod\_authnz\_ctf* authentication and authorization module now supports third party authentication, and handles load balanced/proxy setups gracefully.
- More time taken during project creation is fixed.
- Upgrade to *Git/Gerrit 17.4.11-2.13.9* to open review requests from the *Reviews* tab without any issues and to get rid of the thread leakage issue from RabbitMQ.
- The changes are updated in *site-options.conf* file as expected after modifying `MAX_WWW_CLIENTS` option.
- The comments section is displayed properly on the **View Artifact** page for the MST (GMT-7:00) time zone.
- You can now create a bulleted list with proper indentation as expected in the artifact description while creating an artifact.
- Fixed the issue in which the user even with sufficient permission is unable to create the "Artifact Open/Close Chart (Multiple Trackers)" report.
- Fixed the issue in which an empty "gitroot" folder is created while provisioning during the upgrade.
- When the tokens *SYSTEM\_EMAIL* and *ADMIN\_EMAIL* are updated in *site-options.conf* file, the changes are not taking effect in *runtime-options.conf* and the respective property files. This is fixed.
- Users were not able to log on to TeamForge after password reset typically on sites with LDAP and CVS integrations. This is fixed.

## Apply 17.4 Update

Click [here](#) to apply the latest 17.4 update.

# CollabNet TeamForge 17.4 Update 2 release notes

---

Released on: January 12, 2018

## TeamForge 17.4 version history

- TeamForge 17.4 GA: 17.4.770-27 (05/08/2017)
- TeamForge 17.4 Update 1: 17.4.798-66 (08/14/2017)
- TeamForge 17.4 Update 2: 17.4.824-110 (01/12/2018)

## About TeamForge 17.4 Update 2

### Fixed issues

- An error was thrown while creating planning folders that included the keywords such as tracker, docman, planning, frs, taskmgr, scm, wiki, scm, and discussion as part of their names (Example: "Metric tracker"). As a result, the planning folder id was not shown in the **Jump to ID** field as well. This issue is fixed now.
- During an upgrade from TeamForge 16.x, images were not displayed for custom project web pages as only the relative path of the image files were shown. This issue is fixed and now the images are displayed with their relative path as well.
- While sorting documents by document id or document name, the document icons were not sorted properly along with the documents. This issue is fixed now.
- After clicking any link in the custom project home page, a teamforge system error was thrown even though the linked page showed up. This issue is fixed.

## Apply 17.4 Update

Click [here](#) to apply the latest 17.4 update.