



HelixSwarm

Helix Swarm Guide

2022.2
August 2022

PERFORCE

www.perforce.com



Copyright © 2013-2022 Perforce Software, Inc..

All rights reserved.

All software and documentation of Perforce Software, Inc. is available from www.perforce.com. You can download and use Perforce programs, but you can not sell or redistribute them. You can download, print, copy, edit, and redistribute the documentation, but you can not sell it, or sell any documentation derived from it. You can not modify or attempt to reverse engineer the programs.

This product is subject to U.S. export control laws and regulations including, but not limited to, the U.S. Export Administration Regulations, the International Traffic in Arms Regulation requirements, and all applicable end-use, end-user and destination restrictions. Licensee shall not permit, directly or indirectly, use of any Perforce technology in or by any U.S. embargoed country or otherwise in violation of any U.S. export control laws and regulations.

Perforce programs and documents are available from our Web site as is. No warranty or support is provided. Warranties and support, along with higher capacity servers, are sold by Perforce.

Perforce assumes no responsibility or liability for any errors or inaccuracies that might appear in this book. By downloading and using our programs and documents you agree to these terms.

Perforce and Inter-File Branching are trademarks of Perforce.

All other brands or product names are trademarks or registered trademarks of their respective companies or organizations.

Any additional software included within Perforce is listed in "[License statements](#)" on page 1397.

Contents

How to use this guide	19
Syntax conventions	19
Feedback	20
Other documentation	20
Earlier versions of this guide	20
Helix Swarm User Workflow Guide	20
1 About Helix Swarm	21
Swarm videos	21
Related topics	21
What's new	23
Major new functionality	23
Added support for Helix Authentication Service as an SSO provider	23
Minor new functionality	23
Improved performance of the Swarm review page	23
Improved performance of My Dashboard	23
Added support for Helix Core Server 2023.2	23
Added support for P4PHP 2023.2	23
Important information	24
Upcoming changes to SSO configuration	24
PHP 7.X versions will no longer be supported in Swarm 2024 release	24
Project level test and deploy code features to be deprecated in a later Swarm release	24
Swarm 2022.3 or later only works with Helix Core Visual Client (P4V) 2021.3 or later	24
The emulate_ip_protections configurable now defaults to false	24
API version updated to v11 for Swarm 2022.1	25
Swarm no longer supports Microsoft Internet Explorer	25
Swarm no longer supports CentOS 8	25
Swarm no longer supports Ubuntu 16.04	25
Swarm no longer supports CentOS 6 and RHEL 6	25
Upgrading from Swarm 2019.1 and earlier	25
Known limitations	26
After an upgrade to Swarm 2022.3, you may see partial text labels or unexpected page artifacts	26
No activity event is created when a project setting is updated	26
Issue with userids containing a / character	26

Swarm shelvedel trigger can fail for a specific Windows use case	26
Unsupported characters in user names and group names in Swarm 2019.2 and later	26
Swarm support for the "Private editing of streams" feature in Helix Core Server 2019.1 and later	27
Multiple Helix Core Server instances on a single Swarm instance	27
Global Dashboard does not support Single Sign-On (Helix Authentication Service)	27
Project Commits tab can fail to show some Helix Core Server commits in the top level view	27
Task Stream Reviews	28
2 Quickstart	29
Swarm quick reference	29
Menu quick reference	29
Quick URLs quick reference	34
My Dashboard page quick reference	36
Global dashboard page quick reference	38
Reviews list page quick reference	43
Review page quick reference	49
Swarm user tasks	69
Start a code review	69
Edit the reviewers on a review	71
Check for code reviews you need to act on	72
List reviews you are involved with	74
Contribute comments or code changes to a code review	74
Get a local copy of the code in a review for evaluation	77
Fix errors that cannot be merged in a review	79
Change the author of a review	80
Change your user notification settings	81
Unfollow all projects and users you are following	81
Swarm administration tasks	83
Change notification settings for a group	83
Manage project branches	84
Unfollow all projects and users for another user	90
Obliterate a review	91
Change the logging level	92
Check the queue workers	92
Automatically deploy code within a review	94
3 Install and upgrade Swarm	96

Important restrictions	96
Review the runtime dependencies	96
Swarm and Helix Core Server installation considerations	96
Choose the installation process	97
Upgrade Swarm	98
Runtime dependencies	98
Recommended operating systems	99
Apache web server	100
PHP	102
Helix Core Server requirements	106
Helix Core Server event notification	107
Worker dependencies	109
Redis server	110
Supported Helix Core Visual Client (P4V)	110
Supported Helix Authentication Service (HAS)	110
Supported web browsers	111
Workflow prerequisites	111
Optional dependencies	111
Security-enhanced Linux (SELinux)	112
Choose the installation process	113
Install and configure Swarm from a package (recommended)	115
Install and configure Swarm on Ubuntu	115
Install and configure Swarm on CentOS/RHEL	128
Install and configure Swarm on Amazon Linux 2	152
Run Swarm using a Docker container	169
Prerequisites	170
Helix Swarm images	170
Running the Docker container	171
Persisting Docker containers for a production environment	174
Restarting the Docker container	174
Migrating Swarm to a Docker container	176
Advanced configuration options	177
Example of running Swarm using docker-compose	179
Install and configure Swarm manually from a Tarball	181
Redis configuration	182
Apache configuration	186

PHP configuration	190
Swarm configuration	193
Configuring Helix Core Server event notification	197
Set up a recurring task to spawn workers	220
Common post-installation options	224
Post-install configuration options	224
Validate your Swarm installation	239
Review not created	240
Review cannot be updated when using triggers	241
Upgrading Swarm	241
Upgrading a Ubuntu package installation	242
Upgrading a CentOS/RHEL package installation	254
Upgrading a Docker container	274
Upgrading a tarball installation	275
Moving your Swarm instance	297
Summary	297
Installing Swarm on your new server	297
Copying your existing Swarm configuration	298
Configuring Helix Core Server event notification	298
Replacing your original Swarm instance with your new instance	303
Validating your new Swarm instance	305
4 Basics	306
My Dashboard	306
Filtering	307
Review summaries	308
Global Dashboard	310
Toolbar	312
Sidebar	314
Helix Core Server dashboards	317
Navigating directly to a specific Helix Core Server instance in Swarm	320
Activity	321
Files	322
Browsing deleted files and folders	323
File display	324
File edit	329
Download files as a ZIP archive	330

Supported syntax highlighting in the Review page	331
Commits	332
Range filter	333
File Commits	334
Remote depot commits	334
Jobs	335
Edit the Job columns	335
Job display	336
Add jobs	337
Unlinking jobs	338
Changelists	338
Changelist display	339
Diffs	341
Viewing a diff	344
Comments	348
Adding comments	349
Editing comments	350
Tasks	351
Comment features	355
Mark comments as read	360
Mark comments as unread	361
Archiving comments	362
Restoring comments	363
Users	363
Viewing your user profile	364
Viewing another user's profile	371
Viewing another user's profile when you have admin or super user privileges	372
Groups	373
Listing groups	374
Viewing a group	375
Projects	381
Listing projects	381
Viewing a project	382
Private projects	387
Workflows	388
Listing workflows	389

Viewing a workflow that you own	390
Viewing a shared workflow that you do not own	391
Viewing the global workflow	392
Work-in-progress tag	393
Summary	394
Using the work-in-progress tag	394
Tests	395
Listing tests	395
Viewing a test that you own	397
Viewing a shared test that you do not own	397
Notifications	398
@mention notifications	402
(1) Committed change notifications	403
(2) Review start notifications	404
(3) Moderator notifications	404
(4) Group member notifications	404
(5) Disable notifications	405
(6) Comment author notifications	405
(7) Tests have finished	405
Log in/Log out	405
Log in with a password	405
Log in with SSO	406
Log out of Swarm	406
require_login	407
Notable minor features	407
Quick URLs	407
Links in descriptions and comments	409
Search	412
Jira integration	412
Avatars	413
Following	413
Time	413
Keyboard shortcuts	413
About Swarm	414
Custom error pages	414
Short links	414
Markdown	414

Common text styles	415
Text styling example	417
Markdown in comments and review descriptions	420
Markdown in projects	420
5 Code reviews	422
Benefits	422
Facilities	422
Workflow	423
Models	424
Pre-commit model	424
Post-commit model	424
Internal representation	424
Reviews list	428
Review filters	430
Review display	433
Review summary	435
Review state	436
Vote buttons	436
Change state button	437
Review actions button	437
Review description	443
Information panel	445
Files tab	453
Comments tab	463
Activity tab	464
Activities	465
Start a review	465
Update a review	467
Fetch a review's files	468
Deleting shelves	469
Edit reviewers	470
Add a changelist to a review	471
Append a pending changelist to a review	473
Replace review with a pending changelist	474
Replace review with a committed changelist	475
Responsibility	476

Authors	476
Moderators	477
Required reviewers	478
Add yourself as a reviewer	479
Remove yourself as a reviewer	479
Review workflow	479
Basic review workflow	480
Additional workflow tasks	481
Review creation and modification outside of Swarm	483
States	485
Self-approval by review authors	486
State change restrictions with moderation	486
State change restrictions without moderation	487
Required reviewers	488
Change review state	488
6 Groups	491
Add a group	491
Edit a group	497
Delete a group	497
7 Projects	498
Add a project	498
Project settings	499
Participants tab	499
Branches tab	504
Integrations tab	511
General Settings tab	511
Edit a project	514
Membership	514
Add a member	515
Remove a member	516
Owners	517
Moderators	517
Default reviewers	518
Retain default reviewers	519
Delete a project	521
8 Workflows	523

Why should I use Swarm workflow?	523
Workflow overview	526
Workflow rules	526
Workflow basics	529
Example workflows	530
Merging multiple workflows	532
Add a workflow	537
Edit a workflow	542
Delete a workflow	542
9 Tests	544
Add a test	544
Pass special arguments to your test suite	548
Edit a test	551
Delete a test	551
10 Integrations	552
Jira	552
Jira integration	553
Jira integration and Perforce job links	556
LibreOffice	560
Limitations	561
Installation	561
Zip archive	562
Installation	562
Configuration	562
Download files as a Zip archive	563
TeamCity	564
11 Administration	565
Automated deployment for reviews	566
Automated testing for reviews	567
Configuring Jenkins for Swarm integration	570
Avatars	575
Disable avatar lookups	576
Backups	577
Changelist files limit	577
Client integration	578

Customized Swarm installations	579
Comment attachments	580
Comment attachment storage	580
Maximum comment attachment size	581
Comment mentions	581
Regular expressions in exclude lists	583
Comments	585
Comment notification delay	585
Comment threading	586
Commit credit	586
Commit-edge deployment	587
Swarm commit-edge configuration	587
P4V Authentication	588
Commit timeout	589
Configuration overview	589
Diff configuration	602
max_diffs	602
max_size_chunk	602
Email configuration	602
Sender	604
Transport	605
Recipients	605
notify_self	605
Use BCC	606
Use Reply-To	606
Email subject prefix	606
Save all messages to disk	607
Email headers	607
Email thread indexing	608
Emoji	608
Environment	609
mode	610
hostname	610
external_url	611
base_url	611
logout_url	612

emoji_path	613
Excluding Users from Activity Streams	613
Extensions Management	614
List	614
Delete	615
Install	616
Config	617
Tokens	620
Disable	620
Enable	620
Save	621
Load	621
Ping	622
Files configuration	622
max_size	622
download_timeout	623
allow_edits	623
Groups configuration	624
Ignored users for reviews	625
Job keywords	625
Adding a new job keyword	626
License	628
Localization & UTF8 character display	629
Localization	629
Non-UTF8 character display	630
UTF8 conversion	631
Logging	631
Web server logging	632
Helix Core Server logs	632
Swarm logs	632
Logrotate	633
Reference ID	634
Trigger Token Errors	634
Performance logging	635
Mainline branch identification	636
Regular expressions	637

Project mainline and README check	638
Markdown	639
Menu helpers	640
Add a menu item to a menu	640
Make Groups menu item visible to admin-user and above only	643
Create a menu item using only the custom module name	644
Add a menu item to the project menu of a specific project	644
Multiple Helix Core Server instances	646
Set up the Swarm configuration file for the Helix Core Servers	647
Configure event notification for each Helix Core Server	650
Configure a cron job for each Helix Core Server instance	658
Further information	660
Notifications	662
Global settings	664
Obliterate Review	668
When you obliterate a review	669
Obliterate a review	669
P4TRUST	669
If a certificate changes	670
Projects	670
Restrict project name and branch definition editing to administrators	670
Limit adding projects to administrators	671
Limit adding projects to members of specific groups	671
Project readme	672
Projects tab initial fetch	673
Allow project members to view project settings	673
Redis server	673
Swarm connection to Redis	674
Redis server configuration file	676
Manually verify the Redis caches	677
Review cleanup	678
Review keyword	680
Review keyword configuration	680
Create a review	682
Add a changelist to a review	682
Remove mentions from review descriptions	685

Reviews filter	686
filter-max	686
result_sorting	687
date_field	687
Reviews	687
Allow author change	687
Allow author obliterate review	688
Disable approve for reviews with open tasks	688
Disable self-approval of reviews by authors	689
Disable tests on approve and commit	689
Protected end states	689
Expand all file limit	690
Expand group reviewers	691
Maximum initial reviewers and reviewer groups in secondary navigation bar	691
Moderator behavior when a review spans multiple branches	692
More context lines	692
Process shelf file delete when	693
Synchronize review description	694
Review complexity	695
Maximum files limit	696
Search	696
maxlocktime	697
p4_search_host	697
Comment attachment thumbnail and blur generation	697
Security	698
Require login	698
Prevent login	699
Auto-create new users	700
Sessions	700
X-Frame-Options header	702
Disable commit	702
Restricted Changes	703
Limit adding projects to administrators	704
Limit adding projects to members of specific groups	704
IP address-based protections emulation	704
Disable system info	706
HTTP client options	706

Strict HTTPS	708
Apache security	708
PHP security	709
proxy_mode	710
Short links	711
Single Sign-On PHP configuration	712
Helix Authentication Service	712
Swarm SAML 2.0 settings	713
Skip activity	718
Slack integration	719
Slack configuration	720
Setup project configurations	721
Slack notification trigger	723
swarm_root	724
System Information	725
Perforce	725
Log	726
PHP Info	727
Queue Info	727
Cache Info	731
Tag processor	732
wip tag configuration	733
Test definitions	734
project_and_branch_separator	734
Trigger options	735
Command-line options	735
Configuration items	740
Unapprove modified reviews	742
Uninstall Swarm	742
Background	743
Uninstall steps	743
Upgrade index	745
status_refresh_interval	745
batch_size	746
Users	746
Dashboard refresh interval	746

Display full names	747
Review preferences	747
Time preferences	749
Workers	749
Worker status	750
Worker configuration	750
Manually start workers	751
Manually restart workers	751
Workflow global rules	752
Workflow prerequisites	752
workflow	752
Global workflow	753
12 Extending Swarm	761
Resources	761
jQuery	761
JavaScript resources	761
PHP resources	761
Laminas Framework resources	762
Development mode	762
To enable development mode:	762
To disable development mode:	763
Modules	763
Module overview	764
Task details	773
Example email module	776
CSS & JavaScript	788
Sample JavaScript extensions	789
Sample CSS customizations	790
CSS customization when Swarm is connected to multiple Helix Core Server instances	794
Swarm API	796
API versions	796
Current API versions	796
Unsupported API versions	797
Get Swarm version information	798
Swarm API basics	798
Authentication	798

Requests	799
Pagination	801
Responses	801
API endpoints (v9)	802
Activity : Swarm Activity List	802
Cache: Swarm Cache	813
Changes : API controller providing a service for changes	825
Comments : Swarm Comments	830
Groups : Swarm Groups	847
Login : Swarm Login API	864
Projects : Swarm Projects	878
Reviews : Swarm Reviews	907
Servers : Swarm Servers API	954
Users : Swarm Users	955
Workflows : Controller for querying, creating and updating workflows	959
API (v9) examples	985
API endpoints (v10)	1003
Activity: Swarm activity	1003
Attachments: Swarm comment attachments	1020
Changes: Swarm changes	1028
Comments: Swarm comments	1044
Files: Swarm files	1099
Groups: Swarm groups	1117
Jobs: Swarm Perforce jobs	1122
Login: Swarm login	1126
Menus: Swarm main menu items	1140
Projects : Swarm Projects	1147
Reviews: Swarm reviews	1174
Search: Swarm search	1291
Specs: Spec fields	1302
Testdefinitions: Swarm Test definitions	1309
Testruns: Swarm Test Integration	1320
Users: Swarm users	1343
Workflows : Swarm workflows	1349
Glossary	1377
Getting help	1396
License statements	1397

How to use this guide

This guide tells you how to use Helix Swarm for collaboration and code review for teams using Helix Core Server. It is intended for anyone using Swarm to perform code review tasks with Helix Core Server.

This section provides information on typographical conventions, feedback options, and additional documentation.

Syntax conventions

Helix documentation uses the following syntax conventions to describe command line syntax.

Notation	Meaning
<code>literal</code>	Must be used in the command exactly as shown.
<i>italics</i>	A parameter for which you must supply specific information. For example, for a <i>serverid</i> parameter, supply the ID of the server.
<code>-a -b</code>	Both <i>a</i> and <i>b</i> are required.
<code>{-a -b}</code>	Either <i>a</i> or <i>b</i> is required. Omit the curly braces when you compose the command.
<code>[-a -b]</code>	Any combination of the enclosed elements is optional. None is also optional. Omit the brackets when you compose the command.
<code>[-a -b]</code>	Any one of the enclosed elements is optional. None is also optional. Omit the brackets when you compose the command.
<code>...</code>	Previous argument can be repeated.

- `p4 [g-opts] streamlog [-l -L -t -m max] stream1 ...`
means 1 or more stream arguments separated by a space
- See also the use on `...` in [Command alias syntax](#) in the *Helix Core Command-Line (P4) Reference*

Tip

`...` has a different meaning for directories. See [Wildcards](#) in the *Helix Core Command-Line (P4) Reference*.

Feedback

How can we improve this manual? Email us at manual@perforce.com.

Other documentation

See <https://www.perforce.com/support/self-service-resources/documentation>.

Tip

You can also search for Support articles in the [Perforce Knowledgebase](#).

Earlier versions of this guide

Earlier versions of this guide: [2022.1](#), [2021.2](#)

To find even earlier versions of this guide, use the following URL and replace *v18.2* with the version number you are looking for: <https://www.perforce.com/manuals/v18.2/swarm/index.html>

Helix Swarm User Workflow Guide

For a step-by-step walk-through of setting up a workflow, adding it to a project, and using it to progress through a review, see the [Helix Swarm User Workflow Guide](#).

1 | About Helix Swarm

Swarm enables collaboration and code review for teams using Helix Core Server, helping ship quality software faster. Swarm is a collaboration tool for all types of intellectual property. It unites teams to foster brainstorming, formalize content reviews, and keep projects moving forward. Contributors share files, comment, suggest tasks, vote up or down, and submit work directly within its elegant, web-based interface. Swarm automates the entire process via notifications and makes it easy to monitor progress.

Swarm videos

Watch the following videos for an overview of Swarm and to understand its basic operation:

[Introduction to Swarm](#)

[Create a Project in Swarm](#)

[Create a Swarm Review](#)

[Collaborate on a Review in Swarm](#)

Related topics

Getting started:

- If you are experienced with code review but unfamiliar with Swarm, start with the [Quick reference](#) and [Quickstart](#) sections.
- If you are unfamiliar with code review, start with the [Basics](#) section.
- If you are unfamiliar with the Swarm workflow feature, start with the [Workflow overview](#) section.
- For a step-by-step walk-through of setting up a workflow, adding it to a project, and using it to progress through a review, see the [Helix Swarm User Workflow Guide](#).

Users:

- [Code reviews](#)
- [Projects](#)
- [Groups](#)
- [Workflows](#)

Administrators:

- [Swarm administration](#)
- [Swarm Integrations](#)
- [Extending Swarm](#)

- [Install and configure Swarm](#)
- ["Upgrading Swarm" on page 241](#)

What's new

This section provides a summary of the notable changes in Swarm for the 2022.2 release. Full details are available in the distribution's **RELNOTES.txt** file.

Major new functionality

To share your feedback about Swarm, see ["Getting help"](#) on page 1396.

Added support for Helix Authentication Service as an SSO provider

Swarm now supports Helix Authentication Service (HAS) as a Single Sign-On (SSO) provider. This helps to simplify configuration and create a more robust SSO solution. If your Helix Core Server is configured to use HAS, Swarm will automatically authenticate with the Helix Core Server using HAS. When HAS is configured, the option to log in using SSO is displayed on the Swarm login page. See ["Log in with SSO"](#) on page 406.

Minor new functionality

Improved performance of the Swarm review page

Performance of the **Manage reviewers** dialog when handling large sets of users and groups is improved. See ["Edit reviewers"](#) on page 470.

Improved performance of My Dashboard

Performance of **My Dashboard** has been improved when loading the page. See ["My Dashboard"](#) on page 306.

Added support for Helix Core Server 2023.2

Swarm now supports Helix Core Server 2023.2. See ["Helix Core Server requirements"](#) on page 106.

Added support for P4PHP 2023.2

Swarm now supports P4PHP 2023.2 for the Swarm package and tarball installations. See ["Required PHP extensions"](#) on page 103.

Important information

Upcoming changes to SSO configuration

Swarm now supports Helix Authentication Service (HAS) as a Single Sign-On (SSO) provider. This helps to simplify configuration and create a more robust SSO solution. The [SAML](#) configuration will be removed from Swarm in a later release.

PHP 7.X versions will no longer be supported in Swarm 2024 release

Swarm will drop support for PHP 7.X in 2024. This is part of our commitment to move away from using versions of platforms that have reached End-of-Life (EOL).

PHP version 8.X addresses important security issues and helps improve Swarm's performance. You can upgrade to PHP 8.x when you upgrade Swarm.

Project level test and deploy code features to be deprecated in a later Swarm release

The project level test and deploy code features will be deprecated in a later Swarm release. We recommend you use test integration to automatically deploy code within a review. For more information, see "Add a test" on page 544.

Swarm 2022.3 or later only works with Helix Core Visual Client (P4V) 2021.3 or later

When upgrading to Helix Swarm 2022.3 or later, ensure that you upgrade Helix Core Visual Client (P4V) to 2021.3 or later.

The emulate_ip_protections configurable now defaults to false

From Swarm 2022.2, the default setting of the `emulate_ip_protections` configurable in the `SWARM_ROOT/data/config.php` file is now set to false. If you are upgrading Swarm and have `emulate_ip_protections` set to true, the upgrade will not change your setting.

API version updated to v11 for Swarm 2022.1

For Swarm 2022.1, we introduced a new set of v11 APIs. These refine and extend the v10 APIs, and further standardize the endpoint and response pattern. The v10 APIs will continue to be available for some time to come

- Any improvements made in the migration to v11 will not be backported to v9 or v10.
- New API endpoints will be created as v11 and will not be backported to v9 or v10.

Swarm no longer supports Microsoft Internet Explorer

Swarm 2022.1 and later no longer supports Microsoft Internet Explorer. This is part of our commitment to move away from using platforms that have reached End-of-Life (EOL).

Swarm no longer supports CentOS 8

Swarm 2022.1 and later no longer supports Swarm installation on CentOS 8. This is part of our commitment to move away from using versions of platforms that have reached End-of-Life (EOL).

Swarm no longer supports Ubuntu 16.04

Swarm 2021.1 and later no longer supports Swarm installation on Ubuntu 16.04. This is part of our commitment to focus on supported technology platforms.

Swarm no longer supports CentOS 6 and RHEL 6

Swarm 2020.2 does not support Swarm installation on CentOS 6 and RHEL 6. This is part of our commitment to focus on supported technology platforms.

Upgrading from Swarm 2019.1 and earlier

Swarm 2019.2 introduced a Redis in-memory cache to improve performance and reduce the load on the Helix Core Server. This replaces the file-based cache that was previously used by Swarm.

On Swarm systems with a large number of users, groups, and projects, the initial population of this cache can take some time. If you have a large Swarm system you should read through the Redis server connection and configuration options before installing or upgrading Swarm, see ["Redis server" on page 673](#).

Known limitations

After an upgrade to Swarm 2022.3, you may see partial text labels or unexpected page artifacts

After an upgrade to Swarm 2022.3, if you see partial text labels or unexpected page artifacts, clear your web browser cache and refresh Swarm in your web browser.

No activity event is created when a project setting is updated

When a project setting is updated, Swarm does not create an activity.

Issue with userids containing a / character

Issue: Logging in to Swarm 2020.2 or earlier with a *userid* that contains a forward slash (/) character results in a spinner for the user icon that does not go away.

Fix: This is fixed for new installations of Swarm 2021.1 and later, but is not fixed if you upgrade from an earlier version of Swarm. To fix this issue when you upgrade from Swarm 2020.2 or earlier, add `AllowEncodedSlashes NoDecode` to the `VirtualHost` block of your `/etc/apache2/sites-enabled/perforce-swarm-site.conf` file. For more information about the `VirtualHost` block, see "[Apache configuration](#)" on page 186.

Swarm shelvedel trigger can fail for a specific Windows use case

The Swarm `shelvedel` trigger will fail in the following specific situation:
When the Helix Core Server is hosted on Windows and a user runs the shelf delete command while in the root of one of their local drives.

For example, `c:\p4 -u bruno -c my_shelf shelve -d -c 9`

Unsupported characters in user names and group names in Swarm 2019.2 and later

Swarm does not support the following characters in user names and group names: `:@{}()`

Swarm support for the "Private editing of streams" feature in Helix Core Server 2019.1 and later

Supported: stream specs can be edited in your workspace using the [Private editing of streams](#) feature and they are displayed in reviews.

Not supported: Swarm cannot commit locked stream specs.


Multiple Helix Core Server instances on a single Swarm instance

Issue: Swarm will lose connection to all of the Helix Core Servers if you edit the [base_url](#) configurable value in the `environment` block of `<swarm_root>/data/config.php`. This will stop your system working.

Fix: Remove the `base_url` configurable from the `environment` block of `<swarm_root>/data/config.php`.

Global Dashboard does not support Single Sign-On (Helix Authentication Service)

Issue: If Helix Authentication Service is enabled for Swarm and the **Try to login to all available servers with these credentials** checkbox or the **All available servers** option is selected in a login dialog, Swarm will not try to log in to any of the other Helix Core Server instances that are configured for Helix Authentication Service.

Workaround: Log in to them individually using the instance **Log in** button  in the sidebar or by including the server instance name in the URL, for example:
`https://swarm.company.com/serverA`.

Project Commits tab can fail to show some Helix Core Server commits in the top level view

The **Project Commits** tab top level client view is made up of all of the branches of the project.

For example, Project Alpha:

Branch: QA:
`//depot/alpha/dev/QA/...`

Branch: Dev :
`//depot/alpha/dev/...`
`-//depot/alpha/dev/QA/...`

The project commits tab view is generated by processing the branches in the order that they were created in and from top to bottom for the paths in each of those branches.

For the project Alpha example above:

The first path includes all of the directories and files under `//depot/alpha/dev/QA/` in the project commits tab top level view.

The second path includes all of the directories and files under `//depot/alpha/dev/` in the project commits tab top level view.

The third path excludes all of the directories and files in `//depot/alpha/dev/QA/` from the project commits tab top level view.

Result: commits made to `//depot/alpha/dev/QA/` that should be shown for the **QA** branch are not displayed in the **Project Commits** tab top level view.

When you have a simple branch structure this can be avoided by considering this issue when you create your branches. In the example above, creating the **Dev** branch first and then creating the **QA** branch avoids the problem because `//depot/alpha/dev/QA/` is not excluded from the **Project Commits** tab top level view.

Tip

Individual branch views display the commits correctly.

Task Stream Reviews

Pre-commit reviews in a task stream are not yet supported.

2 | Quickstart

Often, the best way to learn how software works is to try it. If the interface is sufficiently discoverable, you may not need to refer to documentation much, or at all. Sometimes, you just need to see the required steps to complete a task; that's what this chapter is all about.

Swarm quick reference	29
Swarm user tasks	69
Swarm administration tasks	83


Swarm quick reference



This section contains quick reference guides to help to familiarize you with Swarm:

Menu quick reference

The Header, Main, and Project menus are used to quickly navigate around Swarm.

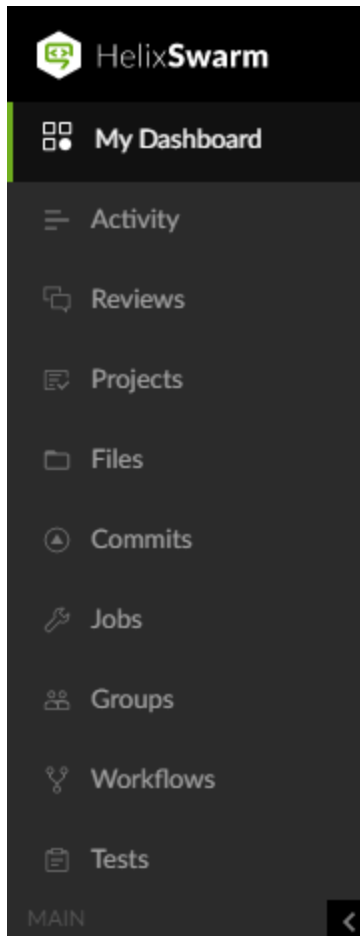
Header



-  click to display a list reviews that need your attention, see the [My Dashboard](#) page.
- **Search box:** search for projects, groups, users, and files that match your search string.
-  click to view the Swarm help.
- **Log in** (only displayed if you are not logged in to Swarm): click to log in to Swarm, see [Log in to Swarm](#).
- **User id** dropdown menu (only displayed if you are logged in to Swarm):

- **My Profile:** click for your activity, shelves, profile settings, and notification settings. For more information about your profile, see "[Viewing your user profile](#)" on page 364.
- **System Information** (*admin* and *super* users only): click to view the Swarm system information page. For details about the information on this page, see "[System Information](#)" on page 725.
- **About Swarm:** click to view Swarm version information.
If you are a user with *super* privileges, Swarm also displays the trigger token required when you configure Helix Core Server for Swarm
- **Log Out:** log out of Swarm, see "[Log out of Swarm](#)" on page 406.

Main menu




- **My Dashboard:** click to display a list reviews that need your attention, see [My Dashboard](#).
- **Activity:** click to display a list of events that have occurred recently in the associated Helix Core Server, see [Activity](#).
- **Reviews:** click to display a list of open and closed reviews. For more information about the reviews list page, see ["Reviews list" on page 428](#).
- **Projects:** click to display a list of Swarm projects. For more information about the projects page, see ["Projects" on page 381](#).
- **Files:** click to display a list of files in the Helix Core Server, starting at the top level. For more information about the files page, see ["Files" on page 322](#).
- **Commits:** click to display a list of commits made in the Helix Core Server. For more information about the commits page, see ["Commits" on page 332](#).
- **Jobs** (only displayed if Perforce jobs are configured): click to display a list of Perforce jobs in the Helix Core Server. For more information about the jobs page, see ["Jobs" on page 335](#).
- **Groups** (displayed by default but can be hidden by administrator, see ["Groups configuration" on page 624](#).): click to display the Helix Core Server groups in Swarm, see ["Groups" on page 491](#).
- **Workflows** (not displayed if the Workflow feature is disabled): click to display a list of Swarm workflows. For more information about the workflows page, see ["Workflows" on page 388](#).
- **Tests** (not displayed if the Workflow feature is disabled): click to display a list of Swarm tests. For more information about the tests page, see ["Tests" on page 395](#).
- **<** (only displayed if the main menu is expanded): click to collapse the menu.
- **>** (only displayed if the main menu is collapsed): click to expand the menu.
- **Helix Core Server instance name** (only displayed if Swarm is connected to multiple Helix Core Servers):



the instance name is displayed in the bottom left of the menu.

Project menu

1. Click the **Projects** link in the Swarm main menu.
2. Select a project from the [Projects page](#).

The project overview page is displayed:



-  click to display a list reviews that need your attention, see the [My Dashboard](#) page.
- **MAIN MENU** click to display the "Main menu" on [page 31](#).
- **Overview** (only displayed if there is a README markdown file in the project's mainline, see ["Markdown in projects" on page 420](#)): click to display an overview of the project. For more information about the project overview page, see ["Overview page" on page 383](#).
- **Activity**: click to display the activity stream for the project. For more information about the project activity page, see ["Activity page" on page 384](#).
- **Reviews**: click to display a list of open and closed code reviews for the project. For more information about the project reviews page, see ["Reviews page" on page 385](#).
- **Files**: click to display a list of files for the project. The project Files page shows a list of files for the project, starting with a folder view representing each branch. Branches are designated with the *branch* icon . For more information about the files page, see ["Files page" on page 385](#).
- **Commits**: click to display a list of changes made to the project. For more information about the commits page, see ["Commits page" on page 385](#).
- **Jobs** (only displayed if Perforce jobs are configured): click to display a list of jobs associated with the project. For more information about the jobs page, see ["Jobs page" on page 386](#).

- **Settings** (only displayed if you have permission to view or edit the project): click to display the project settings. For instructions on how to change the project settings, see "[Project settings](#)" on [page 499](#).
- **<** (only displayed if the project menu is expanded): click to collapse the menu.
- **>** (only displayed if the project menu is collapsed): click to expand the menu.
- **Helix Core Server instance name** (only displayed if Swarm is connected to multiple Helix Core Servers): the instance name is displayed in the bottom left of the menu.

Quick URLs quick reference

Swarm handles URLs intelligently to reduce the amount of information you need to enter to quickly locate what you are looking for:

1. Enter a URL in the following format:
`https://myswarm.url/<identifier>`
2. Swarm checks the `<identifier>` and redirects you to the first match it finds. Swarm checks for the identifier in following order:
 - [Reviews](#)
 - [Changelists](#)
 - [Depot paths](#)
 - [Projects](#)
 - [Jobs](#)
 - [Users](#)
 - [Groups](#)
 - [Depot names](#)

Note

- Changelist and review identifiers must be numeric.
- If you enter an identifier that does not exist for any type of resource, Swarm displays a **Page Not Found** error.

Example usage

Display the latest version of a review

For example, to display review 124: enter the URL `https://myswarm.url/124`, Swarm redirects to `https://myswarm.url/reviews/124`.

Tip

- To display a specific version of a review, enter the full URL including **reviews** and append **/v<n>/** to the URL. For example, to display version 2 of review 124: enter `https://myswarm.url/reviews/124/v2/`
- To display a diff between two versions of a review, enter the full URL including **reviews** and append **/v<n,n>/** the URL. For example, to display the diff between version 2 and 4 of review 124: enter `https://myswarm.url/reviews/124/v2,4/`
- To open a review with a specific tab open, append **#<tab name>** to the end of the URL (**#<tab name>** must be the last item in the URL):
 - **Files tab:** `#files` (default if **#<tab name>** is not specified)
 - **Comments tab:** `#comments`
 - **Activity tab:** `#activity`

Display a changelist

For example, to display changelist 123: enter the URL `https://myswarm.url/123`, Swarm redirects to `https://myswarm.url/changes/123`.

Tip

To open a changelist with a specific tab open, append **#<tab name>** to the end of the URL:

- **Files tab:** `#files` (default if **#<tab name>** is not specified)
- **Comments tab:** `#comments`

Display the latest version of a file

Enter the URL `https://myswarm.url/depot/alpha/readme.txt`, Swarm redirects to `https://myswarm.url/files/depot/alpha/readme.txt`.

Tip

- To display a specific version of a file, enter the full URL including **files** and append **?v=<n>** to the URL. For example, to display version 2 of the `depot/alpha/readme.txt` file: enter `https://myswarm.url/files/depot/alpha/readme.txt?v=2`
- To open a file with a specific tab open, append **#<tab name>** to the end of the URL (**#<tab name>** must be the last item in the URL):
 - **View tab:** `#view` (default if **#<tab name>** is not specified)
 - **Commits tab:** `#commits`

Display the user profile page for jsmith


Enter the URL `https://myswarm.url/jsmith`, Swarm redirects to `https://myswarm.url/users/jsmith`.

My Dashboard page quick reference

The purpose of the dashboard is to allow you to focus on reviews that need to be done, so that other users are not blocked. The dashboard lists the most recently modified reviews first, and shows your role in the review.

Note

Since it is tied to the logged in user, **My Dashboard** is only populated if you are logged in.

1. Log in to Swarm.
2. If the **My Dashboard** page is not already displayed, click **My Dashboard** in the menu or click the Swarm  icon above the menu.

My Dashboard

Project






















Role

Author

There are new changes to your review list

Clear all

Search for a review

Description	Your Role	Last Activity	State	Tests	Complexity	Votes
<div><div></div><div><div>Updated README to have a bit more information.</div><div>Jack Boone requested a  pre-commit review 365 for mercury:dev</div></div></div>	Reviewer	44 minutes ago			 137	 0  0
<div><div></div><div><div>Updated notes for the new year.</div><div>Jack Boone requested a  post-commit review 422 for jplugin:main</div></div></div>	Required reviewer	2 months ago			 18	 0  0
<div><div></div><div><div>Add comment to method.</div><div>Steve Russell requested a  post-commit review 328 for jplugin:main</div></div></div>	Required reviewer	3 months ago			 6	 2  0

All available reviews are displayed

Jump directly to a specific area of the dashboard page using the following links:

- ["Filtering" below](#)
- ["Review summaries" on the facing page](#)

Filtering

Project	Role	Author	There are new changes to your review list  Clear all				

The filters are used to filter your dashboard reviews:

- **Project:** filter by the project the review is part of, limiting results to **My Projects** (projects you are an owner of or a member of) or specific projects. To select multiple projects, click the projects you want to filter by. The **Project** filter will only show projects for which there are reviews in your dashboard.

- **Role:** filter by your specific role in a review, limiting results to reviews for which you are the author, a reviewer, a required reviewer, or a moderator. The **Role** filter will only show roles for which there are reviews in your dashboard.
- **Author:** filter the reviews to only those that have been authored by a certain user. Select the author from the drop down list of users. The **Author** filter will only show authors for which there are reviews in your dashboard.
- **Clear all:** click to reset all dashboard filters back to their defaults.
- **Search:** filter the reviews by searching the review descriptions.

Refresh button


Refresh button (only displayed when new content is available for your dashboard): click **Refresh** to update your dashboard with the new content.


Note

By default, when your dashboard is open, Swarm checks for new content every five minutes.

Review summaries

The dashboard displays a summary for each review.




Description	Your Role	Last Activity	State	Tests	Complexity	Votes
 Updated README to have a bit more information. Jack Boone requested a  pre-commit review 365 for mercury:dev	Reviewer	about an hour ago			 137	 2  0

- **Avatar** displays the avatar of the review author, click to view the profile of the author, see ["Viewing another user's profile" on page 371](#).
- **Description** contains:
 - first line of the review description, click to open the review. If the review description is too long it is truncated, click on the ellipsis  to expand it in the list page.
 - review author, click to view the profile of the author.
 - review type, either a [pre-commit review](#) or a [post-commit review](#).
 - review number, click to open the review, see ["Review display" on page 433](#).
 - review project branches, click to open the project.
- **Your role:** displays your role in the review and is the reason the review is in your dashboard. This can be Author, Reviewer, Required Reviewer, or Moderator.
- **Last activity:** displays the last time that any changes were made to the review, including votes, comments, commits, and file changes.
- **State:** a review can be in one of the following states:

- **Needs review:** The review has started and the changes need to be reviewed.
- **Needs revision:** The changes have been reviewed and the reviewer has indicated that further revisions are required.
- **Approved:** The review has been approved. The changes may need to be committed.

Note

The Approved state only applies to review authors, it is only shown for a review that has been approved but has not been committed.

- **Tests:** shows the test suite state for the review, either tests in progress , tests passed , or tests failed .
- **Complexity:** a traffic light icon and number shows the relative complexity of the review and the total number of lines changed in the review. [Complexity can be configured](#) by your Swarm administrator but, by default:
 - **Red:** ≥ 300 changes
 - **Amber:** < 300 and > 30 changes
 - **Green:** ≤ 30 changes

Tip

- Review complexity is only calculated for a review when the review is updated and the file content has changed.
- Review complexity is only stored for the current version of a review.

Hover over the complexity icon to display more detailed information:



- **Votes:** displays the number of up votes and down votes for the review.

Global dashboard page quick reference

Note

Global dashboard is only available if your Swarm administrator has configured Swarm to connect to more than one Helix Core Server instance.

The purpose of the global dashboard is to allow you to focus on reviews that need to be done, so that other users are not blocked. The global dashboard lists reviews by Helix Core Server according to the most recently modified first, and shows your role in the review.

1. Enter the basic Swarm URL without a Helix Core Server instance name, for example:
`https://swarm.company.com`
2. If the log in dialog is displayed it is the same dialog that is displayed when you click **Log In** from the banner, log in to the global dashboard.

For instructions about logging in to the global dashboard, see "[Log in to one or more Helix Core Server instances](#)" on page 312.

Tip

- Since it is tied to the logged in user, the global dashboard is only populated for the Helix Core Server instances you are logged in to in Swarm.
- If you are already logged in to a Helix Core Server instance in Swarm, the dashboard for that server will be automatically populated when you open or refresh the global dashboard.

The global dashboard is displayed:

The screenshot shows the Global Dashboard interface. On the left, a sidebar lists two instances: 'main' (ssl:1666) and 'assets' (ssl:2666). The main area displays a table of reviews. The table has columns for Author, ID, Description, Project(s), Your Role, State, and Last activity. The reviews are grouped by instance, with 'main' having five reviews and 'assets' having one. Each review entry includes a user profile picture, a review ID, a description, the project name, the reviewer's role, a state icon, a progress bar, and the time since the review was made.


Author	ID	Description	Project(s)	Your Role	State	Last activity
main						
[Profile]	277	Added endpoint for initiating server shutdown.	mercury: dev	Reviewer	[Icon]	24 minutes ago
[Profile]	279	Keep track of all out clients we have active.	mercury: dev	Author	[Icon]	9 days ago
[Profile]	271	Added first set of icons for the application.	mercury: main	Reviewer	[Icon]	14 days ago
[Profile]	135	Tidying up some of the code to fix some of the ed...	jplugin: main	Reviewer Moderator	[Icon]	14 days ago
[Profile]	249	Adding description to P4Credentials & P4Credent...	jplugin: main	Moderator	[Icon]	14 days ago
assets						
[Profile]	4	Updated the README with some context.	mercury: main	Author	[Icon]	29 minutes ago

The global dashboard is made up of three main areas:


- **"Toolbar" on the next page:** log in to Helix Core Server instances, log out of all Helix Core Server instances, and open the Swarm help from the global dashboard toolbar.
- **"Sidebar" on the next page:** Log in/logout of individual Helix Core Server instances and view your profile for any instance that you are logged in to.


- ["Helix Core Server dashboards" below](#): view, filter and search the dashboards of the Helix Core Server instances you are logged in to. Jump directly to a Helix Core Server, review, or project by clicking on a link.




Toolbar

- **Log In** (only displayed in the toolbar if you are not logged in to any Helix Core Server instances): click to log in to one or more Helix Core Server instances, see ["Log in to one or more Helix Core Server instances" on page 312](#).
- **Logout from all instances** (only displayed in the toolbar if you are logged in to at least one Helix Core Server instance): click to logout from all of the Helix Core Server instances, see ["Logout of all Helix Core Server instances" on page 314](#).
- **Help** : click to view the Swarm help.

Sidebar

- **Log in status:**
 - Not logged in to the Helix Core Server instance:







A grey rectangular box with a grey circle icon on the left. To the right of the icon, the text "assets" is displayed in blue, and "ssl:2666" is displayed in grey below it.
 - Logged in to the Helix Core Server instance:

A grey rectangular box with a green circle icon on the left. To the right of the icon, the text "main" is displayed in blue, and "ssl:1666" is displayed in grey below it.
- Helix Core Server instance actions:
 - **Log in**  (only available if you are not logged in to the Helix Core Server instance): hover over the instance name and click the **Log in** button  when it is displayed, see ["Log in to a Helix Core Server instance" on page 315](#).
 - **Logout** (only available if you are logged in to the Helix Core Server instance): hover over the instance name, click the **Gear** button  , and select **Logout** from the dropdown menu to logout of the instance, see ["Logout of a single Helix Core Server instance" on page 317](#).

Helix Core Server dashboards

The global dashboard lists reviews by Helix Core Server according to the most recently modified first, and shows your role in the review.

Example global dashboard showing a number of Helix Core Server dashboards:

All Instances ▾		All Projects ▾		All Roles ▾		Authored by	Search	
Author	ID	Description	Project(s)	Your Role	State	▲/▼	^/▽	Last activity
main								
	277	Added endpoint for initiating server shutdown.	mercury: dev	Reviewer	👁	🔴	0 / 0	24 minutes ago
	279	Keep track of all out clients we have active.	mercury: dev	Author	✍	▲	0 / 0	9 days ago
	271	Added first set of icons for the application.	mercury: main	Reviewer	👁	▲	0 / 0	14 days ago
	135	Tidying up some of the code to fix some of the ed...	jplugin: main	Reviewer Moderator	👁	🔴	1 / 0	14 days ago
	249	Adding description to P4Credentials & P4Credent...	jplugin: main	Moderator	👁	▲	1 / 0	14 days ago
assets								
	4	Updated the README with some context.	mercury: main	Author	✍	▲	0 / 0	29 minutes ago

Jump directly to a specific area of the dashboard page using the following links:

- ["Dashboard filters" below](#)
- ["Dashboard content" on the next page](#)

Dashboard filters

All Instances ▾	All Projects ▾	All Roles ▾	Authored by	Search
-----------------	----------------	-------------	-------------	--------

Used to filter your dashboard reviews, the following filters are available:







- **Instances:** filter by the Helix Core Server instance, limiting results to **All Instances** or to an individual instance.
- **Projects:** filter by the project the review is part of, limiting results to **All Projects** or to an individual project. The **Project** filter will only show projects for which there are reviews in your dashboard.
- **Roles:** filter by your specific role in a review, limiting results to reviews for which you are the author, a reviewer, a required reviewer, or a moderator. The **Role** filter will only show roles for which there are reviews in your dashboard.
- **Authored by:** filter the reviews to only those that have been authored by a certain user. Type in this field to get a drop down list of users to filter by.
- **Reset** (only displayed if one or more filters are set): click the **Reset** button to reset all dashboard filters back to their defaults.
- **Search:** filter the reviews by searching the description and review ID.

Dashboard content

The dashboard for each Helix Core Server shows a summary of the information for each review.

Tip



- Click on the Helix Core Server name in the header bar to open Swarm for that instance in a new tab.
- Click on the header bar to the right of the Helix Core Server name to collapse the dashboard for that instance. Click again to expand the dashboard for the instance .

Author	ID	Description	Project(s)	Your Role	State	 	Last activity
main 							
	277	Added endpoint for initiating server shutdown.	mercury: dev	Reviewer	 	0 / 0	1 hour ago

- **Author** avatar: the review author's avatar, hover over the avatar to see the ID and name of the review author. Click on the avatar to go to the profile of the review author, see ["Viewing another user's profile" on page 371](#).
- **ID**: the unique number used to identify the Swarm review. Click the review ID to display the review, see ["Review display" on page 433](#).
- **Description**: the review description may be truncated if it is too long, in which case click on the review description to expand it.
- **Project(s)**: lists the project branches this review covers. A review may span multiple branches and projects. Click a branch link to navigate to the project page for that branch.
- **Your role**: displays your role in the review and is the reason the review is in your dashboard. This can be Author, Reviewer, Required Reviewer, or Moderator.
- **State**: a review can be in one of the following states:
 - **Needs review**: The review has started and the changes need to be reviewed.
 - **Needs revision**: The changes have been reviewed and the reviewer has indicated that further revisions are required.
 - **Approved**: The review has been approved. The changes may need to be committed.

Note

The Approved state only applies to review authors, it is only shown for a review that has been approved but has not been committed.

- **Type**: displays the type of review, either a [pre-commit review](#)  or a [post-commit review](#) .
- **Votes up/down**: displays the number of up votes and down votes for the review.

- **Last activity:** displays the last time that any changes were made to the review, including votes, comments, commits, and file changes.

Reviews list page quick reference

The **Reviews** page lists code reviews that are in progress, and code reviews that are complete.

1. Log in to Swarm.
2. Click **Reviews** in the main menu.

Description	Created	State	Tests	Complexity	Comments	Votes
Added Windows to list of build platforms. Allison Clayborne requested a post-commit review 396, 31 minutes ago for jplugin:main	31 minutes ago	Open	0	137	1	0
Tidying up some of the code base. Allison Clayborne requested a pre-commit review 391, 3 days ago for jplugin:main	3 days ago	Open	0	18	1	0
Added some state exception checking. Jack Boone requested a pre-commit review 388, 4 days ago for jplugin:main	4 days ago	Open	0	6	2	2
Update to the 12th entry point, which is defined to measure the latency Jack Boone requested a pre-commit review 385, 4 days ago for maker:master	4 days ago	Open	0	1462	0	1
Update the README to use markdown rather than the old format. Jack Boone requested a pre-commit review 382, 4 days ago for blue-book:main	4 days ago	Open	0	2	0	0

Jump directly to a specific area of the **Reviews** page using the following links:

- ["Opened and Closed tabs" below](#)
- ["Review filters" on the next page](#)
- ["Reviews list content" on page 47](#)

Opened and Closed tabs

The **Opened** and **Closed** tabs display:

- **Opened** tab: displays a list of all code reviews that have started, are being reviewed, are awaiting revisions, or need to be committed.
- **Closed** tab: displays a list of all code reviews that have been approved and committed, rejected, or archived.

Tip

- The **Opened** and **Closed** tabs display the number of reviews in each tab. Initially this can be an over estimate which Swarm dynamically corrects as more information becomes available to it.
- The **Closed** tab initially displays a – character. The – is replaced by a number when you navigate to the **Closed** tab.

Review filters

Tip

Swarm updates the URL in your browser to reflect the filter options you have selected. This makes it easy to bookmark or share your review list filter settings with the URL.

If you share a URL with **Comment** or **Vote** filters selected, those filters are applied to the user running the URL. This means that their reviews list page will contain different reviews to your reviews list page.



The following filters are available (from left to right):

- **Branch** (only available on the project ["Reviews page" on page 385](#)): a dropdown menu that lets you filter which reviews to display based on the current project branches:

Branch 

- **Select branch:** all reviews for all of the selected branches are displayed. To select multiple branches, click the branches you want to filter by.
- **Branch search:** an auto-complete search field that allows you to choose one of the branches within the current project. Once specified, only reviews for the selected project branch are displayed.
- **Project** dropdown: filter by the project the review is part of:

Project 

- **My projects:** displays all reviews for all of the projects you are participating in, as a member, owner, moderator, or follower.
- **Select project:** displays all reviews for all of the selected projects. To select multiple projects, click the projects you want to filter by.
- **Project search:** an auto-complete search field that allows you to choose one of the projects defined in Helix Core Server. Once specified, only reviews for the selected project are displayed.

Tip

Swarm creates a projectid based on the project name whenever a project is created. For example, a project called **Test Data** is given a projectid of **test-data**. If the project name changes, the projectid does not change.

If your search finds a match in a projectid but not in the associated project name, the project name for that projectid is returned in the search results. For example, if the project name was changed from **Test Data** to **Sample Data**, a search for *test* will return the **Sample Data** project in the search results.

- **Role** dropdown: filter reviews based on user involvement, options are:

Role 

- **All reviews:** displays all reviews.
- **Reviews I've authored:** displays reviews that you have authored.
- **Reviews I'm participating in:** displays reviews that you are a reviewer of, but not an author of.
- **Reviews I'm an individual reviewer of:** displays reviews that you are an individual reviewer of, but not a group reviewer of, or an author of.
- **Review I've authored or I'm participating In:** displays reviews that you have authored, or are a reviewer of.
- **Authored by:** an auto-complete search field that allows you to choose one of the user accounts defined in the Helix Core Server. Once specified, only reviews authored by the user are displayed.

When you select one of the available options, the list of options updates to match the currently selected filter, and the **Role** dropdown indicates the current filter: **All, Author, Participant, Author or Participant, Individual Reviewer, or user.**

- **Reviewers** buttons, (**Opened** tab only):



- **Has reviewers:** displays reviews that have one or more reviewers.
- **No reviewers:** displays reviews that have no reviewers.

- **States** buttons, (**Opened** tab):




- **Needs Review:** displays reviews that need to be reviewed.
- **Needs Revision:** displays reviews that have been reviewed, but need further revisions before the review can be accepted.
- **Approved:** displays reviews that have been approved, and should be committed.

- **States** buttons, (**Closed** tab):



- **Approved:** displays reviews that have been approved and committed.
 - **Rejected:** displays reviews that have been rejected.
 - **Archived:** displays reviews that have been archived.
- **Tests** buttons (when automated tests are enabled for the associated project):



- **Tests passed** displays reviews that have passed tests.
 - **Tests failed:** displays reviews that have failed tests.
- **Comments** buttons:



- **I have commented on:** displays reviews that you have commented on.
- **I have not commented on:** displays reviews that you are participating in but have not commented on

Note

Filters for commenting only apply to reviews which you are a participant of. Commenting on or voting on a review will automatically add you as a participant. If you leave the review after commenting on it, then this review will not be included in the list.

- **Votes** buttons:



- **I have voted up:** displays reviews that you have voted up.
- **I have voted down:** displays reviews that you have voted down.
- **I have not voted on:** displays reviews that you are participating in but have not voted on.
































Note


Filters for voting only apply to reviews which you are a participant of. Commenting on or voting on a review will automatically add you as a participant. If you leave the review after voting on it, then this review will not be included in the list.


- **Clear all:** click to clear all of the filters.
- **Search for a review:** search for a partial match of review description content and an exact match for *reviewid*, *userid*, and *projectid*. The search filter boxes on the **Opened** and **Closed** tabs are independent.

Reviews list content

The review list displays a summary for each review.

Description	Created	State	Tests	Complexity	Comments	Votes
 Added Windows to list of build platforms. Allison Clayborne requested a  post-commit review 396, 31 minutes ago for jplugin:main	31 minutes ago			 137	 1	 0  0
 Tidying up some of the code base. Allison Clayborne requested a  pre-commit review 391, 3 days ago for jplugin:main	3 days ago			 18	 1	 0  0
 Added some state exception checking. Jack Boone requested a  pre-commit review 388, 4 days ago for jplugin:main	4 days ago			 6	 2	 0  2
 Update to the 12th entry point, which is defined to measure the latency Jack Boone requested a  pre-commit review 385, 4 days ago for maker:master	4 days ago			 1462	 0	 1  0




- **Avatar** displays the avatar of the review author, click to view the profile of the author
- **Description** contains:
 - first line of the review description, click to open the review. If the review description is too long it is truncated, click on the ellipsis  to expand it in the list page.
 - review author, click to view the profile of the author
 - review type, pre-commit or post-commit
 - review number and version, click to open the review
 - review creation date and time
 - review project branches, click to open the project
- **Created** or **Last activity** dropdown: click to change the order the reviews are displayed in, options are:



- **Created:** Reviews sorted by when they were created
- **Last activity:** Reviews sorted by when they were last updated

Note

- If the **Result order** button is not displayed, reviews are sorted by either **Created** or **Activity**, as set by your Swarm administrator, see ["Reviews filter" on page 686](#) for details.
- **Result order** button display is a global setting controlled by the Swarm administrator. See ["Reviews filter" on page 686](#) for details.
- When review results are older than 24 hours they are displayed in numerical order within each day.

- **State:** a review can be in one of the following states:
 - **Needs review:** The review has started and the changes need to be reviewed.
 - **Needs revision:** The changes have been reviewed and the reviewer has indicated that further revisions are required.
 - **Approved:** The review has been approved. The changes may need to be committed.
- **Tests:** displays the test suite state for the review, either tests in progress , tests passed , or tests failed .
- **Complexity:** a traffic light icon and number shows the relative complexity of the review and the total number of lines changed in the review. By default, complexity icon displays:
 - **Red:** ≥ 300 changes
 - **Amber:** < 300 and > 30 changes
 - **Green:** ≤ 30 changes

Tip

- Review complexity is only calculated for a review when the review is updated and the file content has changed.
- Review complexity is only stored for the current version of a review.

Hover over the complexity icon to display more detailed information:



- **Comments:** displays the number of open (non-archived) comments that are associated with the review. The icon is filled if there are comments on the review.

- **Votes:** displays the number of up votes and down votes for the review. The appropriate vote icon is filled if you have voted on the review, vote icons with only an outline show votes by others.

Review page quick reference

Tip


The Swarm review page is described in this documentation.


The old review page is now referred to as the classic review page. To use the classic review page, use the **Classic view** toggle switch at the top of the review page.

The Swarm Classic view option will be removed in a future Swarm release.

Help for the classic Swarm review page is available in the Swarm 2021.2 documentation. See [Review display](#).

The Swarm review page is used when reviewing changes:

1. Log in to Swarm.
2. To display the Swarm home page, click the Swarm  icon above the menu.
3. If your dashboard is not already displayed, click the **Dashboard** tab.
4. To open a review, click the **ID** of the review.



The screenshot displays the Swarm review page for review 644. The page is titled "Reviews/644" and shows the review details, including the reviewer "Allison Clayborne" and the review type "pre-commit review". The review is currently in the "Needs review" state. The page includes a "Change state" button and a "Comments (0)" section. Below the review details, there is a list of files being reviewed, including "db.c#1", "est/ risus/ jira/ xgen/ rmi.c", "magis/ scm/ vendor_x86_64/ page32.c", "metadata.c", "stub.c", and "users.c#1". The sidebar on the right shows the "Blocking Approval (1)" status, a list of reviewers (Claire Brevia, Jack Boone, Steve Russell), and a list of tasks and tests.

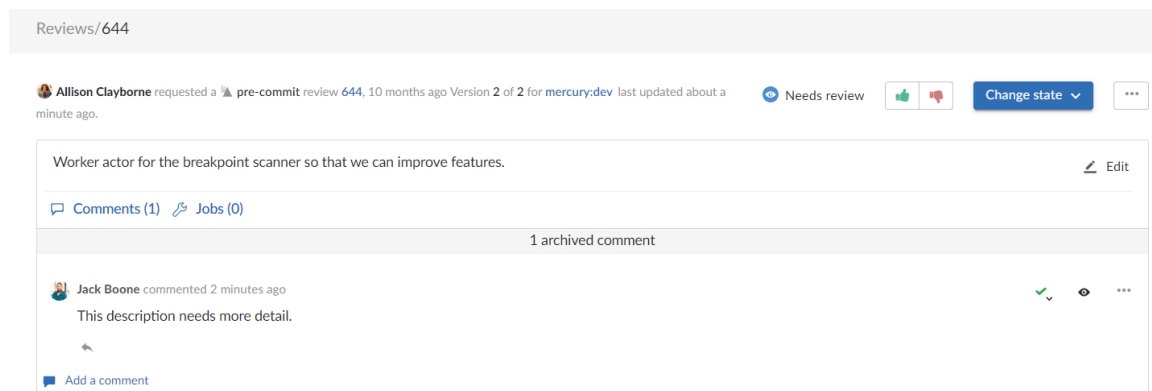
By default, when you open the Review page for the first time, the file list panel and the information panel are hidden. Use the **Show more context** buttons to view the file list panel and the information panel.

For more information about the show more context buttons, see ["File diff panel" on page 61](#).

Jump directly to a specific area of the review page using the following links:


- ["Review description" below](#)
- ["Information panel" on page 53](#)
- ["Files tab" on page 57](#)
 - ["File diff panel" on page 61](#)
- ["Comments tab" on page 65](#)
 - ["Add a comment" on page 68](#)
- ["Activity tab" on page 68](#)


Review description



The review description is made up of the following elements:

- **Review description header:**
 - **Review information:** Displays the following details:
 - **Review author** avatar and name: The avatar and name of the review author. Hover over the avatar to see the name of the review author and hover over the name to see the ID of the review author. Click on the avatar or name to go to the profile of the review author. See ["Viewing another user's profile" on page 371](#).
 - **Review type:** Displays the review type, reviews can be [pre-commit](#) or [post-commit](#).
 - **Review ID:** The unique number used to identify the Swarm review.
 - **Review raised:** When the review was requested.

- **Version:** The version of the review being viewed and the total number of review versions.
- **Project branch:** The project branch the review files are in, click to go to the project page.
- **Last updated:** When the review was last updated.
- **Review state icon:** a review can be in one of the following states:
 - **Needs review:** The review has started and the changes need to be reviewed.
 - **Needs revision:** The changes have been reviewed and the reviewer has indicated that further revisions are required.
 - **Approved:** The review has been approved. The changes may need to be committed.
 - **Rejected:** The review has been completed. The changes are undesirable and should not be committed.
 - **Archived:** The review has been completed for now but it is not rejected, or approved. The review has been filed away in case it is needed in the future.
- **Voting buttons:** click a voting button to vote the review up or down.
- **Change state button:** Click to select a new state from the review dropdown menu. State change options are only displayed if you are authorized to make the state change:
 - **Needs revision:** Select to request changes to the files in the review.
 - **Needs review:** Select to request further review of the changes.
 - **Approve** (only available if the voting requirements for the review are satisfied. For information on voting requirements, see ["Required reviewers" on page 478](#)): select to approve the review.
 - **Commit** (only available for [pre-commit](#) reviews that have been approved): Select to commit the review.
 - **Approve and commit** (only available for unapproved [pre-commit](#) reviews when the voting requirements for the review are satisfied. For information on voting requirements, see ["Required reviewers" on page 478](#).): Select to approve and commit the review in a single step. See ["Approve and commit" on page 489](#).
 - **Reject:** Select to reject the review.
 - **Archive:** Select to archive the review.
- **Review actions**  button: click the **Review actions** button to:
 - **Add change:** Select to add a changelist to the review. Options available depend on whether the review is [pre-commit](#) or [post-commit](#). For information about adding a changelist to a review, see ["Add change" on page 438](#).

- **Download zip** (if configured): Select to download a compressed archive of all of the files in the review, see ["Download files as a ZIP archive" on page 78](#).
 - **Change author** (if configured): Select to change the author of the review, see ["Change the review author" on page 440](#).
 - **Obliterate review** (by default, only available for users with *admin* or *super* user rights): See ["Obliterate Review" on page 668](#).
 - **Join review** (only if you are not a member of the review): Select to join the review. If you vote on a review, you automatically join the review and become a reviewer.
 - **Leave review** (only if you are a member of the review): Select to leave the review.
 - **Mark all comments read**: Click to mark all of the comments on this review as read.
 - **Mark all comments unread**: Click to mark all of the comments on this review as unread.
 - **Refresh projects**: Select to check if the review files are associated with any projects created or updated after the review was last updated. If any projects are found that are not already associated with the review, the review is linked to them. See ["Refresh projects" on page 442](#).
 - **Disable notifications** (only supported in the original Swarm review page): Select to disable notifications for this review.
 - **Try it out** (only supported in the original Swarm review page and if configured): If your deployment system provides a URL to the deployment, you can select it to see the deployment results.
- **Review description**: The review description is automatically copied from the changelist that was originally used to create the review.
- When the review description size exceeds the 25% of the page, only the initial part of the review description is displayed. Select the **Expand description** button to view the full review description and select the **Collapse description** button to hide a part of the review description.
- To change the review description, click the  **Edit** button. Markdown content is displayed in the review description, but Markdown support is limited to prevent execution of raw HTML and JavaScript content. For information about Markdown, see ["Markdown in comments and review descriptions" on page 420](#)
- **Update pending changelist** checkbox in the edit description dialog (only available if you are editing the description of a pre-commit review and you are the original author of the changelist that created the review): Select the checkbox to also apply your review description changes to the original changelist description.
- **Comments (n)**: Click to add a comment to the review description, view existing description comments, or hide existing description comments, see ["Commenting on a review description" on page 349](#).

- **Jobs (n)** (if configured): Perforce jobs can be linked to the review. Click to add a Perforce job to the review, view jobs linked to the review, unlink a job from the review, or hide jobs on the review. For more information on Perforce jobs, see ["Jobs" on page 335](#).
 - **Add a job:** Click to add a job to the review. Select the job from the dialog that is displayed.
 - **jobnnnnnn** (where **nnnnnn** is the job number): Click the job number to view details of the Perforce job. For more information about Perforce jobs, see ["Job display" on page 336](#).
 - **X** button: To unlink a job from the review, click the **X** button of the job you are unlinking from the review.
- **Send All Notifications (n)** (where **n** is the number of delayed notifications): Click to manually send all of your delayed notifications for the review. For more information about comment notification delay, see ["Comment notification delay" on page 355](#).

Information panel

By default, when you open the Review page for the first time, the file list panel and the information panel are hidden. Use the **Show more context** buttons to view the file list panel and the information panel.

For more information about the show more context buttons, see ["File diff panel" on page 61](#).

The information panel on the right side of the review page contains the following elements:

- **Blocking Approval or Blocking Commit:**

Once a review has been approved, the **Blocking Approval** section is not displayed.

For pre-commit reviews, if a user adds a new requirement to an approved review, then that requirement is displayed in the **Blocking Commit** section. For example, if a project initially has the **Minimum up votes** requirement set to one and this requirement is satisfied, and the review is in approved state. Now, if the **Minimum up votes** requirement changes to two then this requirement is displayed in the **Blocking Commit** section. Only when this requirement is satisfied the user can do a commit.

- **Pending Comment Tasks:** Provides a list of pending review comments that are flagged as *tasks*. The pending comment tasks are shown only when the `disable_approve_when_tasks_open` configurable is set to true in the `SWARM_ROOT/data/config.php` file.

Note

If Swarm is configured to prevent approval of reviews with open tasks and a review has open tasks, the **Approve**, and **Approve and commit** options will not be available for the review. This option is configured by an administrator. See ["Disable approve for reviews with open tasks" on page 688](#).

To approve, or approve and commit a review with open tasks, you must address the tasks first and then set them to **Task Addressed**, or **Not a Task**. See ["Set a task to Task addressed or Not a task" on page 352](#) for details.

For more information about approving a review with pending tasks, see ["Approve a review with open tasks" on page 354](#).

- **Required Reviewers:** Provides a list of all the Individual and group reviewers that are required to vote on the review. A review can not be approved or committed until the required reviewers have voted up on the review.

A required ["Review page quick reference" on page 49](#) has a star badge over their avatar. A required ["Group reviewer" on page 447](#) either has a star badge or a star badge with a 1 over their avatar.

- **Blocking Tests:** Provides a list of workflow tests that have **Failed** or have an **In progress** status that is blocking approval for a review. A review can not be approved or committed until the review author fixes all the failed tests. For more information about a failed test, see ["Review page quick reference" on page 49](#).
- **Minimum Up votes:** Provides the number of **Minimum Up votes** required for each project that you have access to. This includes the maximum number of **Minimum Up votes** required for all branches in a project blocking approval or commit for a review. For more information about **Minimum up votes** for a project see, [General Settings tab](#). For more information about **Minimum up votes** for a project branch, see [Branches tab](#).


For [private project](#), the maximum number of **Minimum Up votes** for all private projects and their branches that is blocking approval for a review is displayed.

- **Moderator Votes:** Provides a list of branch moderators that are blocking approval or commit for a review. Depending on the configuration of `moderator_approval` configurable in the `SWARM_ROOT/data/config.php` file, the branch moderator approval is required. Ensure that the branch moderator approves the review and not just up votes it.


By default, when a review spans multiple branches that have different moderators, only one moderator from any one of the branches needs to approve the review.

Swarm can be configured to require that one moderator from each branch must approve the review, this is a global setting. If a moderator belongs to more than one of the branches spanned by the review, their approval will count for each of the branches they belong to. For instructions on how to configure moderator behavior, see ["Moderator behavior when a review spans multiple branches" on page 692](#).





■ Reviewers:


- **Edit**  button (if enabled): click to edit the reviewers for the review
- **Up** vote and **Down** vote count: indicates the number of up votes and down votes the review has.
- **Groups**: lists groups that are reviewers for the review. When at least one person in the group has voted, the avatar displays a badge indicating whether the group, as a whole, has voted up or down. Click on the group to see who has voted, and how they have voted, see ["Group reviewer" on page 447](#).
- **Individuals**: lists individuals that are reviewers for the review. When an individual has voted on the review, their avatar displays a badge indicating whether they voted up or down. For more information about individual reviewers, see ["Individual reviewer" on page 449](#).

■ Tasks










- **Task Summary** icons: summary of the number and status of comments flagged as tasks for the review. For more information about tasks, see ["Tasks" on page 351](#).
 - **Red Flag** icon: indicates the numbers of open tasks on the review.
 - **Green Check Mark** icon: indicates the numbers of tasks that have been addressed on the review.
 - **Blue Double-Check Mark** icon: indicates the number of tasks that have been addressed and verified on the review.
- **Show Task details**  button: click to display the **Tasks** dialog listing all tasks associated with the review. The tasks can be filtered by reporter and task type, and you can change the task state directly from the **Task** dialog. See ["Task details" on page 353](#).

- **Tests** (if configured): displays the test results for the current review version in the information panel

-  **Passed:** Displays the number of the tests that have passed for the review version.
-  **Failed:** Displays the number of the tests that have failed for the review version.
-  **Manual:** Displays the number of **On Demand** tests that can be run manually for the review version.
-  **In Progress:** Displays the number of tests that are running for the review version.









Show test details  button: Click to view test run information for tests associated with the review version:


Tests

Status	Title		
 Passed	main-codeSniff		Rerun test 
 Passed	main-EsLint		Rerun test 
 Failed	main-Jest		Rerun test 


Messages

Error contacting server

 Passed	main-jsLint		Rerun test 
 Manual	main-library		Run test 
 In progress	main-modules		



If your continuous integration system calls back to Swarm with a URL, the test run is linked to the URL provided. If the test result contains any messages, click the **Show more**  button to the left of the test to view them

When the content of a review is unchanged between review versions, the tests are not rerun because the files have not changed. In this case, the earlier test results are displayed and marked with the word **(Copy)**. You can manually run these tests if required by using the



Run test button.

The following actions are available from the **Tests** dialog for tests called by a workflow:

- **Rerun test:** If a test has completed you can rerun the test, click  to rerun the test for the most recent version of the review.
- **Run test:** If the review workflow has a test set to run **On Demand**, click  to manually run the test for the most recent version of the review.

Tip

- The **Run test** and **Rerun test** buttons are only available if you are logged in to Swarm, viewing the latest version of the review, and the test was called by a workflow.
- The **Rerun test** button is not available for tests that are in progress.
- When you click **Run test** or **Rerun test**, all of the buttons in the list are temporarily disabled while Swarm starts the test and checks if any other tests are in progress.
- The **Rerun test** and **Run test** buttons are only available for tests that are run by a workflow. They are not available for tests that are run because they are configured on a project.

Note


Private projects: if a test for a private project is added to your review because **Iterate tests for affected projects and branches** is selected for the test, Swarm honors the private project's permissions and displays it as **Private project** in the test list to users that do not have permission to view it.

- For information about iterating tests, see "[Iterate tests for affected projects and branches checkbox](#)" on page 546.
- For information about private projects, see "[Private projects](#)" on page 387.

- < (only displayed if the information panel is expanded): click to collapse the panel.
- > (only displayed if the information panel is collapsed): click to expand the panel.

Files tab


Use the **Files** tab to view the files in the review and to see how they have changed using the Swarm diff view.





Swarm supports stream specs in your workspace using the [Private editing of streams](#) feature. If a changelist or review contains a stream spec, it will be displayed first in **Files** with the prefix **stream: //**, for example: **stream: //MyStreamDepotName/MyStreamSpecLocationName**. A changelist/review can only contain one stream spec.

The **Files** tab is made up of the following elements (from left to right):

- **Review version selectors:** Select which version of the review you want to diff. For details on using the version selectors, see ["Select review versions to view" on page 454](#).
- **File listing header:**
 - **Review version and common file path summary:**
 - **Comparing the files in the latest version of the review to Base or Head:** Displays the current revisions of the files in the review, the changelist that the review is based on, and the common path for the review files.
 - **Comparing the files in two versions of a review:** Displays which two versions of the review are being compared, which changelists the files are in, and the common path for the files in both versions of the review.
 - **Total number of files:** Indicate the number of files that are marked as read out of the total number of files in a review.
 - **File change summary** icons: Indicate the number of files that have been added, edited, and deleted in this version of the review.



- 
Expand the Diff view to full screen: Select the button to expand the diffs view to full screen. Once the diffs view is expanded to full screen, select the button a second time to exit the full screen mode.
- Diff actions**  button: select the **Diff actions** button to:
 - Expand or collapse all files** (expands or collapses all the files within a review if one or more files are in the collapsed or expanded state)
 - Expand all files** (only displayed if all or at least one file is collapsed) Select to expand all the files within a review.
 - Collapse all files** (only displayed if all files are expanded) Select to collapse all the files within a review.
 - Show or hide all in-line comments for all files in a review** (Shows the total number of comments across all files in a review)
 - Hide in-line comments** (only displayed if inline comments are expanded): Select to collapse all of the inline file comments.
 - Show in-line comments** (only displayed if inline comments are collapsed): Select to expand all of the inline file comments.
 - Show diffs in-line or side-by-side** (only works for text files)
 - Show diffs in-line:** Select to display the diffs in inline format.
 - Show diffs side-by-side:** Select to display the diffs in side-by-side format. The oldest file revision is shown in the left pane, and the newer one is shown in the right pane.
 - Show or hide whitespace and tab characters for all files in a review**
 - Show whitespace and tab characters** (display of CRLF line endings is only supported in the original Swarm review page): Shows whitespace and tab characters for all text files. The tab characters are padded to show the correct alignment of tabs.
 - Hide whitespace and tab characters:** Hides whitespace and tab characters for all text files.
 - Show or hide whitespace diffs**
 - Show whitespace diffs:** Whitespace changes are highlighted, makes it easier to identify changes in file types where whitespace is important.

- **Hide whitespace diffs:** Whitespace changes are not highlighted, this makes it easier to see the important changes in file types where whitespace changes are not important.

- **Navigate to the next file**

Select this option to open the next file in the Diff view or select **Alt+N** on your keyboard.

- **Navigate to the previous file**

Select this option to open the previous file in the Diff view or select **Alt+P** on your keyboard.

- **File list:**

Important

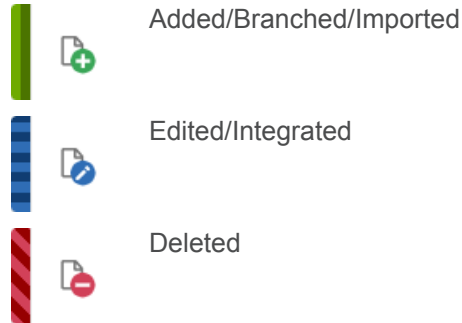
- For a Swarm review containing large number of files, using **Ctrl+F** to search for a file does not work. You can use the file filter to search for a specific file in the review. To manually select a file from the file stack, scroll through the Review page.
- By default, when you open the Review page for the first time, the file list panel and the information panel are hidden. Use the **Show more context** buttons to view the file list panel and the information panel.

For more information about the show more context buttons, see ["File diff panel" on the facing page](#).

- **File filter:** Filter by file name, file extension, or folder name to find the files you are interested in. You can see the diff view for the matched search result in the File stack section.
- **File list resize:** Change the file list width by dragging the split bar left or right. Alternatively, click the split bar to collapse or expand the file list. If you change the file list width for a review, Swarm saves the width locally for that review. The saved width is used next time you view the review.
- **View complete file list:** Use the scroll bar to view the complete list of files.

- **File stack:**

- **File change type** icons: Each file is marked with a ribbon and an icon indicating whether the file was:



- **File content panel header:** The file content panel header contains the following elements from left to right:



- **File name:** The name and revision of the file you are currently viewing.
- **Total number of file comments:** The total number of comments for a file. The comments count includes all inline comments and file-level comments. It does not include the archived comments.
- **Show Full Context** button (only displayed for text file diffs): Click to toggle between displaying only the portions of the file that have changed and the full file.
- **Open in a new tab** button: Click to open the file in a new tab.
- **Mark file as read** button: Click to mark the file as read, click again to mark the file as unread.

File diff panel

By default, all files larger than 1 MB in size are truncated. When a file is truncated a message is displayed. Use the `"max_size"` on page 622 configurable in the `SWARM_ROOT/data/config.php` file to increase or decrease this limit.





- **Text file diff:**

When you view a diff, the changes are highlighted:

- Red indicates lines that have been removed.
- Blue indicates lines that have been modified.
- Green indicates lines that have been added.



For more information on the supported extensions for syntax highlighting in the Review page, see ["Supported syntax highlighting in the Review page" on page 331](#).

Example inline diff view:



 ./db.c#1 4   


42 lines hidden above

43	43	* perferendis consequuntur qui quidem. Et et aut est ex.
44	44	*/
45	45	void * initVoluptatem() {
46	46	int loop = 1;
47	47	
	48	// Abbatia hodie conubia litora nunc.
	49	dolor = flag / 1;
	50	printf("Femina sem bis non dolor.");
	51	printf("Wisi, mando gladius.");

 **Steve Russell** commented on review 644 (version 1) - about a year ago ✓  ...

@allison.clayborne Is this really necessary to spam the console with this?

 Add a comment

	52	
48	53	// Harum quaerat in distinctio vitae adipisci impedit distinctio.
49	54	printf("Dicta voluptas mollitia est placeat aut.\n");
50	55	et = tempora + 1;
51	56	voluptas = atoi(qui)?quas:"eos";
52	57	runOccaecati(ut, sed);

34 lines hidden

87	92	popFugiat(aperiam, quo);
88	93	addBeataeDevice(iste, ut);
89	94	}
90	95	
91	96	
	97	/**
	98	* Herbam class tortor justo, augue justo congue. Volutpat a magnis
	99	* cursus pupula. Aliquet diam peristo ullamcorper cultura. Suspendisse
	100	* nostra.
	101	*/

Example side-by-side diff view:

db.c#1

4

<div> <div>42 lines hidden above</div> <div> <div>43</div> <div>* perferendis consequuntur qui quidem. Et et aut es t ex.</div> <div>44</div> <div>*/</div> <div>45</div> <div>void * initVoluptatem() {</div> <div>46</div> <div>int loop = 1;</div> <div>47</div> </div> </div>	<div> <div>42 lines hidden above</div> <div> <div>43</div> <div>* perferendis consequuntur qui quidem. Et et aut es t ex.</div> <div>44</div> <div>*/</div> <div>45</div> <div>void * initVoluptatem() {</div> <div>46</div> <div>int loop = 1;</div> <div>47</div> <div>48</div> <div>// Abbatia hodie conubia litora nunc.</div> <div>49</div> <div>dolor = flag / 1;</div> <div>50</div> <div>printf("Femina sem bis non dolor.");</div> <div>51</div> <div>printf("Wisi, mando gladius.");</div> </div> </div>
<div> <div>48</div> <div>// Harum quaerat in distinctio vitae adipisci impe dit distinctio.</div> <div>49</div> <div>printf("Dicta voluptas mollitia est placeat au t.\n");</div> <div>50</div> <div>et = tempora + 1;</div> <div>51</div> <div>voluptas = atoi(qui)?quas:"eos";</div> <div>52</div> <div>runOccaecati(ut, sed);</div> </div>	<div> <div>52</div> <div>// Harum quaerat in distinctio vitae adipisci impe dit distinctio.</div> <div>53</div> <div>printf("Dicta voluptas mollitia est placeat au t.\n");</div> <div>54</div> <div>et = tempora + 1;</div> <div>55</div> <div>voluptas = atoi(qui)?quas:"eos";</div> <div>56</div> <div>runOccaecati(ut, sed);</div> <div>57</div> </div>
<div> <div>34 lines hidden</div> <div> <div>87</div> <div>popFugiat(aperiam, quo);</div> <div>88</div> <div>addBeataeDevice(iste, ut);</div> <div>89</div> <div>}</div> <div>90</div> <div></div> <div>91</div> </div> </div>	<div> <div>34 lines hidden</div> <div> <div>92</div> <div>popFugiat(aperiam, quo);</div> <div>93</div> <div>addBeataeDevice(iste, ut);</div> <div>94</div> <div>}</div> <div>95</div> <div></div> <div>96</div> </div> </div>

Steve Russell commented on review 644

(version 1) - about a year ago

@allison.clayborne Is this really necessary to spam the console with this?

Add a comment

Image file diff:

Image types that are natively supported by your browser are displayed with a vertical slider on the image.

Drag the slider to the left to reveal more of the newer image and to the right to reveal more of the earlier one. Alternatively, click on the image to move the slider.

Example image diff:










■ **Show more context buttons:**

Sometimes, the concise diff view needs to be expanded to fully understand the context of the change, use the **Show More** and **Show All** buttons to display extra lines around the change:

Tip

The following buttons are not available for stream specs.

- **Show the file list panel**  button (only displayed on the vertical bar next to the file list panel): click to view the file list panel.
- **Show the information panel**  button (only displayed at the bottom of the information panel): click to view the detailed information panel.
- **Show All Lines to Start of File**  button (only displayed for the first change in the file): click to show all of the lines up to the start of the file.
- **Show More Lines for the Code Below**  button: click to show 10 more lines above the change, the extra lines are displayed in the pane below the button.
- **Show Entire Section**  button: click to show all of the lines between the changes that are above and below the button, the two changes and the lines between them are displayed in a single pane.
- **Show More Lines for the Code Above**  button: click to show 10 more lines below the change, the extra lines are displayed in the pane above the button.
- **Show All Lines to End of File**  button (only displayed for the last change in the file): click to show all of the lines down to the end of the file.

Comments tab

The **Comments** tab is used to view all of the comments in the review.

Files **6**
Comments **4**
Activity
Send All Notifications (0)

5 archived comments

Steve Russell commented on review 644 (version 1) //projects/mercury/dev/src/api/db.c, line 51 - 11 months ago
✓ read ✓

Allison Clayborne commented on review 644 (version 1) //projects/mercury/dev/src/api/db.c, line 105 - 10 months ago
👁️ ...

```
+ */
+int TalisAliquam(long quandoScelerisqueSem, float count, double litora) {
+ // Peristo velit parco.
+ parco = volutpat / 1;
+ printf("Aquam arcu vehicula faucibus umbra magna.");
```

@steve.russell what do you think of this?

👍

Steve Russell commented 4 months ago
👁️ ...

This looks good to me. I am happy for it to be checked in.

👍

Allison Clayborne commented 4 months ago
👁️ ...

Thanks for the confirmation.

👍










Add a comment


Drag and drop files to attach them or [Choose your files](#)

Post
☐ Flag as Task
☒ Post and notify (0)

Comments are made up of the following elements:

- **Send all notifications (n)** (only supported in the original Swarm review page): Comment notifications are delayed by default, click to manually send the notification immediately for the review. For more information about comment notification delay, see ["Comment notification delay" on page 355](#).
- **Comment detail:**
 - **Commenter** avatar and name: The avatar and name of the user that made the comment. Hover over the avatar to see their name and hover over their name to see their ID. Click on the avatar or name to see their user profile. For information about user profiles, see ["Viewing another user's profile" on page 371](#).
 - **Review ID:** The unique number used to identify the Swarm review.
 - **Version** link (includes the file and line number if the comment is on a line in a file): Click to go to the review version the comment was made on. If the comment is made inline in a file the link will take you to that line in the file.
 - **Filepath** link (only if the comment is made on a file): Link to the file the comment is made on, includes line number if the comment was made on line in the file.
 - **Comment raised:** When the comment was made.

- **Tasks:** Flagging review comments as tasks is a lightweight workflow within a review that helps authors and reviewers prioritize review feedback. Any comment on a review can be flagged as a task, indicating to the review's author that the described issue needs to be addressed, and that the review is unlikely to be approved without a fix. For information about working with tasks, see ["Tasks" on page 351](#).
 - **No flag displayed:** The comment has not been flagged a task. To flag the comment as a task, click the **Comment actions**  button and select **Flag as Open task**.
 - **Flagged as a Task**  button: the comment has been flagged as a task. Click to confirm that the task has been addressed, or to remove the task flag from the comment.
 - **Task Addressed**  button: the comment task has been addressed. Click to verify that the task has been addressed correctly, to reopen the task if it has not been addressed correctly, or to verify and archive the task (only supported in the original Swarm review page).
 - **Task Verified**  button: the comment task has been addressed and has been verified as correct. Click to reopen the task you think it has not been addressed correctly.
- **Mark comment as read**  button (unread comments only): Click to mark a comment as read, the comment will only be marked as read for you. For more information on marking a comment as read or unread, see ["Mark comments as read" on page 360](#).
- **Restore**  button (archived root level comments only): Click to restore an archived comment. For more information about archiving and restoring comments, see ["Archiving comments" on page 362](#) and ["Restoring comments" on page 363](#).
- **Comment actions**  button:
 - **Flag as Open task:** Select flag a comment as a task. For information about working with tasks, see ["Tasks" on page 351](#).
 - **Archive:** (only available for root level comments): Select to archive a comment and any replies to that comment. For more information about archiving and restoring comments, see ["Archiving comments" on page 362](#) and ["Restoring comments" on page 363](#).
- **Comment context** (only if the comment is made inline in a file): Displays several lines of code before the line of code the comment is attached to. This helps makes sense of the comments should later changes remove those lines.
- **Comment** content: This can be text, a URL, or an attachment (only supported in the original Swarm review page). For more information about the content of comments, see ["Comment features" on page 355](#).
- **Reply**  button: Click to reply to the comment. For information about replying to comments, see ["Reply to comments" on page 357](#).
- **Edit**  button (only available for comments that you have made): Click to edit the comment. For information about editing comments, see ["Editing comments" on page 350](#).

- **Thumbs Up**  button: Click to like a comment. For information about liking comments, see ["Reacting to comments" on page 359](#).

Add a comment

To add a comment, type your comment in the open comment text box at the bottom of the **Comments** tab page.



To reply to a comment, click the **Reply**  button and type your comment in the comment text box.

- **Comment** text box: type your comment in the text box.
- **Flag as a Task** checkbox: Select to flag the comment as a task. For more information about tasks, see ["Flag a comment as a task" on page 351](#).
- **Post** button: Click to post your comment. The comment is posted immediately but the comment notification is delayed, see ["Comment notification delay" on page 355](#).
- **Post and notify (n)** link: Click to post your comment and send the comment notification immediately.
Where **(n)** is the number of delayed comment notifications in the queue waiting to be sent, this number does not include the current comment you are working on.


Activity tab



The **Activity** tab presents a list of the events on this review.

Files **6**
Comments **4**
Activity
Send All Notifications (1)




Allison Clayborne cleared an issue on [review 644](#) for [mercury:dev](#) 6 minutes ago



This needs another review please






Allison Clayborne archived comment on [review 644](#) for [mercury:dev](#) 8 minutes ago



This description needs more detail.






Allison Clayborne replied to a comment on [review 644](#) for [mercury:dev](#) 8 minutes ago

Thanks for the confirmation.





Steve Russell replied to a comment on [review 644](#) for [mercury:dev](#) 9 minutes ago

This looks good to me. I am happy for it to be checked in.



Events in the activity tab include:

- When the review was started
- When a new reviewer joins the review
- When the review's state changes
- When the review's files are updated
- When a reviewer votes on the review

- When someone comments on the review, or one of its files
- When tests pass or fail, provided continuous integration is configured

Swarm user tasks

This section describes a number of the common Swarm user tasks:

Start a code review

Important

If your Helix Core Server is configured as a commit-edge deployment, and your normal connection is to an edge server, Swarm refuses to start reviews for shelved changes that have not been promoted to the commit server.

Within Swarm, this means that the **Request Review** button does not appear for unpromoted shelved changes. Outside of Swarm, attempts to start reviews for unpromoted shelved changelists appear to do nothing. Ask your Helix Core Server administrator for assistance if you cannot start a review.

An administrator of the Helix Core Server can automatically promote shelved changes to the commit server by setting the configurable `dm.shelve.promote` to 1.

Tip

If your changelist only contains a stream spec and its location in the Helix Core Server is not associated with a Swarm project, the review that you create will not have any default reviewers or workflow rules. To associate the review with a project so that it has default reviewers and obeys the project workflow rules, include a file change from the project path in the changelist when you create the review.

1. Use the Helix Core command-line (P4) or a client to create the shelved or committed changelist.
2. Start a code review by using one of the following approaches:

Use Swarm:

1. Use Swarm to view a shelved or submitted changelist.

Tip

To view a shelved or submitted changelist, use a [Quick URL](#). For example, if your change is **54321**, visit the URL: `https://myswarm.url/54321`.

2. Click the **Request Review** button to request a review of that changelist.

Requesting a review on a shelved changelist uses the [pre-commit](#) model and requesting a review on a submitted changelist uses the [post-commit](#) model.

Use Helix Core command-line (P4):

When you are about to shelve or submit files:

1. Include `#review` within your changelist (separated from other text with whitespace, or on a separate line).

Once the review begins, Swarm replaces `#review` with `#review-12345`, where `12345` is the review's identifier.

Note

The `#review` keyword is customizable. For details, see ["Review keyword" on page 680](#).

2. At this time, you can add reviewers to the code review by using `@mention` for users, and `@@mention` for groups in the changelist description for each desired reviewer.

If your `@mention` or `@@mention` includes an asterisk (*) before the *userid* or *groupid*, for example `@*userid`, that user or all of the group members become required reviewers. If your `@@mention` includes an exclamation mark (!) before the *groupid*, for example `@@!groupid`, the members of that group become required reviewers but only one member of the group is required to vote. See ["Required reviewers" on page 478](#) for details.

3. Complete your shelve or submit operation.

Warning

If you shelve a changelist and subsequently edit the description to include `#review`, a review is **not started**. You must re-shelve the files after adding `#review`.

Tip

You can also start a Swarm review with P4V, P4VS, and P4Eclipse. See below for details:

- **P4V:** see the [Swarm integration features](#) section of the [Helix Core Visual Client \(P4V\) Guide](#).
- **P4VS:** see the [Managing files](#) chapter of the [P4VS User Guide](#).
- **P4Eclipse:** see the [Reviewing changes](#) chapter of the [P4Eclipse User Guide](#).


Note

If you are using P4V and its Swarm integration, and you encounter the error **Host requires authentication**, ask your Helix Core Server administrator for assistance. See ["P4V Authentication" on page 588](#) for details.

Once a review has started

Wait for someone else to review your code, or see ["Contribute comments or code changes to a code review" on page 74](#) More [review activities](#) are available.

Edit the reviewers on a review

A review author can edit the reviewers for a review by using the Reviewers **Edit**  button in the review page **Information panel**. Reviewers are able to join or leave reviews, or to change whether their vote is required or optional.

Note


If the review includes content that is part of a project or branch with [Retain default reviewers](#) enabled, the following restrictions apply to the review:

- The voting option for a retained default reviewer can only be changed to a stricter level, you cannot reduce the voting level.
- Retained default reviewers cannot be removed from the review.

Additionally, the following individuals can edit reviewers:

- Users with *admin* or *super* privileges.
- If the review is [moderated](#), the moderators.
- If the review is part of a project, but not moderated, all project members.
- If the review is not part of a project, any authenticated user.

To edit reviewers on a review:

1. Go to the review.
2. From the review **Information panel**, click the **Reviewers** list **Edit**  button.
The **Manage reviewers** dialog is displayed.
3. Click the **Users** tab to edit a user or the **Groups** tab edit a group. This field auto-suggests users, and groups within Helix Core Server as you type (up to a combined limit of 25 entries).
4. **Change the voting requirements of a user or group:**
 - a. Click the dropdown arrow to the right of the user or group you want to change.
 - b. Select one of the following:

- **Users:** select whether their vote is **Required** or **Optional**. A solid star means that their vote is required to approve a review, whereas the outlined star means that their vote is optional.
 - **Groups:** select whether the group vote is **Require all**, **Require one**, or **Optional**. A solid star means that all group member votes are required to approve a review, a solid star with a 1 inside means at least one group member must vote up and no group members vote down to approve a review, and the outlined star means that the group vote is optional.
5. **Add a user or group to a review:**
 - a. Click the **Add user** or **Add group** dropdown icon to display a list of users or groups:
 - b. Start typing the user or group name, the list is filtered as you type.
 - c. Click on the user or group to add to the review.
 - d. Click **+** to add them as a reviewer.
 - e. Set the reviewer type from the dropdown to the right of the user or group.
 6. **Remove a user or group from a review:**
 - a. Click the dropdown icon to the right of the user or group to be removed.
 - b. Select **Remove from the review** from the dropdown menu.
 7. To close the **Manage reviewers** dialog, click **X**.

Check for code reviews you need to act on

The purpose of the dashboard is to allow you to focus on reviews that need to be done, so that other users are not blocked. The dashboard lists the most recently modified reviews first, and shows your role in the review.

Your dashboard is available by clicking **My Dashboard** in the menu or by clicking the Swarm  icon above the main menu.

Note






















Since it is tied to the logged in user, **My Dashboard** is only populated if you are logged in.


1. Log in to Swarm.
2. If the **My Dashboard** page is not already displayed, click **My Dashboard** in the menu or click the Swarm  icon above the menu.

My Dashboard

Project Role Author

There are new changes to your review list [Refresh](#) [Clear all](#)

Description	Your Role	Last Activity	State	Tests	Complexity	Votes
 Updated README to have a bit more information. Jack Boone requested a  pre-commit review 365 for mercury:dev	Reviewer	44 minutes ago			 137	 0  0
 Updated notes for the new year. Jack Boone requested a  post-commit review 422 for jplugin:main	Required reviewer	2 months ago			 18	 0  0
 Add comment to method. Steve Russell requested a  post-commit review 328 for jplugin:main	Required reviewer	3 months ago			 6	 2  0

 All available reviews are displayed

- By default all projects with reviews you need to act on are displayed.
- To filter the reviews displayed by project, select **My Project** or specific projects from the **Project** filter. To select multiple projects, click the projects you want to filter by. The **Project** filter will only show projects for which there are reviews in your dashboard.
- By default all reviews you need to act on are displayed.
- To filter by your role in a review, select **Author**, **Reviewer**, **Required reviewer**, or **Moderator** from the **Role** filter. The **Role** filter will only show roles for which there are reviews in your dashboard.
- By default all authors with reviews you need to act on are displayed.
- To filter by review author, select **All Authors** or a specific author from the **Author** filter. The **Author** filter will only show authors for which there are reviews in your dashboard.
- The **Refresh** button is displayed when new content is available for your dashboard. Click **Refresh** to update your dashboard with the new content.

Note


By default, when your dashboard is open, Swarm checks for new content every five minutes.

Tip

- For more information on the dashboard page, see ["My Dashboard" on page 306](#).
- Click on a review **ID** number to display the [code review](#).

List reviews you are involved with

1. Click **Reviews** in the main menu.



2. By default all projects are displayed.
3. To filter the reviews displayed by project, select **My projects** or specific projects from the **Project** drop down. To select multiple projects, click the projects you want to filter by.
4. By default all reviews are displayed.
5. To filter by your role in a review, select **Reviews I've authored**, **Reviews I'm participating in**, **Reviews I'm an individual reviewer of**, **Reviews I've authored or I'm participating in**, or search for reviews authored by a specific user from the **Role** drop down.
6. Click the **Opened** tab to view open reviews, click the **Closed** tab to view closed reviews. See ["Reviews list" on page 428](#) for more information.

Contribute comments or code changes to a code review

You can contribute to a code review in the following ways:

- Comment on a review
- Comment on a review description
- Comment on a specific line in a file
- Comment on a file in a review
- Reply to a comment
- Edit files in a review

Comments

Comments are the primary feedback mechanism provided by Swarm. Review comments can be flagged as *tasks* for a lightweight workflow within a review that helps authors and reviewers prioritize review feedback. By default, comment notifications are delayed to allow you to add or edit comments as you progress through a review without sending a notification for each individual comment on the review. Comment notifications are rolled up into a single notification that you can either leave to be sent automatically, or you can send manually, see "[Comment notification delay](#)" on page 355.

Tip

For more information about comments and comment features, see "[Comments](#)" on page 348.

Comment on a review

1. From the review page: click **Comments** to view the **Comments** tab.
2. Add your comment in the text area.
3. **Optional:** select the **Flag as Task** checkbox to mark the comment as a task that needs to be addressed.
4. Click **Post**.

Comment on a review description

1. From the review page description area: click **Comments (n)** (where **n** is the number of comments that already exist).
2. Click **Add a comment**.
3. Add your comment in the text area.
4. **Optional:** select the **Flag as Task** checkbox to mark the comment as a task that needs to be addressed.
5. Click **Post**.

Tip

To hide the description comments, click **Comments n** (where **n** is the number of description comments that exist). To display the comments again, click **Comments n**.

Comment on a specific line in a file


1. From the review page: click **Files** to view the **Files** tab.
2. Click on the line you want to comment on.
3. Add your comment in the text area.

4. **Optional:** select the **Flag as Task** checkbox to mark the comment as a task that needs to be addressed.
5. Click **Post**.

Comment on a file in a review

1. From the review page: click **Files** to view the **Files** tab.
2. If there are multiple files, click the file you want to comment on to expand its view.
3. Click the **Add a Comment** link in the footer of the file display.
4. Add your comment in the text area.
5. **Optional:** select the **Flag as Task** checkbox to mark the comment as a task that needs to be addressed.
6. Click **Post**.

Reply to a comment

1. From the review page: click **Comments** to view the **Comments** tab.
2. Navigate to the comment you are replying to.
3. Click **Reply**  below the comment you are replying to.
4. Add your comment in the text area.
5. **Optional:** select the **Flag as Task** checkbox to mark your reply as a task that needs to be addressed.
6. Click **Post**.

Editing review files

Edit files in a review

1. [Get a local copy](#) of the review's files.
2. Edit the files as required.
3. Prepare a changelist with the edited files and include `#review-1234` within the changelist's description (separated from other text with whitespace, or on a separate line), where `1234` is the review's identifier.

Warning

If you use an invalid review identifier, it will appear that nothing happens. Swarm is currently unable to notify you of this situation. If the review has not been correctly updated, use the **Add Change** button in the Swarm review heading to add the changelist to the review, see ["Add a changelist to a review" on page 471](#).

4. Depending on the [model of code review](#) you are using, you would:

- Shelf the files (for pre-commit).
- Submit the files (for post-commit).

Get a local copy of the code in a review for evaluation

Swarm manages one or more changelists containing shelved copies of all of the files belonging to a specific review. You can unshelve the files to receive a copy of the review's code, or you can click the [Download .zip](#) button on the review to download a ZIP archive containing all of the review's files.

Determine the changelist that contains the files in the review

1. Visit the review's page.
2. The current review version's changelist appears in the file list heading.

#3: Change 697707 shelved into //depot/main/swarm

In this example, the changelist is **697707**. You use the identified changelist in place of *shelved changelist* below.

Note

Swarm can version file updates in reviews. For more information, see ["Review display" on page 433](#).

Using P4

For a shelved changelist, use a command-line shell and type:

```
$ p4 unshelve -s shelved changelist
```

For a committed changelist, use a command-line shell and type:

```
$ p4 sync @committed changelist
```

Note

Your client's view mappings need to include the changelist's path.

Using P4V

For a shelved changelist:

1. Select **Search > Go To**.
2. Change the select box to **Pending Changelist**.
3. Type in the shelved changelist number and click **OK**.
4. Select the files in the **Shelved Files** area.
5. Right-click and select **Unshelve**.
6. Click **Unshelve**.

For a committed changelist:

1. Select **Search > Go To**.
2. Change the select box to **Submitted Changelist**.
3. Type in the submitted changelist number and click **OK**.
4. Select the files in the **Files** area.
5. Right-click and select **Get this Revision**.
6. Click **Close**.

Download files as a ZIP archive

The **Download zip** option is used to download a ZIP archive that contains all of the files in the review. The file revisions of the downloaded files are the file revisions in the most recent review version selected in the review version selector, see ["Select review versions to view" on page 454](#).

Note

The **Download zip** option is not displayed if the zip command-line tool is not installed on the Swarm server. For information about installing, and configuring the zip command-line tool, see [Zip archive](#).

When you select the **Download zip** option, Swarm performs the following steps:

1. Scans the files/folders:
 - Checks that you have permission to access their contents, according to the Helix Core Server protections.
 - Checks that the total file size is small enough to be processed by Swarm.
2. Syncs the file contents to the Swarm server from the Helix Core Server.
3. Creates the ZIP archive by compressing the file content.
4. Starts a download of the generated ZIP archive.

Note

- You might not see all of the above steps; Swarm caches the resulting ZIP archives so that repeated requests to the same files/folders can skip the sync and compress steps whenever possible.
- If an error occurs while scanning, syncing, or compressing, Swarm indicates the error.

Fix errors that cannot be merged in a review

The problem can occur when you attempt to **Commit** or **Approve and Commit** via the Swarm UI and the shelved files are out of date.

Helix Swarm cannot currently help with resolving conflicts; you need to use a Helix Core Server client such as p4 or P4V to resolve conflicts.

Resolve via P4

1. Acquire a [local copy](#) of the files.
2. Sync the files to the head version:

```
$ p4 sync
```

3. Begin resolving files with:

```
$ p4 resolve
```

4. Choose an appropriate option to resolve each file. For example:

```
$ p4 resolve
/home/bruno/bruno_ws/dev/main/jam/command.c - merging
//depot/dev/main/jam/command.c#9
Diff chunks: 4 yours + 2 theirs + 1 both + 1 conflicting
Accept(a) Edit(e) Diff(d) Merge (m) Skip(s) Help(?) e:
```

5. Re-shelve the resolved files with:

```
$ p4 shelve
```

Note

Ensure that the changelist description contains **#review-12345** (separated from other text by whitespace, or on a separate line), where **12345** is the identifier of the review you are updating.

Warning

If you use an invalid review identifier, it will appear that nothing happens. Swarm is currently unable to notify you of this situation. If the review has not been correctly updated, use the **Add Change** button in the Swarm review heading to add the changelist to the review, see ["Add a changelist to a review" on page 471](#).

For more information, see "Resolve" in the *Helix Core Command-Line (P4) Guide*.

Resolve via P4V

1. Acquire a [local copy](#) of the files.
2. In P4V, right-click your workspace folder and select **Resolve Files**.

The **Resolve** dialog appears.

3. Choose the appropriate options to resolve each file.
4. Right-click your workspace folder and select **Shelve Files**.

The **Shelve** dialog appears.

Note

Ensure that the changelist description contains **#review-12345** (separated from other text with whitespace, or on a separate line), where **12345** is the identifier of the review you are updating.

Warning

If you use an invalid review identifier, it will appear that nothing happens. Swarm is currently unable to notify you of this situation. If the review has not been correctly updated, use the **Add Change** button in the Swarm review heading to add the changelist to the review, see ["Add a changelist to a review" on page 471](#).

5. Click **Shelve**.

For more information, see "Resolving Files" in the *Helix Core Visual Client (P4V) Guide*.

Change the author of a review


Note


By default you cannot change the author of a review, this option must be enabled by your Swarm administrator. See ["Allow author change" on page 687](#) for details.

If the author of a review is no longer available, or ownership of a review is passed to a different developer, it is useful to be able to change the author of that review.

To change the author of a review:


1. Click **Reviews** in the main menu.
2. Click the **Role** drop down and search for reviews by the author you want to replace. The list will auto-complete with users as you type.



3. Select the author from the list, reviews by that author are displayed.
4. Click the review you want to change the author on.
5. Click the **Review actions**  button.
6. Select **Change author**.
7. Select the new review author from the user list.
8. Click **Save** to change author of the review.

Change your user notification settings

1. Click on your userid in the header and select **My Profile**.



2. Click the **Notifications** tab to display your user notifications settings.
3. Configure which notifications you receive when events occur within Swarm, see the ["Notifications tab" on page 369](#).


Unfollow all projects and users you are following

When you are viewing your own user profile, you can unfollow all projects and users that you are following.

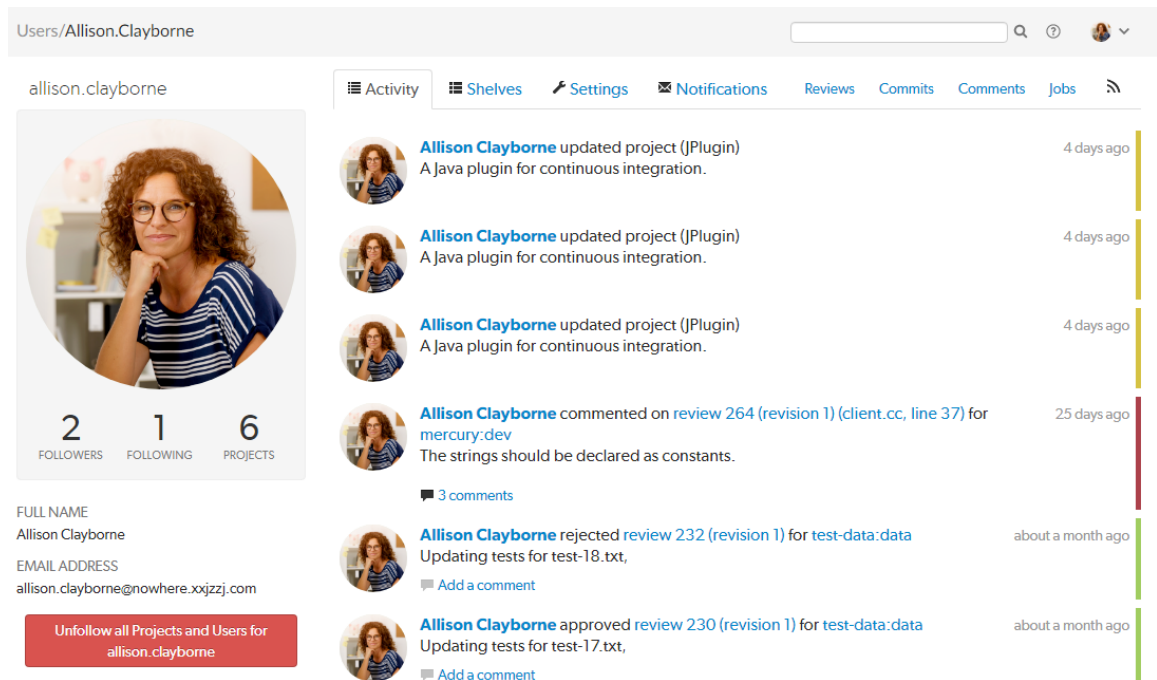
Note

This action cannot be undone.

1. Display your own user profile when you are logged into Swarm by clicking on your *userid* in the header and selecting **My Profile**.



Your user page is displayed:



2. Click the **Unfollow all projects and users for *your-username*** button.
3. Click **OK** when the confirmation dialog is displayed to complete the unfollow action.
4. You will no longer be following any projects or users.

Swarm administration tasks

This section describes a number of the common Swarm administration tasks:

Change notification settings for a group

Important

You must be an owner of the group, have *super* privileges in Helix Core Server (**p4d**), or have *admin* privileges in **p4d** version 2012.1 or later, to edit group notifications. If you do not have sufficient permissions, Swarm does not display the **Notifications** tab for the group.

1. Click **Groups** in the main menu.
2. Click the group you want to edit.
3. Click the group **Notifications** tab.
 - If **Group mailing list** is **Disabled**, emails are sent to the group members individual email addresses and only the following notification options are available:
 - **Email members when a review is requested**: When any member of this group creates a review, this group will be notified.
 - **Email members when a change is committed**: When a change is committed into Perforce, if the owner of the changelist is a member of this group, this group will be notified.

Tip

When a user commits a changelist in Swarm, it is committed on behalf of the changelist owner. If the changelist owner is a member of this group, this group will be notified.

- If **Group mailing list** is **Enabled** notifications are sent to the group email address and additional notifications are available for the group, see [group notifications](#) for details.

When a group mailing list is enabled it overrides the group member's preference set for the `notify_self` configurable in the `SWARM_ROOT/data/config.php` file. So even when the `notify_self` configurable is set to false the group member will receive an email notification.

There is a caveat that if a user is a group member and the same user is added individually to the review and has set the `notify_self` configurable to true, the user receives two email notifications.

For more information about `notify_self` configurable, see ["notify_self" on page 605](#).


Manage project branches

Initial steps:

1. Visit your project page.
2. Click **Settings** in the menu.

A project summary with Participants tab, Branches tab, Integrations tab, and General Settings tab is displayed.

3. Click **Branches** tab and expand the **Add new branch** pane.



The screenshot shows the 'Add new branch' form within the 'Branches' tab. The form has two input fields: 'Name *' and 'Depot paths *'. Below the 'Name *' field is a message: 'A name is required for a branch'. Below the 'Depot paths *' field is a message: 'Branch paths are required for a branch'. At the bottom of the form is a button labeled 'Add new branch'.

Adding a branch

1. Follow the [initial steps](#).
2. Enter a short **Name** for your branch.
3. Enter one or more branch paths in the **Depot paths** field using depot syntax, one path per line.

Tip

- Branch paths, and files can be excluded by putting a minus symbol – at the start of the path.
- Branch paths are processed in order, starting with the first file path in the list.

For example:

```
//depot/main/swarm/...  
-//depot/main/swarm/test/...  
//depot/main/swarm/test/ResultSummary.html
```


The first path includes all of the directories and files under `//depot/main/swarm/` in the project branch.

The second path excludes all of the files in `-//depot/main/swarm/test/` from the project branch.

The third path includes the `ResultSummary.html` file from the previously excluded `//depot/main/swarm/test/` directory.

Note

- Wildcards should not be used; the only exception is that the branch path can end with the Helix Core Server wildcard `...`
- Branch paths are case sensitive.
- Branch paths, and files are not checked to see if they are valid when you save the branch:
 - If you enter an invalid path ending in the wildcard `...`, the path will not be displayed in the [project file browser](#) until the path is created. This allows you to specify a path before it has been created.
 - If you enter a path that ends with a file that has not been committed, or a non-existent file, Swarm displays a 404 error when you navigate the path to the file with the [project file browser](#).

Note

The [Project Commits](#) tab can fail to show some Helix Core Server commits in the top level view. Individual branch views display the commits correctly:

The **Project Commits** tab top level client view is made up of all of the branches of the project.

For example, Project Alpha:

Branch: QA:

`//depot/alpha/dev/QA/...`

Branch: Dev :

`//depot/alpha/dev/...`

`-//depot/alpha/dev/QA/...`

The project commits tab view is generated by processing the branches in the order that they were created in and from top to bottom for the paths in each of those branches.

For the project Alpha example above:

The first path includes all of the directories and files under `//depot/alpha/dev/QA/` in the project commits tab top level view.
 The second path includes all of the directories and files under `//depot/alpha/dev/` in the project commits tab top level view.
 The third path excludes all of the directories and files in `//depot/alpha/dev/QA/` from the project commits tab top level view.

Result: commits made to `//depot/alpha/dev/QA/` that should be shown for the **QA** branch are not displayed in the **Project Commits** tab top level view.

When you have a simple branch structure this can be avoided by considering this issue when you create your branches. In the example above, creating the **Dev** branch first and then creating the **QA** branch avoids the problem because `//depot/alpha/dev/QA/` is not excluded from the **Project Commits** tab top level view.

A new branch is created. Expand the branch pane to view and further edit the following branch settings:

4. **Optional** (not displayed if the Workflow feature is disabled): associate a workflow with the project branch.

To use the parent project workflow, select **Inherit from project** from the **Workflow** dropdown list. This is the default when you create a new branch.

To associate a workflow, select the workflow from the **Workflow** dropdown list, or enter the workflow name in the search field. The field auto-suggests workflows as you type. Only workflows that you own and shared workflows are shown in the dropdown list.

Tip

- When a workflow is associated with a project, the workflow is used for all of the branches in that project.
- When a project branch is associated with a workflow, the workflow of the parent project is ignored and the branch workflow is used.

For more information about workflows and how project workflows interact with branch workflows, see ["Workflow basics" on page 529](#).

5. **Optional:** set the **Minimum up votes** required for reviews associated with this branch.

To inherit the parent project setting, leave **Minimum up votes** set to **Inherit from project**. This is the default when you create a new branch.

To use the branch setting and ignore the parent project setting, set **Minimum up votes** to **1** or more on the branch.

A review cannot be approved until all of the [Required reviewers](#) have voted up the review and the **Minimum up votes** specified has been satisfied.

- If a review spans projects/branches, the Minimum up votes for **each** of the projects and branches must be satisfied before you can approve the review.
- Required reviewers are included when up votes are counted.
- When **Count votes up from** is set to **Members** for a workflow associated with a project/branch, only the up votes of members of the project contribute to satisfying the **Minimum up votes** for a project/branch. For more information about the **Count votes up from** rule, see ["Workflow rules" on page 526](#).

Important

If the Workflow feature is disabled, all votes are counted not just votes from project members.

Important

If the **Minimum up votes** required is set higher than the number of reviewers that exist for a review, approval will be blocked for that review. This is true even if all the reviewers on the review have voted up the review.

6. **Optional:** Only Moderators can approve or reject reviews. Click **Manage moderators** button to add a user or group to the moderators list.

Select the **Users** tab or **Groups** tab and type in the name. The field auto-suggests groups and users within the Helix Core Server as you type.

If a group is specified as a moderator, all of the members of that group have the same moderator privileges for that project branch as if they were added individually.

Once the branch specification is complete and the project has been saved, changing the state of any review associated with this moderated branch is restricted as follows:

- Only moderators can approve or reject the review. Moderators can also transition a review to any other state.

Important

Moderators prevent the automatic approval of reviews, for more information about automatically approving reviews using workflow rules see ["Workflow rules" on page 526](#).

Note

By default, when a review spans multiple branches that have different moderators, only one moderator from any one of the branches needs to approve the review.

Swarm can be configured to require that one moderator from each branch must approve the review, this is a global setting. If a moderator belongs to more than one of the branches spanned by the review, their approval will count for each of the branches they belong to. For instructions on how to configure moderator behavior, see ["Moderator behavior when a review spans multiple branches" on page 692](#).

- The review's author, when not a moderator, can change the review's state to **Needs review**, **Needs revision**, **Archived**, and can attach committed changelists.
Normally, the review's author cannot change the review's state to **Approved** or **Rejected** on moderated branches. However, authors that are also moderators have moderator privileges, and may approve or reject their own review.
When `disable_self_approve` is enabled, authors who are moderators (or even users with admin privileges) cannot approve their own reviews.
- Project members can change the review's state to **Needs review** or **Needs revision**, and can attach committed changelists. Project members cannot change the review's state to **Approved**, **Rejected**, or **Archived**.
- Users that are not project members, moderators, or the review's author cannot transition the review's state.
- For the review's author and project members, if a review is not in one of their permitted states, for example if the review's state is **Rejected**, they cannot transition the review to another state.

These restrictions have no effect on who can start a review.

7. Default reviewers:

- a. **Optional:** specify **Default Reviewers** for the project branch.

This field auto-suggests users, and groups within Helix Core Server as you type (up to a combined limit of 25 entries). Click on the user or group to add them as a default reviewer. Each time a new review is created, the default reviewers will be added to the review.

- **Users:** click the star icon to the left of the *userid* to toggle whether their vote is required or not. A solid star means that their vote is [required](#) to approve a review, whereas the outlined star means that their vote is optional.
- **Groups:** click the star icon to the left of the *groupid*, and select whether the group is a [required reviewer \(one vote\)](#), a [required reviewer \(all votes\)](#), or an optional reviewer. A solid star means that all group member votes are required to approve a review, a solid star with a 1 inside means at least one group member must vote up and no group members vote down to approve a review, and the outlined star means that the group vote is optional.

Click the **X** icon to the right of the *userid* or *groupid* to remove that default reviewer from the default reviewers list.

Important

When a review is part of multiple projects/project branches:

- The default reviewer lists for all of the projects and project branches the review is part of are combined and added to the review.
- If a default reviewer has different reviewer options set on projects and project branches that the review is part of, the strictest reviewer option is used for the review.

Example: A review is created and it is part of **Project A**, **Project B**, and **Project Branch b**.

Project A: default reviewer X is an Optional reviewer

Project B: default reviewer X is an Optional reviewer

Project Branch b: default reviewer X is a Required reviewer

Result: default reviewer X is added to the review as a Required reviewer

Note

If users or groups are [@mentioned](#) in a new changelist description that includes **#review**, they will be added to the review as reviewers. If any of these reviewers are already specified as default reviewers they will not be added to the review again, the reviewer's most restrictive reviewer option is used for the review.

Note

If a default reviewer is deleted from Helix Core Server they will not be added to new reviews.

- Optional:** click the **Retain default reviewers for reviews associated with this branch** checkbox to prevent default reviewers being removed from reviews associated with this branch.

For more information about retained default reviewers, see [Retain default reviewers](#).

Note

The project name does not need to be included in the branch name; Swarm displays the project name with the branch name when appropriate.

Editing a branch

1. Follow the [initial steps](#).
2. Expand the branch pane you want to edit.
3. Revise the **Name**, **Depot paths**, **Branch moderators**, or **Default reviewers** as required.
4. Click the green tick icon to save your branch specifications.

Removing a branch

1. Follow the [initial steps](#).
2. Expand the branch pane you want to delete.
3. Click **Delete Branch**.
A confirmation dialog appears.
4. Click **Yes** to delete the branch.

Unfollow all projects and users for another user

Note


You must be logged in as a user with *admin* or *super* user privileges and viewing the user's profile page to perform this action.

When a user has been removed from the Helix Core Server but they are still following projects and users, it is useful to be able to remove all of their follows. This helps to keep the project and user follower lists up to date.

Note

This action cannot be undone.

1. Use the **Search** box in the header to search for the user, the search field will auto-complete as you type.



2. Select the user from the search results to display the user's profile:

Users/Alex.Randolph

alex.randolph

Activity Shelves Settings Notifications

Reviews Commits Comments Jobs

Follow

0 FOLLOWERS 1 FOLLOWING 4 PROJECTS

FULL NAME
Alex Randolph

EMAIL ADDRESS
alex.randolph@nowhere.xxjzzj.com

Unfollow all Projects and Users for alex.randolph

Alex Randolph committed change 295 into mercury:main
A few minor changes.
Add a comment 4 months ago

Alex Randolph approved review 177 (revision 1)
Adding basics.notifications to documentation.
#review @@!docs
Add a comment about a year ago

Alex Randolph approved review 165 (revision 1)
Adding basics.files to documentation.
#review @@!docs
1 comment about a year ago

Alex Randolph approved review 131 (revision 1)
Adding admin.search to documentation.
#review @@!docs
2 comments about a year ago

Alex Randolph approved review 129 (revision 1)
Adding admin.reviews to documentation.
#review @@!docs
3 comments about a year ago

Alex Randolph approved review 179 (revision 1)
Adding basics.projects to documentation.
#review @@!docs
2 comments about a year ago

3. Click the **Unfollow all projects and users for *username*** button located below the user's email address.
4. Click **OK** when the confirmation dialog is displayed to complete the unfollow action.
5. The user will no longer be following any projects or users.

Obliterate a review

Note

- By default, you must be a user with *admin* or *super* user rights to obliterate a review.
- Optional:** Swarm can be configured to allow users to obliterate reviews that they have authored. This can be configured by your Swarm administrator. See ["Allow author obliterate review" on page 688](#).


Obliterate is used to permanently delete reviews that have been created by mistake. For instance, if a review is associated with the wrong changelist, or a review contains sensitive information that should not be openly available.

For information on what happens to a review when it is obliterated, see ["When you obliterate a review" on page 669](#).

Important

Obliterate must be used with care, the review and all of its associated metadata are permanently deleted. An obliterated review cannot be reinstated, not even by Perforce Support.

To obliterate a review:

1. Navigate to the review.
2. Click the **Review actions**  button and select **Obliterate Review**.
3. Click **Yes** on the confirmation dialog to complete the obliterate action.
4. The review is obliterated.

Change the logging level

Swarm logs various activities to the data/log file. Change the logging level to increase or decrease the volume of log data by editing a configuration file.

An example configuration, in the `SWARM_ROOT/data/config.php` file:

```
<?php
    // this block should be a peer of 'p4'
    'log' => array(
        'priority' => 3, // 7 for max, defaults to 3
    ),
```

The maximum value for the log priority is 7; higher values do not result in increased logging. The minimum value is 0, which means no logging; lower values do not result in further logging reductions. For more information, see ["Logging" on page 631](#).

Check the queue workers

Note

- Swarm package installations are automatically configured with a cron job to spawn workers when you run the `configure-swarm.sh` post-installation configuration script during installation.
- Swarm tarball installations should be setup with a cron job to ensure that workers are running to process events, this is part of the normal installation process. For information on setting up a cron job to spawn workers, see ["Set up a recurring task to spawn workers" on page 220](#).

Helix Swarm uses a custom queue system to process events, provide notifications, and more. The queue system is required to handle the potentially large volume of events from a busy Helix Core Server. A queue manager ensures that sufficient queue workers are available to process items.

You can check your queue workers in the following ways:

- If you do not have *admin* or *super* user permissions, see ["HTTP request" on the facing page](#)
- If you have *admin* or *super* user permissions, see ["System information page" on the facing page](#)

HTTP request

Check the status of the queue by making an HTTP request to `/queue/status`. The response is formatted in JSON and looks like this:

```
{"tasks":0,"futureTasks":1,"workers":3,"maxWorkers":3,"workerLifetime": "595s"}
```

This response indicates that the queue has no current tasks, there is 1 task scheduled for processing later, there are 3 queue workers available, at most 3 workers are created, and queue workers run for at most 10 minutes before self-terminating.

If the queue manager has stopped for some reason, start a new one by making an HTTP request to `/queue/worker`. No response is provided for this request.

System information page

The **Queue info** tab on the **System Information** page displays Swarm basic queue worker information. As well as displaying worker information you can also display the task queue, start a worker, and start a temporary worker.

To navigate to the **Queue Info** tab:

1. Click your **User id** avatar, to the right of the Swarm header.
2. Click **About Swarm**.
3. Click **System Information**.
4. Click the **Queue Info** tab.

System Information

Perforce
Log
PHP Info
Queue Info
Cache Info

Tasks	0
Future Tasks	4
Workers	1
Max Workers	3
Worker Lifetime	595s
Ping Error	false

Start a Worker
Refresh Tab
Show Task Queue
Start a Temp Worker

- For information about the **Queue Info** tab, see "Queue Info" on page 727.

Automatically deploy code within a review

Important

The project level test and deploy code features will be deprecated in a later Swarm release. We recommend you use test integration to automatically deploy code within a review. For more information, see ["Add a test" on page 544](#).

Deploying code in a code review automatically involves enabling **Automated Deployment** in your project's configuration and providing a *trigger URL*. When the *trigger URL* is requested, Swarm expects a deployment program to be executed.

When the deployment processing ends, Swarm expects either a *success callback URL* or *failure callback URL* to be requested by your deployment program. These callback URLs should include a url parameter (either via GET or POST); when a valid-looking URL is included, clicking the deployment status indicator directs the user to the specified URL. This is intended to facilitate easy viewing of the successfully deployed review, or a report indicating why the deployment failed. The url parameter is mandatory for successful deployments, but is optional for failures.

- Navigate to the project page.
- Click the project **Settings** tab to display the **Project Settings** page.

3. **Automated Deployment** checkbox: select **Enable** to display the configuration field:

Automated Deployment ☒ **Enable**

`http://deploy-server/deploy?change={change}`

A URL that will trigger a deployment when reviews are created or updated.
Some special [arguments](#) are supported. [See help for more details.](#)

4. Provide a URL that triggers your deployment execution.

Special arguments are available to inform your deployment program of various details from Swarm:

Note

Helix Plugin for Jenkins 1.10.11 and later: Swarm must send the parameters for the build to Jenkins as a POST request. To do this, enter the parameters in the **Post Body** and select **URL Encoded**.

{change}
The change number

{status}
Status of the change, *shelved* or *submitted*

{review}
The review's identifier

{project}
The project's identifier

{projectName}
The project's name

{branch}
The branch identifier(s), comma-separated

{branchName}
The branch name(s), comma-separated

{success}
Deployment successful callback URL

{fail}
Deployment failure callback URL

3 | Install and upgrade Swarm

This chapter covers the initial installation and configuration of Swarm as well as upgrading an existing Swarm installation.

Important restrictions

Do not prefix group names, project names, user names, or client-names with "swarm-", this is a reserved term used by Swarm. Prefixing a name with "swarm-" will result in unexpected and unwanted behavior in Swarm.

For example:

Prefixing a group name with "swarm-project-" will result in, but is not limited to, the following issues:

- Swarm notifications will not be processed correctly for the group.
- The group will not be visible in Swarm.

Review the runtime dependencies

First, review the runtime dependencies before you install Swarm, see ["Runtime dependencies" on page 98](#).

Swarm and Helix Core Server installation considerations

Before installing Swarm and Helix Core Server you should consider the following:

- To ensure that all characters, including Unicode characters, are displayed and handled correctly by Swarm, configure your Helix Core Server in Unicode mode. For information on configuring your Helix Core Server in Unicode mode, see [Setting up and managing Unicode installations](#) in the *Helix Core Server Administrator Guide*.
- For a small system, you can run Swarm and Helix Core Server on the same machine.
- For larger systems, we recommend that Swarm and Helix Core Server are run on separate machines. The machines should be close to each other to maximize network performance.

Swarm needs to know about a number of Helix Core Server events to operate correctly. To enable this to happen, Helix Core Server Extensions (recommended) or Helix Core Server Triggers are installed on the Helix Core Server machine and they need to talk to Swarm. In either case, performance can be negatively affected if network lag between Swarm and the Helix Core Server is high.

- The Swarm and Helix Core Server machines do not need to have the same operating system. For example, Helix Core Server could be on a Windows server and Swarm could be on a CentOS server.

Note

You cannot install Swarm on a Windows machine.

Choose the installation process

Once you have reviewed the ["Runtime dependencies" on the facing page](#) and know that you can satisfy them, there are a number of ways to install Swarm.

Note

We recommend the package installation method to install Swarm whenever possible, see ["Install and configure Swarm from a package \(recommended\)" on page 115](#). Package installs ensure that all of the Swarm dependencies are installed and this is the easiest way to install Swarm. For a list of recommended operating systems for Swarm, see ["Recommended operating systems" on page 99](#).

Choose one of the following installation methods (we recommend the package installation method whenever possible):

- **Swarm RPM or Debian packages (recommended):**

1. Follow the steps provided in ["Install and configure Swarm from a package \(recommended\)" on page 115](#).
2. If you installed Helix Core Server Extensions in the previous step, skip this step and go to step 3:
Configure Helix Core Server for Swarm, see ["Installing Triggers" on page 204](#).
3. Review the post-install configuration options to customize your Swarm installation, see ["Post-install configuration options" on page 224](#).
4. Validate that Swarm is working correctly before using it in production, see ["Validate your Swarm installation" on page 239](#).

You are all set to start using Swarm. Enjoy!

Tip

To get started with Swarm, see the ["Quickstart" on page 29](#) chapter.

- **Swarm Docker container:**

1. Follow the steps provided in [Run Swarm using a Docker container](#).
2. If you installed Helix Core Server Extensions in the previous step, skip this step and go to step 3:

Configure Helix Core Server for Swarm, see ["Installing Triggers" on page 204](#).

3. Review the post-install configuration options to customize your Swarm installation, see ["Post-install configuration options" on page 224](#).
4. Validate that Swarm is working correctly before using it in production, see ["Validate your Swarm installation" on page 239](#).

You are all set to start using Swarm. Enjoy!

Tip

To get started with Swarm, see the ["Quickstart" on page 29](#) chapter.

■ Swarm.tgz (Tarball):

1. Follow steps provided in the ["Install and configure Swarm manually from a Tarball" on page 181](#).
2. Configure Redis, see ["Redis configuration" on page 182](#).
3. Configure Apache, see ["Apache configuration" on page 186](#).
4. Configure PHP, see ["PHP configuration" on page 190](#).
5. Configure Swarm, see ["Swarm configuration" on page 193](#).
6. Configure the Helix Core Server to notify Swarm about Helix Core Server events, see ["Configuring Helix Core Server event notification" on page 197](#).
7. Set up a recurring task to spawn workers, see ["Set up a recurring task to spawn workers" on page 220](#).
8. Review the post-install configuration options to customize your Swarm installation, see ["Post-install configuration options" on page 224](#).
9. Validate that Swarm is working correctly before using it in production, see ["Validate your Swarm installation" on page 239](#).

You are all set to start using Swarm. Enjoy!

Tip

To get started with Swarm, see the ["Quickstart" on page 29](#) chapter.

Upgrade Swarm

If you already have a working Swarm installation and you want to upgrade Swarm to a newer release, see ["Upgrading Swarm" on page 241](#).

Runtime dependencies

In order to successfully install, configure, and deploy Swarm, the following dependencies are required:

- A recommended operating system, see ["Recommended operating systems" below](#)
- An Apache server with the modules required by Swarm installed, see ["Apache web server" on the facing page](#)
- A supported version of PHP with the modules required by Swarm installed, see ["PHP" on page 102](#).
- A supported Helix Core Server deployment, and the ability to connect to it from the system hosting Swarm, see ["Helix Core Server requirements" on page 106](#)

Note

"Helix Core Server deployment" can refer to a running **p4d** or a proxy, replica, edge server, or commit server.

- Helix Core Server Extensions or Triggers to notify Swarm about Helix Core Server events, see ["Helix Core Server event notification" on page 107](#).
- **curl** or **wget** for Swarm worker operation, see ["Worker dependencies" on page 109](#)
- If Security-enhanced Linux (SELinux) is installed it must be configured correctly, see ["Security-enhanced Linux \(SELinux\)" on page 112](#)

Optional dependencies:

- **LibreOffice:** required to view office-type documents. For more information, see ["LibreOffice" on page 111](#).
- **zip:** command-line archiving tool: Required to download zip archives of files and folders. For more information, see ["Zip" on page 112](#).
- **Helix Authentication Service SSO:** enables Swarm to authenticate with SSO when the Helix Core Server is configured for Helix Authentication Service, see ["Helix Authentication Service SSO" on page 112](#).
- **Sendmail or equivalent:** required to use Swarm email notifications. For more information, see ["Swarm email notifications" on page 112](#).

Recommended operating systems

Perforce recommend that Swarm is installed on one of the following operating systems:

- **Ubuntu 20.04 LTS and 22.04 LTS** latest stable release should be used.
- **CentOS 7** latest point release with PHP libraries installed from the Remi repository. See ["PHP" on page 102](#).
- **RHEL 7, 8, and 9** latest stable release with PHP libraries installed from the Remi repository. See ["PHP" on page 102](#).
- **Rocky Linux 8 and 9** latest stable release with PHP libraries installed from the Remi repository. See ["PHP" on page 102](#).

- **Amazon Linux 2:** latest stable release with PHP libraries installed from Amazon Linux Extras. See "PHP" on page 102.

Note

You cannot install Swarm on a Windows machine.

Other Linux distributions

Important

Swarm has not been tested on Linux distributions that are not on our recommended list. Because of this, our Support team may not be able to help you to the same extent as they would for the recommended operating systems.

Swarm will probably run on any Linux distribution based on RHEL 7, 8, or 9. Swarm includes binary versions of P4PHP, the Perforce extension for PHP. This P4PHP dependency typically sets the limit to what Linux distributions you can install Swarm on.

You might be able to get Swarm running on another platform if you build P4PHP yourself and satisfy the other runtime dependencies. Instructions on how to obtain and build P4PHP from source can be found in the [P4PHP Release notes](#).

Important

P4PHP does not support threaded operation. If you compile P4PHP from source, ensure that the version of PHP you compile for is non-threaded.

Apache web server

Swarm requires Apache HTTP Server 2.4 or newer:

- <https://httpd.apache.org>

Important

The following notes are applicable for CentOS/RHEL7, RHEL 8, and RHEL 9:

- **CentOS/RHEL 7:** Use the Remi repository configuration package (**remi-release-7.rpm**) to give Swarm access to PHP 7/PHP 8 and to LibreOffice which is part of the optional package install. Use the **epel-release-latest-7.noarch.rpm** repository configuration package to give Swarm access to EPEL packages.

- **RHEL 8:** Use the Remi repository configuration package (**remi-release-8.rpm**) to give Swarm access to PHP 7/PHP 8 and to LibreOffice which is part of the optional package install. Use the **epel-release-latest-8.noarch.rpm** repository configuration package to give Swarm access to EPEL packages.
- **RHEL 9:** Use the Remi repository configuration packages (**remi-release-8.rpm** and **remi-release-9.rpm**) to give Swarm access to PHP 7/PHP 8 and to LibreOffice which is part of the optional package install. Use the **epel-release-latest-8.noarch.rpm** and **epel-release-latest-9.noarch.rpm** repository configuration package to give Swarm access to EPEL packages.
- **Swarm 2019.1 to 2020.1:** these versions of Swarm used PHP packages from the SCL repositories for CentOS/RHEL 7. This was to provide access to more recent versions of PHP than are available as standard on CentOS/RHEL. This required the use of the **httpd24-httpd** package for Apache. These were all installed into **/opt/rh**

Swarm 2020.2 and later: these versions of Swarm uses the Remi repository for CentOS/RHEL 7, RHEL 8, and RHEL 9. This provides PHP 7.4 installed in the standard file system structure. This means that the old **httpd24-httpd** version of Apache is no longer needed, and the standard system version of Apache is being used again.

The SCL Apache site configuration file was installed at this location for Swarm 2019.1 to 2020.1:

```
/opt/rh/httpd24/root/etc/httpd/conf.d/perforce-swarm-site.conf
```

If this exists when Swarm is upgraded to 2020.2 and later, this file is copied to **/etc/httpd/conf.d/perforce-swarm-site.conf** if there is no file at the destination. It is also re-written to change references from **/var/log/httpd24** to **/var/log/httpd**

If a site configuration file for Swarm already exists in **/etc/httpd**, the copy and re-write is not performed.

After upgrade, **httpd24-httpd** is disabled.

- To avoid seeing the Apache HTTP server Linux test page when you start the Apache server, comment out the content of the **welcome.conf** file located in the **/etc/httpd/conf.d/** directory.

- To avoid loading the Apache HTTP server example configuration instead of the Swarm configuration when the Apache server starts, rename the `autoindex.conf` file located in the `/etc/httpd/conf.d/` directory to `z-autoindex.conf` or similar. This is required because Apache runs the first conf file it finds in the `/etc/httpd/conf.d/` directory (alphabetical order) and that must be the `performer-swarm-site.conf` file.

Swarm also requires the following Apache modules:

- **Ubuntu:** `mod_php x` for interacting with PHP (usually installed with PHP)
Where x is the version of PHP you are running, for example 7. For supported versions of PHP, see "PHP" below.
- **CentOS/RHEL:** `php-fpm` for interacting with PHP (usually installed with PHP)
- `mod_rewrite` URL rewriting engine
http://httpd.apache.org/docs/2.4/mod/mod_rewrite.html

Important

Only the `prefork` MPM is supported. Use of the worker or event MPMs is not supported and is likely to cause problems because P4PHP does not support threaded operation.

For more information on the prefork MPM, see: <https://httpd.apache.org/docs/2.4/mod/prefork.html>

PHP

Swarm requires PHP 7.2, 7.4, 8.0, 8.1 or 8.2:

- <https://secure.php.net>

Important

- PHP must be non-threaded because P4PHP does not support threaded operation.
- **CentOS/RHEL 7:** Use the Remi repository configuration package (`remi-release-7.rpm`) to give Swarm access to PHP 7/PHP 8 and to LibreOffice which is part of the optional package install. Use the `epel-release-latest-7.noarch.rpm` repository configuration package to give Swarm access to EPEL packages.

- **RHEL 8:** Use the Remi repository configuration package (**remi-release-8.rpm**) to give Swarm access to PHP 7/PHP 8 and to LibreOffice which is part of the optional package install. Use the **epel-release-latest-8.noarch.rpm** repository configuration package to give Swarm access to EPEL packages.
- **RHEL 9:** Use the Remi repository configuration packages (**remi-release-8.rpm** and **remi-release-9.rpm**) to give Swarm access to PHP 7/PHP 8 and to LibreOffice which is part of the optional package install. Use the **epel-release-latest-8.noarch.rpm** and **epel-release-latest-9.noarch.rpm** repository configuration package to give Swarm access to EPEL packages.

Required PHP extensions

Swarm requires the following PHP extensions:

Tip

If you install Swarm from a package, all of these PHP extensions are automatically pulled in as dependencies on all platforms.

- **iconv** (character encoding converter)
<https://secure.php.net/iconv>
This is typically enabled by default with most PHP distributions
- **JSON** (JavaScript Object Notation)
<https://secure.php.net/json>
This is typically enabled by default with most PHP distributions, although recent distributions are making this optional.
- **Session** (session handling)
This is typically enabled by default with most PHP distributions
- **P4PHP** version 2019.1 or later (the Perforce PHP Extension)
This release of Swarm contains P4PHP 2023.2 with the Swarm package and tarball installations.

Note

Swarm package and tarball installations: two versions of P4PHP are supplied for each PHP 7 and PHP 8 version supported by Swarm. They are located in the **p4-bin/bin.linux26x86_64** directory.

- **perforce-php7x.so** and **perforce-php8x.so** compatible with systems using SSL 1.0.2
- **perforce-php7x-ssl1.1.1.so** and **perforce-php8x-ssl1.1.1.so** compatible with systems using SSL 1.1.1 (by default Ubuntu 20.04 uses SSL 1.1.1)
- **perforce-php8x-ssl3.so** compatible with systems using SSL 3.0.0 (by default Ubuntu 22.04 uses SSL 3.0.0)

Where **x** is the version of PHP 7 or PHP 8.

If the **perforce.ini** file is not pointing at the correct version of P4PHP and you connect to an SSL enabled Helix Core Server:

- The Swarm web-page will not load and you might see a **Connection Reset** error.
- There might be an **undefined symbol: SSLey** message in the Apache error log

Upgrading Swarm:

- **Swarm package:** the latest P4PHP version is installed automatically.
- **Swarm tarball installation:** you must configure Swarm to use the version of P4PHP in the new Swarm tarball. For swarm tarball upgrade instructions, see "[Upgrading a tarball installation](#)" on page 275.
- **php-xml** (DOM API for XML manipulation, the Swarm RSS feed will not work if it is not installed)
- **php-mbstring** (multi-byte character strings, the Swarm RSS feed will not work if it is not installed)
- **php-redis** (PHP extension for Redis, the Swarm cache will not work if it is not installed)
- **php-gd** (PHP extension for creating the blur images used for image icons on comments. If **php-gd** is not installed, images attached to comments display the file extension name instead of a blur or thumbnail image.)

Recommended PHP extension

Swarm greatly benefits from the following PHP extension:

- **Imagick** (integrates ImageMagick into PHP to improve Swarm's ability to preview graphics formats that web browsers typically cannot display and create to comment attachment thumbnails) <https://secure.php.net/imagick>

Installation instructions for [Imagick](#).

Note

If you see an error similar to the one below for **PS**, **PS2**, **EPS**, **PDF**, or **XPS** files, the **imagick policy.xml** file needs to be edited. For example, the error for a **PDF** file will be similar to:

```
convert-im6.q16: attempt to perform an operation not allowed
by the security policy `PDF' @
error/constitute.c/IsCoderAuthorized/408
```

If you see the error above:

1. Open the ImageMagick `policy.xml` file for editing:
`/etc/ImageMagick-6/policy.xml`
2. Disable lines with the pattern below for the following `<filetype>`s: PS, PS2, EPS, PDF, or XPS:
`<policy domain="coder" rights="none" pattern="<filetype>" />`
3. Save the `policy.xml` file.
4. Restart Apache for the changes to become active.

Building PHP from source (not recommended)

If you build PHP from source, the following dependencies are required:

- `openssl`
- `mcrypt`
- `zlib`
- `gettext`
- `curl`
- `apxs` (Apache extension tool)

See the source PHP documentation for details on how to include these modules in your PHP build. There is normally a `--with-xxxx` option that defines where the dependency is loaded from.

For example for `apxs`, `zlib`, and `openssl`:

```
./configure --with-apxs2=<path_to_apxs> --with-zlib=<path_to_zlib_
library> --with-openssl
```

Helix Core Server requirements

Note

Helix Core Server 2020.2 and later: any new file being shelved that has the same content as an existing shelved file refers to the existing archive file instead of creating a duplicate archive file. No Helix Core Server or Swarm configuration is required for this feature.

This Helix Core Server feature automatically reduces the space required for the Swarm-managed shelved review changelists. Swarm creates these changelists for its own internal use. Helix Core Server only updates new shelves, it does not retrospectively update your existing shelves.

For more information on the Swarm-managed changelists, see ["Internal representation" on page 424](#).

Tip

Swarm should work with newer versions of the Helix Core Server, but this is not tested.

Swarm works with any [supported versions](#) of the Helix Core Server (Standard Maintenance). This release of Swarm is tested against and supports the following Helix Core Server versions:

- 2022.1
- 2022.2
- 2023.1
- 2023.2

Download Helix Core Server from here: <https://www.perforce.com/downloads/helix-core-p4d>

Important

- Swarm does not support Helix Core Servers configured to use **P4AUTH**, see [Centralized authentication server \(P4AUTH\)](#) in the *Helix Core Server Administrator Guide*.
- Swarm does not support the Helix MFA Authenticator.

Helix Core Server automated user requirements for Swarm

Swarm requires an automated user with at least admin privileges in the Helix Core Server to enable Swarm to run against the Helix Core Server. The Swarm admin user also needs to be a [standard user type](#). This can be an existing user, or a new user created specifically to support Swarm.

Swarm does take up a license but if you are paying for the licenses you can request a background user for Swarm. This is a normal user that is provided so you can use Swarm without consuming one of your paid licenses. The form to request the background user can be found here:

<https://www.perforce.com/support/vcs/helix-core-request-background-user>

Important

If Helix Authentication Service is configured for your Helix Core Server, the user account running Swarm must not use the Helix Authentication Service.

For more information about setting up Helix Core Server, see [Installing and upgrading the server](#) in *Helix Core Server Administrator Guide*.

Swarm and Helix Core Server installation considerations

Before installing Swarm and Helix Core Server you should consider the following:

- To ensure that all characters, including Unicode characters, are displayed and handled correctly by Swarm, configure your Helix Core Server in Unicode mode. For information on configuring your Helix Core Server in Unicode mode, see [Setting up and managing Unicode installations](#) in the *Helix Core Server Administrator Guide*.

- For a small system, you can run Swarm and Helix Core Server on the same machine.
- For larger systems, we recommend that Swarm and Helix Core Server are run on separate machines. The machines should be close to each other to maximize network performance.

Swarm needs to know about a number of Helix Core Server events to operate correctly. To enable this to happen, Helix Core Server Extensions (recommended) or Helix Core Server Triggers are installed on the Helix Core Server machine and they need to talk to Swarm. In either case, performance can be negatively affected if network lag between Swarm and the Helix Core Server is high.

- The Swarm and Helix Core Server machines do not need to have the same operating system. For example, Helix Core Server could be on a Windows server and Swarm could be on a CentOS server.

Note

You cannot install Swarm on a Windows machine.

Helix Core Server event notification

Swarm needs to know about a number of Helix Core Server events to operate correctly, this can be done by using Helix Core Server Extensions (recommended) or Helix Core Server Triggers. Swarm installs include the Swarm Helix Core Server extension file and trigger scripts required for Swarm to get the events it needs from your Helix Core Server.

Triggers are still supported, but we recommend you use Helix Core Server Extensions. Helix Core Server Extensions are easier to install and maintain than Triggers.

Use one of the following:

- **Recommended:** Helix Core Server Extensions, see "[Helix Core Server Extensions dependencies](#)" below
- Helix Core Server Triggers, see "[Trigger dependencies](#)" below

Helix Core Server Extensions dependencies

Warning

If you are using the Swarm Helix Core Server extension, Swarm Helix Core Server Triggers must not be installed.

Helix Core Server Extensions must be installed on your Helix Core Server to complete your Swarm installation. The Swarm Extensions script is included with the Swarm product download and is copied to your Helix Core Server during configuration.

To install the Swarm Helix Core Server extension you need:

A compatible version of Helix Core Server for Helix Core Server Extensions:

- **Linux:** Helix Core Server 2021.2 and later. If you are using an earlier version of Helix Core Server, you must use triggers.
- **Windows:** Helix Core Server 2021.2 and later. If you are using an earlier version of Helix Core Server, you must use triggers.

Trigger dependencies

Warning

If you are using Swarm Helix Core Server Triggers, the Swarm Helix Core Server extension must not be installed.

Tip

We recommend you use Helix Core Server Extensions, Helix Core Server Extensions are easier to install and maintain than triggers. See "[Helix Core Server Extensions dependencies](#)" above.

Swarm triggers must be installed on the Helix Core Server to complete the Swarm installation.

The Swarm triggers require **perl 5.08+**:

<https://www.perl.org/get.html>

On the Windows platform, we have tested Swarm against Strawberry Perl. There are two Perl modules that are also required which may not be part of a minimal Perl installation.

- **HTTP::Tiny** is required to make calls to the Swarm server. If this is not present, then the trigger will attempt to use the command line curl program. This module is standard on Strawberry Perl on Windows, and available as a package with the version of Perl provided on CentOS 7 and Ubuntu 20.04.

IO::Socket::SSL is required if the Swarm server is configured to use SSL and **HTTP::Tiny** is present. This is provided as standard by Strawberry Perl, and available on Linux.

Warning

If the **HTTP::Tiny** module is not available the triggers require the use of **curl**. This must be installed for the triggers to function. On CentOS, for example, this can be done using the **yum** package installer using **yum install curl**.

Swarm triggers also require the following perl modules to be installed:

■ Windows:

- **JSON** is required to exchange data between the browser and the server and is included by default with Strawberry Perl.

■ Ubuntu:

- **JSON** is required to exchange data between the browser and the server and must be installed.
- **libjson-perl** is required to manipulate JSON formatted data and must be installed.

■ CentOS/RHEL:

- **JSON** is required to exchange data between the browser and the server and must be installed.
- **perl-JSON** is required to manipulate JSON formatted data and must be installed.

Swarm package installation on the same machine as P4D

If Swarm is installed from a package on the same machine as P4D, the triggers require the following perl modules to be installed:

■ Ubuntu:

- **Perl 5.08+**
- **libio-socket-ssl-perl** is required if the Swarm server is configured to use SSL.
- **libjson-perl** is required to manipulate JSON formatted data and must be installed.

■ CentOS/RHEL:

- **Perl 5.08+**
- **perl-IO-Socket-SSL** is required if the Swarm server is configured to use SSL.
- **perl-JSON** is required to manipulate JSON formatted data and must be installed.

Worker dependencies

Swarm uses short-lived workers to process the Swarm queue, new workers are regularly spawned by a recurring task.

One of the following must be installed to ensure new workers are regularly spawned:

- **curl**, download from: <https://curl.haxx.se/download.html>
- **wget**, download from: <https://ftp.gnu.org/gnu/wget/>

For more information about Swarm workers, see "Set up a recurring task to spawn workers" on page 220.

Redis server

Swarm requires Redis to manage its caches. Swarm caches data from the Helix Core Server to improve the performance of common searches in Swarm and to reduce the load on the Helix Core Server.

Redis is included with the Swarm package and Tarball installations:

- **Swarm package:** Redis is automatically installed on the Swarm machine and Swarm is automatically configured to use Redis.
- **Swarm Tarball installation:** see the [Redis installation and configuration](#) section for instructions about installing Redis.

Tip

When Swarm starts it verifies the Redis cache, during this time you cannot log in to Swarm. The time taken to verify the Redis cache depends on the number of users, groups, and projects Swarm has. Start-up time can be improved by persisting the memory cache. You can persist the memory cache by disabling background saves and enabling append saves in the `redis-server.conf` file, see "[Redis server configuration file](#)" on page 676.

Optional:

If you prefer to use your own Redis server, you must edit the Redis server connection configurable in Swarm. For information on setting your Redis server connection, see "[Use your own Redis server](#)" on page 235.

Supported Helix Core Visual Client (P4V)

Helix Swarm 2022.3 or later only works with Helix Core Visual Client (P4V) 2021.3 or later.

Supported Helix Authentication Service (HAS)

Helix Swarm works with any [supported versions](#) of Helix Authentication Service (Standard Maintenance).

Supported web browsers

The following browsers are supported for use with Swarm:

- Apple Safari, latest stable version
- Google Chrome, latest stable version
- Mozilla Firefox, latest stable version
- Microsoft Edge, latest stable version

Other web browsers might also work, including prior, development or beta builds of the above web browsers, but are not officially supported.

Note

- JavaScript and cookies must be enabled in the web browser for Swarm to operate.
- Mobile web browsers are not supported by Swarm.

Workflow prerequisites

For Swarm 2019.2 and later, the workflow feature is enabled by default.

Tip

If you are upgrading from an earlier version you will need to update your triggers, see ["Upgrading Swarm" on page 241](#).

The workflow feature requires:

- Workflow must be enabled, see the ["workflow " on page 752](#) system configurable.
- Your `swarm-trigger.conf` file must be configured correctly, see ["Using triggers to push events to Swarm" on page 204](#).
- Your Helix Core Server trigger table must contain the workflow trigger lines, see ["Using triggers to push events to Swarm" on page 204](#).

Optional dependencies

LibreOffice

Swarm can display previews of office-type documents when LibreOffice is installed on the Swarm server. Installation is not required, but when LibreOffice is installed Swarm automatically detects its presence.

For more information about LibreOffice, see ["LibreOffice" on page 560](#).

Zip

You can download a ZIP archive of files/folders when the zip command-line tool is installed on the Swarm server.

For more information about installing and configuring Zip, see ["Zip archive" on page 562](#).

Helix Authentication Service SSO

Swarm support for SSO using Helix Authentication Service with Helix Core Server requires Helix Core Server 2019.1 or later

For more information about configuring Swarm to authenticate with Helix Authentication Service, see ["Single Sign-On PHP configuration" on page 712](#).

Swarm email notifications

Sendmail or an equivalent is required to use Swarm email notifications. By default, the configuration in `php.ini` relies on SendMail being installed. For more information about configuring email for Swarm, see ["Email configuration" on page 602](#).

Security-enhanced Linux (SELinux)

Swarm supports SELinux on CentOS/RHEL and Amazon Linux 2. SELinux is an advanced access control mechanism that improves security for Linux distributions.

SELinux operates in one of three modes:

- **enforcing**: this mode blocks and logs any actions that do not match the defined security policy.
- **permissive**: this mode logs actions that do not match the defined security policy but these actions are not blocked.
- **disabled**: in this mode SELinux is off, actions are not blocked and are not logged.

Tip

To check the mode SELinux is operating in, view the `/etc/selinux/config` file with vi or a similar editor:

```
root $ vi /etc/selinux/config
```

For instructions on configuring SELinux:

- **CentOS/RHEL**: See ["SELinux configuration" on page 149](#).
- **Amazon Linux 2**: See ["SELinux configuration" on page 166](#).

If you see a Swarm configuration error similar to the error shown below, SELinux has not been correctly configured for Swarm. Check you have configured SELinux on [CentOS/REHL](#) or [Amazon Linux 2](#) correctly.

Swarm has detected a configuration error

Problem detected:

- The data directory (/opt/perforce/swarm/data) is not writeable.

Please investigate the below PHP error below:

Cannot write to cache directory ('/opt/perforce/swarm/data/cache'). Check permissions.

/opt/perforce/swarm/module/Application/Module.php on line 329

For more information, please see the [Install and upgrade Swarm](#) documentation; in particular:

- [Runtime dependencies](#)
- [Installation](#)
- [PHP configuration](#)
- [Swarm configuration](#)

If you are using SELinux: check it is configured correctly, see [Security-enhanced Linux \(SELinux\)](#).

Please restart your web server after making any PHP changes.

If you change your configuration you must delete your Swarm config cache, see [Swarm config cache file delete](#).

Choose the installation process

Once you have reviewed the "Runtime dependencies" on page 98 and know that you can satisfy them, there are a number of ways to install Swarm.

Note

We recommend the package installation method to install Swarm whenever possible, see "Install and configure Swarm from a package (recommended)" on page 115. Package installs ensure that all of the Swarm dependencies are installed and this is the easiest way to install Swarm. For a list of recommended operating systems for Swarm, see "Recommended operating systems" on page 99.

Choose one of the following installation methods (we recommend the package installation method whenever possible):

■ Swarm RPM or Debian packages (recommended):

1. Follow the steps provided in "Install and configure Swarm from a package (recommended)" on page 115.
2. If you installed Helix Core Server Extensions in the previous step, skip this step and go to step 3:
Configure Helix Core Server for Swarm, see "Installing Triggers" on page 204.

3. Review the post-install configuration options to customize your Swarm installation, see ["Post-install configuration options" on page 224](#).
4. Validate that Swarm is working correctly before using it in production, see ["Validate your Swarm installation" on page 239](#).

You are all set to start using Swarm. Enjoy!

Tip

To get started with Swarm, see the ["Quickstart" on page 29](#) chapter.

■ Swarm Docker container:

1. Follow the steps provided in [Run Swarm using a Docker container](#).
2. If you installed Helix Core Server Extensions in the previous step, skip this step and go to step 3:
Configure Helix Core Server for Swarm, see ["Installing Triggers" on page 204](#).
3. Review the post-install configuration options to customize your Swarm installation, see ["Post-install configuration options" on page 224](#).
4. Validate that Swarm is working correctly before using it in production, see ["Validate your Swarm installation" on page 239](#).

You are all set to start using Swarm. Enjoy!

Tip

To get started with Swarm, see the ["Quickstart" on page 29](#) chapter.

■ Swarm.tgz (Tarball):

1. Follow steps provided in the ["Install and configure Swarm manually from a Tarball" on page 181](#).
2. Configure Redis, see ["Redis configuration" on page 182](#).
3. Configure Apache, see ["Apache configuration" on page 186](#).
4. Configure PHP, see ["PHP configuration" on page 190](#).
5. Configure Swarm, see ["Swarm configuration" on page 193](#).
6. Configure the Helix Core Server to notify Swarm about Helix Core Server events, see ["Configuring Helix Core Server event notification" on page 197](#).
7. Set up a recurring task to spawn workers, see ["Set up a recurring task to spawn workers" on page 220](#).
8. Review the post-install configuration options to customize your Swarm installation, see ["Post-install configuration options" on page 224](#).

9. Validate that Swarm is working correctly before using it in production, see ["Validate your Swarm installation" on page 239](#).

You are all set to start using Swarm. Enjoy!

Tip

To get started with Swarm, see the ["Quickstart" on page 29](#) chapter.

Install and configure Swarm from a package (recommended)

This chapter covers the initial installation and configuration of Swarm from a package. Helix Swarm is available in two distribution package formats: Debian (`.deb`) for Ubuntu systems and RPM (`.rpm`) for CentOS, RedHat Enterprise Linux (RHEL), and Amazon Linux 2.

Using distribution packages greatly simplifies the installation, updating, and removal of software, as the tools that manage these packages are aware of the dependencies for each package.

Swarm packages are available for:

- **Ubuntu 20.04 LTS and 22.04 LTS** latest stable release should be used.
- **CentOS 7** latest point release with PHP libraries installed from the Remi repository. See ["PHP" on page 102](#).
- **RHEL 7, 8, and 9** latest stable release with PHP libraries installed from the Remi repository. See ["PHP" on page 102](#).
- **Rocky Linux 8 and 9** latest stable release with PHP libraries installed from the Remi repository. See ["PHP" on page 102](#).
- **Amazon Linux 2:** latest stable release with PHP libraries installed from Amazon Linux Extras. See ["PHP" on page 102](#).

For instructions on installing Swarm from a package:

Install and configure Swarm on Ubuntu

Swarm is available as a package for Ubuntu 20.04 LTS and 22.04 LTS. Using distribution packages greatly simplifies the installation, updating, and removal of software, as the tools that manage these packages are aware of the dependencies for each package.

Important

We recommend that the latest stable release of Ubuntu 20.04 LTS or 22.04 LTS is used.

Important

Swarm does not support Helix Core Servers configured to use **P4AUTH**, see [Centralized authentication server \(P4AUTH\)](#) in the *Helix Core Server Administrator Guide*.

Note

- Swarm can be connected to Helix Core Servers (P4D) and commit servers.
To configure Swarm to connect to more than one Helix Core Server (P4D), see ["Multiple Helix Core Server instances"](#) on page 646.
To configure Swarm to connect to a Helix Core Server configured to use *commit-edge architecture*, see ["Commit-edge deployment"](#) on page 587.
- Swarm must not be connected to Helix Broker, Helix Proxy, Helix Edge, forwarding replica, or read-only replica servers.

Installation

Important

- Review the runtime dependencies before you install Swarm, see ["Runtime dependencies"](#) on page 98.
- Review the PHP requirements before you upgrade Swarm, see ["PHP"](#) on page 102.
- Review the Helix Core Server requirements before you install Swarm, see ["Helix Core Server requirements"](#) on page 106.

1. Configure the Perforce package repository, on the server to host Swarm and on the server hosting your Helix Core Server.

Important

If the server hosting your Helix Core Server cannot use packages, for example when it is running Windows, skip this step on that server.

Follow the instructions for your OS distribution:

- **Ubuntu 20.04:** As root, create the file `/etc/apt/sources.list.d/perforce.list` with the following content:
`deb http://package.perforce.com/apt/ubuntu/ focal release`
- **Ubuntu 22.04:** As root, create the file `/etc/apt/sources.list.d/perforce.list` with the following content:
`deb http://package.perforce.com/apt/ubuntu/ jammy release`

2. Import the Perforce package signing key, on the server to host Swarm and the server hosting your Helix Core Server.

Important

- If the server hosting your Helix Core Server cannot use packages, for example when it is running Windows, skip this step on that server.
- If you are using an older version of OpenSSL that does not support TLSv1.2 or higher, replace 'https' with 'http' in the **wget** command.

Run the following commands as root for your OS distribution:

■ Ubuntu 20.04:

```
wget -qO - https://package.perforce.com/perforce.pubkey |
sudo apt-key add -
sudo apt-get update
```

■ Ubuntu 22.04:

```
wget -qO - https://package.perforce.com/perforce.pubkey |
gpg --dearmor | sudo tee /usr/share/keyrings/perforce.gpg
sudo apt-get update
```

For information about how to verify the authenticity of the signing key, see:

<https://www.perforce.com/perforce-packages>

3. Install the main Swarm package on the server to host Swarm.

```
sudo apt-get install helix-swarm
```

Important

When the Swarm installation has completed, you are prompted to run the **configure-swarm.sh** post-installation script.

Do not run this script until you have completed the rest of these **Installation** instructions. Instructions for running the **configure-swarm.sh** post-installation script are in the **Post-Installation configuration** section referenced in the final step of the **Installation** instructions.

4. Swarm needs to know about some Helix Core Server events to operate correctly. Use Helix Core Server Extensions (recommended) or Helix Core Server Triggers to notify Swarm about these events. The Swarm Helix Core Server extension can be installed automatically by the Swarm **configure-swarm.sh** post-installation script, but Triggers must be manually installed. To use Helix Core Server Extensions, ignore this step and skip to the next step.

Trigger installation only (not recommended): Install the Swarm triggers package on the server hosting your Helix Core Server. This might be the server hosting Swarm or elsewhere on your network.

Important

- If the server hosting your Helix Core Server cannot use packages, for example when it is running Windows, you need to copy the appropriate Swarm trigger script from `/opt/perforce/swarm/p4-bin/scripts` to the server hosting your Helix Core Server. The `swarm-trigger.pl` is for both Linux and Windows systems. Once copied, the trigger script needs to be configured. See ["Installing Triggers" on page 204](#) for details.
- The package installs a config file at `/opt/perforce/etc/swarm-trigger.conf` that you will need to modify. See ["Installing Triggers" on page 204](#) for more details on configuring that file.

Install the Swarm triggers package on the server hosting your Helix Core Server

```
sudo apt-get install helix-swarm-triggers
```

5. **Optional:** Install the Swarm optional package, on the server hosting Swarm.

While not required, installing this package installs the dependencies required to use the ImageMagick and LibreOffice Swarm modules. These modules provide previews of a variety of image and office documents.

```
sudo apt-get install helix-swarm-optional
```

6. Complete the ["Post-installation configuration" below](#) steps.

Post-installation configuration

Once the helix-swarm package has been installed, additional configuration is required. Perform the following steps:

1. Use the Swarm post-installation script to configure Swarm on the server hosting Swarm.

Important

- If your Helix Core Server is configured for Helix Authentication Service, the Helix Core Server must be temporarily configured to allow fall-back to passwords while you establish a connection to the Helix Core Server. Run the following command on the Helix Core Server to enable fall-back to passwords:

```
p4 configure set auth.sso.allow.passwd=1
```

- To use Helix Core Server Extensions, make sure that Extensions are installed and configured on your Helix Core Server before running the configuration script. For information about Helix Core Server Extensions, see [p4 extension](#) in [Helix Core Command-Line \(P4\) Reference](#) and [Extensions overview](#) in [Helix Core Extensions Developer Guide](#).

Note

The Swarm post-installation configuration script can be used in a few different ways. The steps below outline the most straightforward configuration using an interactive install, but you can review the options by running:

```
sudo /opt/perforce/swarm/sbin/configure-swarm.sh -h
```

Run the interactive post-installation configuration script:

```
sudo /opt/perforce/swarm/sbin/configure-swarm.sh
```

The configuration script displays the following summary:

```
-----
configure-swarm.sh: Thu Aug 26 11:29:49 PDT 2021: commencing
configuration of Swarm
Summary of arguments passed:
Interactive?      [yes]
Force?           [no]
P4PORT           [(not specified, will prompt)]
Swarm user       [(not specified, will prompt, will suggest
swarm)]
Swarm password   [(not specified, will prompt)]
Email host       [(not specified, will prompt)]
Create depot     [(not specified, will prompt)]
Bypass Lock      [(not specified)]
Use Extensions?  [(not specified, will prompt)]
Swarm host       [(not specified, will prompt, will suggest
myhost)]
Swarm port       [(default (80))]
Swarm base URL   [(default (empty))]
Create Swarm user? [yes]
Super user       [(not specified, will prompt)]
Super password   [(not specified, will prompt)]
```

2. Provide information to the configuration script.

After the summary, the configuration script prompts for the following information:

- a. Specify a value for **P4PORT** in the form: *my-helix-core-server:1666*

```
No P4PORT specified
```

```
Swarm requires a connection to a Helix Core Server.  
Please supply the P4PORT to connect to.
```

```
Helix Core Server address (P4PORT):
```

Specify the hostname and port for your Helix Core Server. If defined, the value for P4PORT is used as the default. The configuration script verifies that it can connect:

```
-response: [myp4host:1666]
```

```
Checking P4PORT [myp4host:1666]...
```

```
-P4 command line to use: [/opt/perforce/bin/p4 -p  
myp4host:1666]
```

```
Attempting connection to [myp4host:1666]...
```

```
-connection successful:
```

```
  Server address: myp4host:1666
```

```
  Server version: P4D/LINUX26X86_64/2021.1/2179737  
(2021/04/24)
```

```
  Server license: 10000 users (support ends 2022/05/16)
```

```
  Server license-ip: 192.168.0.1
```

Important

If your Helix Core Server is deployed using the commit-edge architecture, ensure that the Swarm port value points to the commit server.

For more information, see the [Commit-edge](#) chapter in the *Helix Core Server Administrator Guide*.

- b. Specify the userid and password of a normal user with admin-level privileges in the Helix Core Server.

```
Checking Swarm user credentials...
```

```
No Swarm user specified
```

```
Swarm requires a Helix user account with 'admin' rights.  
Please provide a username and password for this account.  
If this account does not have 'admin' rights, it will
```

be set for this user.

Helix username for the Swarm user [swarm]:

Enter the userid. The default is *swarm*.

Note

To allow users to clean up reviews created by other users when the review is committed, the Helix Core Server user account must have *super* permissions. See ["Review cleanup" on page 678](#).

-response: [swarm]

Helix password or login ticket for the Swarm user (typing hidden):

Enter the login ticket, or password, for the userid.

Important

You must use a long-lived login ticket for the Swarm user.

Note

You can obtain a login ticket by running (in another shell):

```
p4 -p myp4host:1666 -u userid login -p
```

If the login ticket you provide expires in less than a year, you will receive a warning. If you receive this warning, ask your Helix Core Server administrator to make sure the **swarm** user is in a Helix Core Server group that has the **Timeout** field set to a year or more. See [p4 group](#) in the [Helix Core Command-Line \(P4\) Reference](#).

Checking Swarm user credentials...

-checking if user [swarm] exists in [myp4host:1666]...

-user exists

Obtaining Helix login ticket for [swarm] in [myp4host:1666]...

-login ticket obtained

Checking user [swarm]'s ticket against [myp4host:1666]...

-login ticket is good

Checking user [swarm] has at least access level [admin]...

-user has maximum access level [admin]

-user meets minimum access level [admin]

c. Specify the hostname for the Swarm UI.

swarm needs a distinct hostname that users can enter into their browsers to access Swarm. Ideally, this is a fully-qualified domain name, for example 'swarm.company.com', but it can be just a hostname, for example 'swarm'.

Whatever hostname you provide should be Swarm-specific and not shared with any other web service on this host.

Note that the hostname you specify typically requires configuration in your network's DNS service. If you are merely testing Swarm, you can add a hostname->IP mapping entry to your computer's hosts configuration.

Hostname for this Swarm server [myhost]:

Note

The default is the current hostname. The configuration script does not verify that the hostname actually works (DNS configuration may not exist yet).

d. Specify a mail relay host or leave empty to use your local mail handler.

Swarm can use a mail relay host to send email notifications. Leave empty if you want to use the local mail handler (for example Sendmail, Postfix etc), or enter a hostname (for example mx.yourdomain.com) to use a relay host.

Mail relay host:

Note

The configuration script does not verify that the mail relay host you provide actually accepts SMTP connections.

- e. To enable file attachments on comments, Swarm must have a depot location to save the attachment files in. Swarm can create this depot and set the depot protections for you automatically or you can manually set it up later.

Important

You must be a user with *super* user permissions to create the **.swarm** depot and set protections.

```
- checking depot
```

```
Swarm has the ability to store attachments against review
comments. To
do this, it needs to have a depot where they are stored. By
default,
this is //.swarm. We can create the depot for you
automatically, and
set the protections on it to the following:
```

```
list user * * -//.swarm/...
admin user swarm * //.swarm/...
super user super * //.swarm/...
```

```
If you want to enable attachments, and for the depot and
protections
to be set up for you, then say yes here. If you say no, then
you can
still do this manually later.
Do you want to create a .swarm depot and set protections?
(y/n) [n]
```

When prompted to automatically create the depot, select one of the following:

- Type **y** to automatically create the `.swarm` depot and set protections.
- Type **n** skip depot creation. You can create the depot and set the protections manually later if you want to, see ["Comment attachments" on page 580](#).

If you choose to automatically create the depot, the script will create the depot, set the protections, and report when completed.

```
Depot //.swarm successfully created. Protections have been set
so
that only the 'swarm' and 'super' have permissions to
access it directly.
```

- f. Swarm needs to know about the exclusive locks to work with exclusively opened files. To enable exclusive locks, set the `filetype.bypasslock` to 1. For more information about exclusive locks, see ["Handling Exclusive Locks" on page 238](#).

Tip

If this setting is not enabled in Helix Core Server, Swarm will report exceptions when working with exclusively opened files similar to **Cannot unshelve review (x). One or more files are exclusively open**, and that you must have the `filetype.bypasslock` configurable enabled.

- g. Swarm needs to know about some Helix Core Server events to operate correctly. Use Helix Core Server Extensions (recommended) or Helix Core Server Triggers to notify Swarm about these events. Helix Core Server Extensions are easier to install and maintain than triggers.

Important

You must be a user with *super* user permissions to install and configure Helix Core Server Extensions.

```
Do you want to use Swarm's Helix Core server extension?
Configuring Server extensions requires super user access to
the Helix Server.
If you install the Swarm server extension, do not install the
Swarm triggers.
Server extensions are supported for:
* Linux: Helix server 19.2 and later.
* Windows: Helix server 21.2 and later.

Use server extensions?
```


When prompted to **Use server Extensions**, choose one of the following:

- **Recommended:** Type **y** to use Helix Core Server Extensions. The configuration script will try and install and configure the Swarm Helix Core Server extension.
- Type **n** to use Helix Core Server Triggers. Triggers must be installed manually, Swarm will prompt you to do this when the Swarm configuration script completes.

Note

If you choose to install the Swarm Helix Core Server Extensions, the script will:

- check your Helix Core Server supports Helix Core Server Extensions
- check if Helix Core Server Extensions are installed and configured on your Helix Core Server. If they are, Swarm does not need to do anything
- install and configure the Swarm Helix Core Server extension on your Helix Core Server

If any of these checks fail, Swarm will not install the Swarm Helix Core Server extension and will report the issues on the configuration summary screen.

```

.....
-restarting Apache...
-Apache restarted
configure-swarm.sh: Thu Aug 26 11:31:36 PDT 2021: completed
configuration of Helix Swarm

::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::
:::::::::
::
::  Swarm is now configured and available at:
::
::      http://myhost/
::
::  You may login as the Swarm user [swarm] using the password
::  you specified.
::
::  Ensure that you have configured the Swarm hostname in your
::  network's DNS, or have added an IP address-to-hostname
::  mapping to your computer's hosts configuration so that you
::  can access Swarm.
::
::  Server side extensions are installed and configured
::  on your P4D server.
::
::  Documentation for optional post-install configuration,
such as
::  configuring Swarm to use HTTPS, operate in a sub-folder,
or on a
::  custom port, is available:
::

```

```

::
https://www.perforce.com/perforce/doc.current/manuals/swarm/se
tup.post.html
::

.....
:

```

Note

If you have installed Swarm on a host that does not provide other web services, you may wish to disable Apache's default site configuration. Doing so means that regardless of the hostname a user might use to reach the web server hosting Swarm, Swarm would be presented.

Be aware that disabling Apache's default site configuration could disable existing web services or content.

Disabling Apache's default site configuration on Ubuntu hosts is easy. Run:

```
$ sudo a2dissite 000-default
```

To add Swarm as an HTML tab in the Helix Core Visual Client (P4V), see [HTML Tabs](#) section in the *P4VJS Developer Guide*.

3. The basic Swarm configuration is now complete.

Important

If your Helix Core Server is configured for Helix Authentication Service, you can force all of your users to authenticate via your Identity Provider (IdP) by disabling fall-back to passwords. To disable fall-back to passwords on the Helix Core Server, run the following command:

```
p4 configure set auth.sso.allow.passwd=0
```

4. Do one of the following:
 - **If the Swarm Helix Core Server extension was installed by the configuration script:** Review the post-install configuration options to customize your Swarm installation, see ["Post-install configuration options" on page 224](#).
 - **If the Swarm Helix Core Server extension was not installed, you must install triggers:** Configure Helix Core Server for Swarm, see ["Installing Triggers" on page 204](#).

Upgrading

See ["Upgrading a Ubuntu package installation" on page 242](#) for details.

Uninstall

1. Depending on how you have configured your Helix Core Server events, do one of the following:
 - Uninstall the Swarm Helix Core Server extension. To do this, run the following command on your Helix Core Server:


```
p4 extension --delete Perforce::helix-swarm --yes
```
 - Uninstall the Swarm triggers. As a *super* user, run the `p4 triggers` command from your Helix Core Server and manually remove all the Swarm trigger code lines.
2. Uninstall the Swarm main package, the Swarm triggers package, and the Swarm optional package. The following example assumes that the packages are all installed on the same server:

```
sudo apt-get remove helix-swarm helix-swarm-triggers helix-swarm-optional
```

3. Remove the Swarm trigger scripts from the server hosting your Helix Core Server.

Important

If you manually installed the trigger script, perhaps because the server hosting your Helix Core Server cannot use packages (e.g. Windows), manually remove the script.

Install and configure Swarm on CentOS/RHEL

Swarm packages can be used to install Swarm on CentOS 7, RHEL 7, RHEL 8, and RHEL 9. Using packages greatly simplifies the installation, updating, and removal of software, because the tools that manage packages are aware of the dependencies for each package.

Do one of the following depending on your operating system:

- ["Install Swarm on CentOS/RHEL 7 " on the next page](#)
- ["Install Swarm on RHEL 8" on page 132](#)
- ["Install Swarm on RHEL 9" on page 136](#)

Related topics

- ["Post-installation configuration " on page 138](#)
- ["SELinux configuration" on page 149](#)
- ["Upgrading" on page 152](#)
- ["Uninstall" on page 152](#)

Install Swarm on CentOS/RHEL 7

Important

- We recommend that the latest point release is used. At the time of this Swarm release, the latest stable CentOS/RHEL 7 release is 7.9.
- Swarm does not support Helix Core Servers configured to use **P4AUTH**, see [Centralized authentication server \(P4AUTH\)](#) in the *Helix Core Server Administrator Guide*.

Note

- Swarm can be connected to Helix Core Servers (P4D) and commit servers.
To configure Swarm to connect to more than one Helix Core Server (P4D), see ["Multiple Helix Core Server instances"](#) on page 646.
To configure Swarm to connect to a Helix Core Server configured to use *commit-edge architecture*, see ["Commit-edge deployment"](#) on page 587.
- Swarm must not be connected to Helix Broker, Helix Proxy, Helix Edge, forwarding replica, or read-only replica servers.

Installation

Important

- Review the runtime dependencies before you install Swarm, see ["Runtime dependencies"](#) on page 98.
- Review the PHP requirements before you upgrade Swarm, see ["PHP"](#) on page 102.
- Review the Helix Core Server requirements before you install Swarm, see ["Helix Core Server requirements"](#) on page 106.

1. Configure the Perforce package repository, on the server to host Swarm and on the server hosting your Helix Core Server.

Important

If the server hosting your Helix Core Server cannot use packages, for example when it is running Windows, skip this step on that server.

As root, create the file `/etc/yum.repos.d/perforce.repo` with the following content:

```
[Perforce]
name=Perforce
baseurl=http://package.perforce.com/yum/rhel/7/x86_64/
```

```
enabled=1
gpgcheck=1
```

2. Import the Perforce package signing key, on the server to host Swarm and the server hosting your Helix Core Server.

Important

If the server hosting your Helix Core Server cannot use packages, for example when it is running Windows, skip this step on that server.

Run the following command as root:

```
rpm --import https://package.perforce.com/perforce.pubkey
```

For information about how to verify the authenticity of the signing key, see:

<https://www.perforce.com/perforce-packages>

3. Install the main Swarm package on the server to host Swarm (run these commands as root):

- a. Deploy the `epel-release-latest-7.noarch.rpm` repository configuration package to give Swarm access to EPEL packages:

```
yum install https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm
```

- b. Deploy the Remi repository configuration package to give Swarm access to PHP 7.x (only required the first time you upgrade to PHP 7.x):

```
yum install https://rpms.remirepo.net/enterprise/remi-release-7.rpm
```

Tip

If you don't deploy the Remi repository, you will see dependency errors when you do the next steps.

- c. Install the `yum-utils` package to give access to the `yum-config-manager` command:

```
yum install yum-utils
```

- d. **RHEL only:** Enable the optional channel for some dependencies:

```
subscription-manager repos --enable=rhel-7-server-optional-rpms
```

- e. Install the **Default/Single version** of PHP:

- i. Disable `remi-php*`:

```
yum-config-manager --disable 'remi-php*'
```

- ii. Enable PHP 7.4:

```
yum-config-manager --enable remi-php74
```

- iii. Run an upgrade for PHP:

```
yum update
```

- f. Install Swarm and accept the prompts to import the GPG keys for Remi and EPEL when requested:

```
yum install helix-swarm
```

Important

When the Swarm installation has completed, you are prompted to run the **configure-swarm.sh** post-installation script.

Do not run this script until you have completed the rest of these **Installation** instructions. Instructions for running the **configure-swarm.sh** post-installation script are in the **Post-Installation configuration** section referenced in the final step of the **Installation** instructions.

Note

The firewall configuration may need to be adjusted to allow access to the web server.

```
sudo firewall-cmd --permanent --add-service=http  
sudo systemctl reload firewalld
```

If you subsequently wish to enable HTTPS, run (as root):

```
sudo firewall-cmd --permanent --add-service=https  
sudo systemctl reload firewalld
```

4. Swarm needs to know about some Helix Core Server events to operate correctly. Use Helix Core Server Extensions (recommended) or Helix Core Server Triggers to notify Swarm about these events. The Swarm Helix Core Server extension can be installed automatically by the Swarm **configure-swarm.sh** post-installation script, but Triggers must be manually installed. To use Helix Core Server Extensions, ignore this step and skip to the next step.

Trigger installation only (not recommended): Install the Swarm triggers package on the server hosting your Helix Core Server. This might be the server hosting Swarm or elsewhere on your network.

Important

- If the server hosting your Helix Core Server cannot use packages, for example when it is running Windows, you need to copy the appropriate Swarm trigger script from `/opt/perforce/swarm/p4-bin/scripts` to the server hosting your Helix Core Server. The `swarm-trigger.pl` is for both Linux and Windows systems. Once copied, the trigger script needs to be configured. See ["Installing Triggers" on page 204](#) for details.
- The package installs a config file at `/opt/perforce/etc/swarm-trigger.conf` that you will need to modify. See ["Installing Triggers" on page 204](#) for more details on configuring that file.

Install the Swarm triggers package on the server hosting your Helix Core Server (run this command as root):

```
yum install helix-swarm-triggers
```

5. **Optional:** Install the Swarm optional package, on the server hosting Swarm.

While not required, installing this package installs the dependencies required to use the ImageMagick and LibreOffice Swarm modules. These modules provide previews of a variety of image and office documents. Run this command as root:

```
yum install helix-swarm-optional
```

6. Complete the ["Post-installation configuration " on page 138](#) steps.

Install Swarm on RHEL 8

Important

- We recommend that the latest stable release is used. At the time of this Swarm release, the latest stable RHEL 8 release is 8.8.
- Swarm does not support Helix Core Servers configured to use **P4AUTH**, see [Centralized authentication server \(P4AUTH\)](#) in the *Helix Core Server Administrator Guide*.

Note

- Swarm can be connected to Helix Core Servers (P4D) and commit servers.
To configure Swarm to connect to more than one Helix Core Server (P4D), see ["Multiple Helix Core Server instances" on page 646](#).
To configure Swarm to connect to a Helix Core Server configured to use *commit-edge architecture*, see ["Commit-edge deployment" on page 587](#).
- Swarm must not be connected to Helix Broker, Helix Proxy, Helix Edge, forwarding replica, or read-only replica servers.

Installation

Important

- Review the runtime dependencies before you install Swarm, see ["Runtime dependencies" on page 98](#).
- Review the PHP requirements before you upgrade Swarm, see ["PHP" on page 102](#).
- Review the Helix Core Server requirements before you install Swarm, see ["Helix Core Server requirements" on page 106](#).

1. Configure the Perforce package repository, on the server to host Swarm and on the server hosting your Helix Core Server.

Important

If the server hosting your Helix Core Server cannot use packages, for example when it is running Windows, skip this step on that server.

As root, create the file `/etc/yum.repos.d/perforce.repo` with the following content:

```
[Perforce]
name=Perforce
baseurl=http://package.perforce.com/yum/rhel/8/x86_64/
enabled=1
gpgcheck=1
```

2. Import the Perforce package signing key, on the server to host Swarm and the server hosting your Helix Core Server.

Important

If the server hosting your Helix Core Server cannot use packages, for example when it is running Windows, skip this step on that server.

Run the following command as root:

```
rpm --import https://package.perforce.com/perforce.pubkey
```

For information about how to verify the authenticity of the signing key, see:

<https://www.perforce.com/perforce-packages>

3. Install the main Swarm package on the server to host Swarm (run these commands as root):
 - a. Deploy the `epel-release-latest-8.noarch.rpm` repository configuration package to give Swarm access to EPEL packages:

```
dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-8.noarch.rpm
```

- b. Deploy the Remi repository configuration package to give Swarm access to PHP 7.x (only required the first time you upgrade to PHP 7.x):

```
dnf install https://rpms.remirepo.net/enterprise/remi-release-8.rpm
```

Tip

If you don't deploy the Remi repository, you will see dependency errors when you do the next steps.

- c. Install the **yum-utils** package to give access to the **yum-config-manager** command:

```
dnf install yum-utils
```

- d. Install the **Default/Single version** of PHP:

- i. Enable the module stream for PHP 7.4:

```
dnf module reset php
```

- ii. Install PHP 7.4:

```
dnf module install php:remi-7.4
```

- iii. Run an upgrade for PHP:

```
dnf update
```

- e. Install Swarm and accept the prompts to import the GPG keys for Remi and EPEL when requested:

```
yum install helix-swarm
```

Important

When the Swarm installation has completed, you are prompted to run the **configure-swarm.sh** post-installation script.

Do not run this script until you have completed the rest of these **Installation** instructions. Instructions for running the **configure-swarm.sh** post-installation script are in the **Post-Installation configuration** section referenced in the final step of the **Installation** instructions.

Note

The firewall configuration may need to be adjusted to allow access to the web server.

```
sudo firewall-cmd --permanent --add-service=http
sudo systemctl reload firewalld
```

If you subsequently wish to enable HTTPS, run (as root):

```
sudo firewall-cmd --permanent --add-service=https
sudo systemctl reload firewalld
```

4. Swarm needs to know about some Helix Core Server events to operate correctly. Use Helix Core Server Extensions (recommended) or Helix Core Server Triggers to notify Swarm about these events. The Swarm Helix Core Server extension can be installed automatically by the Swarm **configure-swarm.sh** post-installation script, but Triggers must be manually installed. To use Helix Core Server Extensions, ignore this step and skip to the next step.

Trigger installation only (not recommended): Install the Swarm triggers package on the server hosting your Helix Core Server. This might be the server hosting Swarm or elsewhere on your network.

Important

- If the server hosting your Helix Core Server cannot use packages, for example when it is running Windows, you need to copy the appropriate Swarm trigger script from **/opt/perforce/swarm/p4-bin/scripts** to the server hosting your Helix Core Server. The **swarm-trigger.pl** is for both Linux and Windows systems. Once copied, the trigger script needs to be configured. See ["Installing Triggers" on page 204](#) for details.
- The package installs a config file at **/opt/perforce/etc/swarm-trigger.conf** that you will need to modify. See ["Installing Triggers" on page 204](#) for more details on configuring that file.

Install the Swarm triggers package on the server hosting your Helix Core Server (run this command as root):

```
yum install helix-swarm-triggers
```

5. **Optional:** Install the Swarm optional package, on the server hosting Swarm.

While not required, installing this package installs the dependencies required to use the ImageMagick and LibreOffice Swarm modules. These modules provide previews of a variety of image and office documents. Run this command as root:

```
yum install helix-swarm-optional
```

6. Complete the ["Post-installation configuration" on page 138](#) steps.

Install Swarm on RHEL 9

Important

- We recommend that the latest stable release is used. At the time of this Swarm release, the latest stable RHEL 9 release is 9.2.
- Swarm does not support Helix Core Servers configured to use **P4AUTH**, see [Centralized authentication server \(P4AUTH\)](#) in the *Helix Core Server Administrator Guide*.

Note

- Swarm can be connected to Helix Core Servers (P4D) and commit servers.
To configure Swarm to connect to more than one Helix Core Server (P4D), see ["Multiple Helix Core Server instances"](#) on page 646.
To configure Swarm to connect to a Helix Core Server configured to use *commit-edge architecture*, see ["Commit-edge deployment"](#) on page 587.
- Swarm must not be connected to Helix Broker, Helix Proxy, Helix Edge, forwarding replica, or read-only replica servers.

Installation

Important

- Review the runtime dependencies before you install Swarm, see ["Runtime dependencies"](#) on page 98.
- Review the PHP requirements before you upgrade Swarm, see ["PHP"](#) on page 102.
- Review the Helix Core Server requirements before you install Swarm, see ["Helix Core Server requirements"](#) on page 106.

1. Configure the Perforce package repository, on the server to host Swarm and on the server hosting your Helix Core Server.

Important

If the server hosting your Helix Core Server cannot use packages, for example when it is running Windows, skip this step on that server.

As root, create the file `/etc/yum.repos.d/perforce.repo` with the following content:

```
[Perforce]
name=Perforce
baseurl=http://package.perforce.com/yum/rhel/9/x86_64/
```

```
enabled=1  
gpgcheck=1
```

2. Import the Perforce package signing key, on the server to host Swarm and the server hosting your Helix Core Server.

Important

If the server hosting your Helix Core Server cannot use packages, for example when it is running Windows, skip this step on that server.

Run the following command as root:

```
rpm --import https://package.perforce.com/perforce.pubkey
```

For information about how to verify the authenticity of the signing key, see:
<https://www.perforce.com/perforce-packages>

3. Install the main Swarm package on the server to host Swarm (run these commands as root):
 - a. Run an upgrade for PHP, this will also upgrade the Swarm packages:

```
dnf update
```

- b. Install Swarm.

```
yum install helix-swarm
```

Important

When the Swarm installation has completed, you are prompted to run the **configure-swarm.sh** post-installation script.

Do not run this script until you have completed the rest of these **Installation** instructions. Instructions for running the **configure-swarm.sh** post-installation script are in the **Post-Installation configuration** section referenced in the final step of the **Installation** instructions.

Note

The firewall configuration may need to be adjusted to allow access to the web server.

```
sudo firewall-cmd --permanent --add-service=http  
sudo systemctl reload firewalld
```

If you subsequently wish to enable HTTPS, run (as root):

```
sudo firewall-cmd --permanent --add-service=https  
sudo systemctl reload firewalld
```

4. Swarm needs to know about some Helix Core Server events to operate correctly. Use Helix Core Server Extensions (recommended) or Helix Core Server Triggers to notify Swarm about these events. The Swarm Helix Core Server extension can be installed automatically by the Swarm `configure-swarm.sh` post-installation script, but Triggers must be manually installed. To use Helix Core Server Extensions, ignore this step and skip to the next step.

Trigger installation only (not recommended): Install the Swarm triggers package on the server hosting your Helix Core Server. This might be the server hosting Swarm or elsewhere on your network.

Important

- If the server hosting your Helix Core Server cannot use packages, for example when it is running Windows, you need to copy the appropriate Swarm trigger script from `/opt/perforce/swarm/p4-bin/scripts` to the server hosting your Helix Core Server. The `swarm-trigger.pl` is for both Linux and Windows systems. Once copied, the trigger script needs to be configured. See ["Installing Triggers" on page 204](#) for details.
- The package installs a config file at `/opt/perforce/etc/swarm-trigger.conf` that you will need to modify. See ["Installing Triggers" on page 204](#) for more details on configuring that file.

Install the Swarm triggers package on the server hosting your Helix Core Server (run this command as root):

```
yum install helix-swarm-triggers
```

5. **Optional:** Install the Swarm optional package, on the server hosting Swarm.

While not required, installing this package installs the dependencies required to use the ImageMagick and LibreOffice Swarm modules. These modules provide previews of a variety of image and office documents. Run this command as root:

```
yum install helix-swarm-optional
```

6. Complete the ["Post-installation configuration "](#) below steps.

Post-installation configuration

Important

The following notes are applicable for CentOS/RHEL7, RHEL 8, and RHEL 9:

- **Swarm 2019.1 to 2020.1:** these versions of Swarm used PHP packages from the SCL repositories for CentOS/RHEL 7. This was to provide access to more recent versions of PHP than are available as standard on CentOS/RHEL. This required the use of the `httpd24-httpd` package for Apache. These were all installed into `/opt/rh`

Swarm 2020.2 and later: these versions of Swarm uses the Remi repository for CentOS/RHEL 7, RHEL 8, and RHEL 9. This provides PHP 7.4 installed in the standard file system structure. This means that the old `httpd24-httpd` version of Apache is no longer needed, and the standard system version of Apache is being used again.

The SCL Apache site configuration file was installed at this location for Swarm 2019.1 to 2020.1:

```
/opt/rh/httpd24/root/etc/httpd/conf.d/perforce-swarm-site.conf
```

If this exists when Swarm is upgraded to 2020.2 and later, this file is copied to `/etc/httpd/conf.d/perforce-swarm-site.conf` if there is no file at the destination. It is also re-written to change references from `/var/log/httpd24` to `/var/log/httpd`

If a site configuration file for Swarm already exists in `/etc/httpd`, the copy and re-write is not performed.

After upgrade, `httpd24-httpd` is disabled.

- To avoid seeing the Apache HTTP server Linux test page when you start the Apache server, comment out the content of the `welcome.conf` file located in the `/etc/httpd/conf.d/` directory.
- To avoid loading the Apache HTTP server example configuration instead of the Swarm configuration when the Apache server starts, rename the `autoindex.conf` file located in the `/etc/httpd/conf.d/` directory to `z-autoindex.conf` or similar. This is required because Apache runs the first conf file it finds in the `/etc/httpd/conf.d/` directory (alphabetical order) and that must be the `perforce-swarm-site.conf` file.

Once the helix-swarm package has been installed, additional configuration is required. Perform the following steps:

1. Use the Swarm post-installation script to configure Swarm on the server hosting Swarm.

Important

- If your Helix Core Server is configured for Helix Authentication Service, the Helix Core Server must be temporarily configured to allow fall-back to passwords while you establish a connection to the Helix Core Server. Run the following command on the Helix Core Server to enable fall-back to passwords:

```
p4 configure set auth.sso.allow.passwd=1
```

- To use Helix Core Server Extensions, make sure that Extensions are installed and configured on your Helix Core Server before running the configuration script. For information about Helix Core Server Extensions, see [p4 extension](#) in *Helix Core Command-Line (P4) Reference* and [Extensions overview](#) in *Helix Core Extensions Developer Guide*.

Note

The Swarm post-installation configuration script can be used in a few different ways. The steps below outline the most straightforward configuration using an interactive install, but you can review the options by running:

```
sudo /opt/perforce/swarm/sbin/configure-swarm.sh -h
```

Run the interactive post-installation configuration script:

```
sudo /opt/perforce/swarm/sbin/configure-swarm.sh
```

The configuration script displays the following summary:

```
-----
configure-swarm.sh: Thu Aug 26 11:29:49 PDT 2021: commencing
configuration of Swarm
Summary of arguments passed:
Interactive?      [yes]
Force?           [no]
P4PORT           [(not specified, will prompt)]
Swarm user       [(not specified, will prompt, will suggest
swarm)]
Swarm password   [(not specified, will prompt)]
Email host       [(not specified, will prompt)]
Create depot     [(not specified, will prompt)]
Bypass Lock      [(not specified)]
Use Extensions?  [(not specified, will prompt)]
Swarm host       [(not specified, will prompt, will suggest
myhost)]
Swarm port       [(default (80))]
Swarm base URL   [(default (empty))]
Create Swarm user? [yes]
Super user       [(not specified, will prompt)]
Super password   [(not specified, will prompt)]
```


2. Provide information to the configuration script.

After the summary, the configuration script prompts for the following information:

- a. Specify a value for **P4PORT** in the form: *my-helix-core-server:1666*

```
No P4PORT specified
```

```
Swarm requires a connection to a Helix Core Server.  
Please supply the P4PORT to connect to.
```

```
Helix Core Server address (P4PORT):
```

Specify the hostname and port for your Helix Core Server. If defined, the value for P4PORT is used as the default. The configuration script verifies that it can connect:

```
-response: [myp4host:1666]
```

```
Checking P4PORT [myp4host:1666]...
```

```
-P4 command line to use: [/opt/perforce/bin/p4 -p  
myp4host:1666]
```

```
Attempting connection to [myp4host:1666]...
```

```
-connection successful:
```

```
  Server address: myp4host:1666
```

```
  Server version: P4D/LINUX26X86_64/2021.1/2179737  
(2021/04/24)
```

```
  Server license: 10000 users (support ends 2022/05/16)
```

```
  Server license-ip: 192.168.0.1
```

Important

If your Helix Core Server is deployed using the commit-edge architecture, ensure that the Swarm port value points to the commit server.

For more information, see the [Commit-edge](#) chapter in the *Helix Core Server Administrator Guide*.

- b. Specify the userid and password of a normal user with admin-level privileges in the Helix Core Server.

```
Checking Swarm user credentials...
```

```
No Swarm user specified
```

Swarm requires a Helix user account with 'admin' rights. Please provide a username and password for this account. If this account does not have 'admin' rights, it will be set for this user.

Helix username for the Swarm user [*swarm*]:

Enter the userid. The default is *swarm*.

Note

To allow users to clean up reviews created by other users when the review is committed, the Helix Core Server user account must have *super* permissions. See ["Review cleanup" on page 678](#).

-response: [*swarm*]

Helix password or login ticket for the Swarm user (typing hidden):

Enter the login ticket, or password, for the userid.

Important

You must use a long-lived login ticket for the Swarm user.

Note

You can obtain a login ticket by running (in another shell):

```
p4 -p myp4host:1666 -u userid login -p
```

If the login ticket you provide expires in less than a year, you will receive a warning. If you receive this warning, ask your Helix Core Server administrator to make sure the **swarm** user is in a Helix Core Server group that has the **Timeout** field set to a year or more. See [p4 group](#) in the [Helix Core Command-Line \(P4\) Reference](#).

Checking Swarm user credentials...

-checking if user [*swarm*] exists in [*myp4host:1666*]...

-user exists

Obtaining Helix login ticket for [*swarm*] in [*myp4host:1666*]...

-login ticket obtained

Checking user [*swarm*]'s ticket against [*myp4host:1666*]...

-login ticket is good

```
Checking user [swarm] has at least access level [admin]...  
-user has maximum access level [admin]  
-user meets minimum access level [admin]
```

c. Specify the hostname for the Swarm UI.

swarm needs a distinct hostname that users can enter into their browsers to access Swarm. Ideally, this is a fully-qualified domain name, for example 'swarm.company.com', but it can be just a hostname, for example 'swarm'.

Whatever hostname you provide should be Swarm-specific and not shared with any other web service on this host.

Note that the hostname you specify typically requires configuration in your network's DNS service. If you are merely testing Swarm, you can add a hostname->IP mapping entry to your computer's hosts configuration.

Hostname for this Swarm server [myhost]:

Note

The default is the current hostname. The configuration script does not verify that the hostname actually works (DNS configuration may not exist yet).

d. Specify a mail relay host or leave empty to use your local mail handler.

Swarm can use a mail relay host to send email notifications. Leave empty if you want to use the local mail handler (for example Sendmail, Postfix etc), or enter a hostname (for example mx.yourdomain.com) to use a

```
relay host.
```

```
Mail relay host:
```

Note

The configuration script does not verify that the mail relay host you provide actually accepts SMTP connections.

- e. To enable file attachments on comments, Swarm must have a depot location to save the attachment files in. Swarm can create this depot and set the depot protections for you automatically or you can manually set it up later.

Important

You must be a user with *super* user permissions to create the **.swarm** depot and set protections.

```
- checking depot
```

```
Swarm has the ability to store attachments against review
comments. To
do this, it needs to have a depot where they are stored. By
default,
this is //.swarm. We can create the depot for you
automatically, and
set the protections on it to the following:
```

```
list user * * -//.swarm/...
admin user swarm * //.swarm/...
super user super * //.swarm/...
```

```
If you want to enable attachments, and for the depot and
protections
to be set up for you, then say yes here. If you say no, then
you can
still do this manually later.
```

```
Do you want to create a .swarm depot and set protections?
(y/n) [n]
```

When prompted to automatically create the depot, select one of the following:

- Type **y** to automatically create the **.swarm** depot and set protections.
- Type **n** skip depot creation. You can create the depot and set the protections manually later if you want to, see ["Comment attachments" on page 580](#).

If you choose to automatically create the depot, the script will create the depot, set the protections, and report when completed.

```
Depot //.swarm successfully created. Protections have been set
so
that only the 'swarm' and 'super' have permissions to
access it directly.
```

- f. Swarm needs to know about the exclusive locks to work with exclusively opened files. To enable exclusive locks, set the **filetype.bypasslock** to 1. For more information about exclusive locks, see ["Handling Exclusive Locks" on page 238](#).

Tip

If this setting is not enabled in Helix Core Server, Swarm will report exceptions when working with exclusively opened files similar to **Cannot unshelve review (x). One or more files are exclusively open**, and that you must have the **filetype.bypasslock** configurable enabled.

- g. Swarm needs to know about some Helix Core Server events to operate correctly. Use Helix Core Server Extensions (recommended) or Helix Core Server Triggers to notify Swarm about these events. Helix Core Server Extensions are easier to install and maintain than triggers.

Important

You must be a user with *super* user permissions to install and configure Helix Core Server Extensions.

```
Do you want to use Swarm's Helix Core server extension?
Configuring Server extensions requires super user access to
the Helix Server.
If you install the Swarm server extension, do not install the
Swarm triggers.
Server extensions are supported for:
```

- * Linux: Helix server 19.2 and later.
- * Windows: Helix server 21.2 and later.

Use server extensions?

When prompted to **Use server Extensions**, choose one of the following:

- **Recommended:** Type **y** to use Helix Core Server Extensions. The configuration script will try and install and configure the Swarm Helix Core Server extension.
- Type **n** to use Helix Core Server Triggers. Triggers must be installed manually, Swarm will prompt you to do this when the Swarm configuration script completes.

Note

If you choose to install the Swarm Helix Core Server Extensions, the script will:

- check your Helix Core Server supports Helix Core Server Extensions
- check if Helix Core Server Extensions are installed and configured on your Helix Core Server. If they are, Swarm does not need to do anything
- install and configure the Swarm Helix Core Server extension on your Helix Core Server

If any of these checks fail, Swarm will not install the Swarm Helix Core Server extension and will report the issues on the configuration summary screen.

Once all of the information has been provided, the configuration script configures Swarm. When it has completed the configuration, the configuration summary screen is displayed, for example:

```
.....
-restarting Apache...
-Apache restarted
configure-swarm.sh: Thu Aug 26 11:31:36 PDT 2021: completed
configuration of Helix Swarm

::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::
::::::::::::
::
::  Swarm is now configured and available at:
::
::      http://myhost/
::
::  You may login as the Swarm user [swarm] using the password
::  you specified.
::
::  Ensure that you have configured the Swarm hostname in your
::  network's DNS, or have added an IP address-to-hostname
::  mapping to your computer's hosts configuration so that you
::  can access Swarm.
::
::  Server side extensions are installed and configured
::  on your P4D server.
::
::  Documentation for optional post-install configuration,
such as
::  configuring Swarm to use HTTPS, operate in a sub-folder,
or on a
::  custom port, is available:
::
```

```

::
https://www.perforce.com/perforce/doc.current/manuals/swarm/setu
p.post.html
::

.....
:

```

Note

If you have installed Swarm on a host that does not provide other web services, you may wish to disable Apache's default site configuration. Doing so means that regardless of the hostname a user might use to reach the web server hosting Swarm, Swarm would be presented.

Be aware that disabling Apache's default site configuration could disable existing web services or content.

Disabling Apache's default site configuration on Ubuntu hosts is easy. Run:

```
$ sudo a2dissite 000-default
```

For CentOS hosts, or for non-standard Apache installations, you would need to manually adjust the Apache configuration. Such changes require familiarity with Apache configuration; for more details, see: <https://httpd.apache.org/docs/2.4/configuring.html>

To add Swarm as an HTML tab in the Helix Core Visual Client (P4V), see [HTML Tabs](#) section in the *P4VJS Developer Guide*.

3. Configure SELinux for Swarm, see ["SELinux configuration"](#) on the next page.
4. The basic Swarm configuration is now complete.

Important

If your Helix Core Server is configured for Helix Authentication Service, you can force all of your users to authenticate via your Identity Provider (IdP) by disabling fall-back to passwords. To disable fall-back to passwords on the Helix Core Server, run the following command:

```
p4 configure set auth.sso.allow.passwd=0
```

5. Do one of the following:
 - **If the Swarm Helix Core Server extension was installed by the configuration script:** Review the post-install configuration options to customize your Swarm installation, see ["Post-install configuration options"](#) on page 224.

- **If the Swarm Helix Core Server extension was not installed, you must install triggers:** Configure Helix Core Server for Swarm, see ["Installing Triggers" on page 204](#).

SELinux configuration

Swarm supports SELinux which is an advanced access control mechanism that improves security for Linux distributions.

SELinux operates in one of three modes:

- **enforcing:** this mode blocks and logs any actions that do not match the defined security policy. This is the default mode for SELinux.
- **permissive:** this mode logs actions that do not match the defined security policy but these actions are not blocked.
- **disabled:** in this mode SELinux is off, actions are not blocked and are not logged.

Tip

To check the mode SELinux is operating in, view the `/etc/selinux/config` file with `vi` or a similar editor:

```
root $ vi /etc/selinux/config
```

SELinux must be configured to enable it to work correctly with Swarm, these configuration steps are shown below.

Note

You must complete the Helix Swarm package [Installation](#) steps, and the ["Post-installation configuration"](#) on page 138 steps before configuring SELinux.

Configure SELinux to enforcing mode

Run the following commands as root:

1. Install the package that contains the **semanage** configuration tool, this is used to configure SELinux:

Follow the instructions for your OS distribution:

- **CentOS/RHEL 7:** Install the `policycoreutils-python` package:

```
root $ yum install policycoreutils-python
```
- **RHEL 8:** Install the `policycoreutils-python-utils` package:

```
root $ yum install policycoreutils-python-utils
```
- **RHEL 9:** Install the `policycoreutils-python-utils` package:

```
root $ yum install policycoreutils-python-utils
```

2. Check the current SELinux mode:

```
root $ getenforce
```

3. SELinux will report its mode as; **enforcing**, **permissive**, or **disabled**.
 - a. If the mode is not set correctly edit the **/etc/selinux/config** file with vi or a similar editor.

```
root $ vi /etc/selinux/config
```

- b. Edit the config file so that **SELinux=** is set to **enforcing**.
 - c. Save the config file.
 - d. Reboot the server to complete the SELinux mode change.
4. Allow content in **/opt/perforce/swarm** to be read and written by the **httpd** process:

```
root $ semanage fcontext -a -t httpd_sys_rw_content_t
"/opt/perforce/swarm(/.*)?"
root $ restorecon -R /opt/perforce/swarm
```

5. Allow the **httpd** process to connect to other networked services, for example P4D and Redis:

```
root $ setsebool -P httpd_can_network_connect 1
```

6. Allow comment attachment thumbnails to be created:

```
root $ setsebool -P httpd_tmp_exec 1
```

7. Allow the files in **p4-bin** to be executed by the **httpd** process:

```
root $ semanage fcontext -a -t httpd_sys_script_exec_t
'/opt/perforce/swarm/p4-bin(/.*)?'
root $ restorecon -R -v /opt/perforce/swarm/p4-bin
```

8. Remove the executable constraints on Redis, allowing it to be started by **systemd** at boot time:

```
root $ semanage fcontext -a -t bin_t
/opt/perforce/swarm/sbin/redis-server-swarm
root $ restorecon -v /opt/perforce/swarm/sbin/redis-server-swarm
```

9. Restart the system:

```
root $ systemctl restart httpd
```

10. Check that you can log in to Swarm.
11. **Only if required:** Relabel your filesystem, see note before relabeling:

Important

Relabeling your file system can be a time consuming process, it is recommended that you only do this if you need to. This depends entirely on your SELinux setup, Perforce cannot give you advice on this.

```
root $ touch /.autorelabel
```

12. Reboot the server.
13. Check that you can log in to Swarm.
14. SELinux is now configured for Swarm.

Note

If you can not log in to Swarm it is possible that SELinux is blocking Swarm because its configuration is incorrect. You will need to troubleshoot the SELinux configuration to find any issues.

Install the **setroubleshoot** package, this contains **sealert** which is used when troubleshooting SELinux:

```
root $ yum install setroubleshoot
```

sealert helps you to interpret the contents of the **audit.log**. Run the following command:

```
root $ sealert -a /var/log/audit/audit.log
```

Error message: If you see an error message with a title similar to the message below, it may be because you are running CentOS/RHEL on a Virtual Machine (VM).

```
root $ SELinux is preventing /usr/sbin/ldconfig from write access
on the directory etc.
```

Install **open-vm-tools** on the VM and reboot the VM.

```
root $ yum install open-vm-tools
```

Configure SELinux permissive or disabled mode

Run the following as root:

1. Check the current SELinux mode:


```
root $ getenforce
```
2. SELinux will report its mode as; **enforcing**, **permissive**, or **disabled**.
 - a. If the mode is not set correctly edit the **/etc/selinux/config** file with vi or a similar editor.


```
root $ vi /etc/selinux/config
```
 - b. Edit the config file so that **SELinux=** is set to **permissive** or **disabled** as required.

- c. Save the config file.
- d. Reboot the server to complete the SELinux mode change.
3. Check that you can log in to Swarm.
4. SELinux is now configured for Swarm.

Upgrading

See ["Upgrading a CentOS/RHEL package installation" on page 254](#) for details.

Uninstall

1. Depending on how you have configured your Helix Core Server events, do one of the following:
 - Uninstall the SwarmHelix Core Server extension. To do this, run the following command on your Helix Core Server:


```
p4 extension --delete Perforce::helix-swarm --yes
```
 - Uninstall the Swarm triggers. As a *super* user, run the `p4 triggers` command from your Helix Core Server and manually remove all the Swarm trigger code lines.
2. Uninstall the Swarm main package, the Swarm triggers package, and the Swarm optional package. The following example assumes that the packages are all installed on the same server (run the command as root):


```
yum remove helix-swarm helix-swarm-triggers helix-swarm-optional
```
3. Remove the Swarm trigger scripts from the server hosting your Helix Core Server.

Important

If you manually installed the trigger script, perhaps because the server hosting your Helix Core Server cannot use packages (e.g. Windows), manually remove the script.

Install and configure Swarm on Amazon Linux 2

Swarm packages can be used to install Swarm on Amazon Linux 2. Using packages greatly simplifies the installation, updating, and removal of software, because the tools that manage packages are aware of the dependencies for each package.

Important

We recommend that the latest stable release of Amazon Linux 2 is used.

Important

Swarm does not support Helix Core Servers configured to use **P4AUTH**, see [Centralized authentication server \(P4AUTH\)](#) in the *Helix Core Server Administrator Guide*.

Note

- Swarm can be connected to Helix Core Servers (P4D) and commit servers.
To configure Swarm to connect to more than one Helix Core Server (P4D), see ["Multiple Helix Core Server instances"](#) on page 646.
To configure Swarm to connect to a Helix Core Server configured to use *commit-edge architecture*, see ["Commit-edge deployment"](#) on page 587.
- Swarm must not be connected to Helix Broker, Helix Proxy, Helix Edge, forwarding replica, or read-only replica servers.

Installation

Important

- Review the runtime dependencies before you install Swarm, see ["Runtime dependencies"](#) on page 98.
- Review the PHP requirements before you upgrade Swarm, see ["PHP"](#) on page 102.
- Review the Helix Core Server requirements before you install Swarm, see ["Helix Core Server requirements"](#) on page 106.

1. Configure the Perforce package repository, on the server to host Swarm and on the server hosting your Helix Core Server.

Important

If the server hosting your Helix Core Server cannot use packages, for example when it is running Windows, skip this step on that server.

As root, create the file `/etc/yum.repos.d/perforce.repo` with the following content:

```
[Perforce]
name=Perforce
baseurl=http://package.perforce.com/yum/rhel/7/x86_64/
enabled=1
gpgcheck=1
```

2. Import the Perforce package signing key, on the server to host Swarm and the server hosting your Helix Core Server.

Important

If the server hosting your Helix Core Server cannot use packages, for example when it is running Windows, skip this step on that server.

Run the following command as root:

```
rpm --import https://package.perforce.com/perforce.pubkey
```

For information about how to verify the authenticity of the signing key, see:

<https://www.perforce.com/perforce-packages>

3. Install the main Swarm package on the server to host Swarm.

Run the following commands as root:

- a. Deploy the EPEL repository configuration package to give Swarm access to the EPEL packages.

```
yum install https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm
```

- b. Install the `yum-utils` package to give access to the `yum-config-manager` command:

```
yum install yum-utils
```

- c. Install the Amazon Extras repository if it is not already installed:

```
yum install -y amazon-linux-extras
```

- d. Enable PHP 7.4 from Amazon Linux Extras:

```
amazon-linux-extras enable php7.4
```

- e. Install Swarm and accept the prompts to import the GPG keys for EPEL when requested:

```
yum install helix-swarm
```

Important

When the Swarm installation has completed, you are prompted to run the `configure-swarm.sh` post-installation script.

Do not run this script until you have completed the rest of these **Installation** instructions. Instructions for running the `configure-swarm.sh` post-installation script are in the **Post-Installation configuration** section referenced in the final step of the **Installation** instructions.

Note

The firewall configuration may need to be adjusted to allow access to the web server.

```
sudo firewall-cmd --permanent --add-service=http
sudo systemctl reload firewalld
```

If you subsequently wish to enable HTTPS, run (as root):

```
sudo firewall-cmd --permanent --add-service=https
sudo systemctl reload firewalld
```

4. Swarm needs to know about some Helix Core Server events to operate correctly. Use Helix Core Server Extensions (recommended) or Helix Core Server Triggers to notify Swarm about these events. The Swarm Helix Core Server extension can be installed automatically by the Swarm **configure-swarm.sh** post-installation script, but Triggers must be manually installed. To use Helix Core Server Extensions, ignore this step and skip to the next step.

Trigger installation only (not recommended): Install the Swarm triggers package on the server hosting your Helix Core Server. This might be the server hosting Swarm or elsewhere on your network.

Important

If the server hosting your Helix Core Server cannot use packages, for example when it is running Windows, you need to copy the appropriate Swarm trigger script from **/opt/perforce/swarm/p4-bin/scripts** to the server hosting your Helix Core Server. The **swarm-trigger.pl** is for both Linux and Windows systems. Once copied, the trigger script needs to be configured. See ["Installing Triggers" on page 204](#) for details.

Run this command as root:

```
yum install helix-swarm-triggers
```

Important

The package installs a config file at **/opt/perforce/etc/swarm-trigger.conf** that you will need to modify. See ["Installing Triggers" on page 204](#) for more details on configuring that file.

5. **Optional:** While not required, installing the optional Swarm package installs the dependencies required to use the LibreOffice Swarm module. This enables Swarm to preview office documents.
 - a. Enable LibreOffice in Amazon Linux Extras with:


```
amazon-linux-extras enable libreoffice
```

- b. Install the Swarm optional package, on the server hosting Swarm (run the command as root):

```
yum install helix-swarm-optional
```

6. Complete the "Post-installation configuration" below steps.

Post-installation configuration

Important

- To avoid seeing the Apache HTTP server Linux test page when you start the Apache server, comment out the content of the `welcome.conf` file located in the `/etc/httpd/conf.d/` directory.
- To avoid loading the Apache HTTP server example configuration instead of the Swarm configuration when the Apache server starts, rename the `autoindex.conf` file located in the `/etc/httpd/conf.d/` directory to `z-autoindex.conf` or similar. This is required because Apache runs the first conf file it finds in the `/etc/httpd/conf.d/` directory (alphabetical order) and that must be the `perforce-swarm-site.conf` file.

Once the helix-swarm package has been installed, additional configuration is required. Perform the following steps:

1. Use the Swarm post-installation script to configure Swarm on the server hosting Swarm.

Important

- If your Helix Core Server is configured for Helix Authentication Service, the Helix Core Server must be temporarily configured to allow fall-back to passwords while you establish a connection to the Helix Core Server. Run the following command on the Helix Core Server to enable fall-back to passwords:

```
p4 configure set auth.sso.allow.passwd=1
```

- To use Helix Core Server Extensions, make sure that Extensions are installed and configured on your Helix Core Server before running the configuration script. For information about Helix Core Server Extensions, see [p4 extension](#) in *Helix Core Command-Line (P4) Reference* and [Extensions overview](#) in *Helix Core Extensions Developer Guide*.

Note

The Swarm post-installation configuration script can be used in a few different ways. The steps below outline the most straightforward configuration using an interactive install, but you can review the options by running:

```
sudo /opt/perforce/swarm/sbin/configure-swarm.sh -h
```

Run the interactive post-installation configuration script:


```
sudo /opt/perforce/swarm/sbin/configure-swarm.sh
```

The configuration script displays the following summary:

```
-----
configure-swarm.sh: Thu Aug 26 11:29:49 PDT 2021: commencing
configuration of Swarm
Summary of arguments passed:
Interactive?      [yes]
Force?           [no]
P4PORT           [(not specified, will prompt)]
Swarm user       [(not specified, will prompt, will suggest
swarm)]
Swarm password   [(not specified, will prompt)]
Email host       [(not specified, will prompt)]
Create depot     [(not specified, will prompt)]
Bypass Lock      [(not specified)]
Use Extensions?  [(not specified, will prompt)]
Swarm host       [(not specified, will prompt, will suggest
myhost)]
Swarm port       [(default (80))]
Swarm base URL   [(default (empty))]
Create Swarm user? [yes]
Super user       [(not specified, will prompt)]
Super password   [(not specified, will prompt)]
```

2. Provide information to the configuration script.

After the summary, the configuration script prompts for the following information:

- a. Specify a value for **P4PORT** in the form: *my-helix-core-server:1666*

```
No P4PORT specified
```

```
Swarm requires a connection to a Helix Core Server.
Please supply the P4PORT to connect to.
```

```
Helix Core Server address (P4PORT):
```

Specify the hostname and port for your Helix Core Server. If defined, the value for P4PORT is used as the default. The configuration script verifies that it can connect:

```
-response: [myp4host:1666]

Checking P4PORT [myp4host:1666]...
-P4 command line to use: [/opt/perforce/bin/p4 -p
myp4host:1666]
Attempting connection to [myp4host:1666]...
-connection successful:
  Server address: myp4host:1666
  Server version: P4D/LINUX26X86_64/2021.1/2179737
(2021/04/24)
  Server license: 10000 users (support ends 2022/05/16)
  Server license-ip: 192.168.0.1
```

Important

If your Helix Core Server is deployed using the commit-edge architecture, ensure that the Swarm port value points to the commit server.

For more information, see the [Commit-edge](#) chapter in the *Helix Core Server Administrator Guide*.

- b. Specify the userid and password of a normal user with admin-level privileges in the Helix Core Server.

```
Checking Swarm user credentials...
No Swarm user specified

Swarm requires a Helix user account with 'admin' rights.
Please provide a username and password for this account.
If this account does not have 'admin' rights, it will
be set for this user.

Helix username for the Swarm user [swarm]:
Enter the userid. The default is swarm.
```

Note

To allow users to clean up reviews created by other users when the review is committed, the Helix Core Server user account must have *super* permissions. See ["Review cleanup" on page 678](#).

```
-response: [swarm]
```

Helix password or login ticket for the Swarm user (typing hidden):

Enter the login ticket, or password, for the userid.

Important

You must use a long-lived login ticket for the Swarm user.

Note

You can obtain a login ticket by running (in another shell):

```
p4 -p myp4host:1666 -u userid login -p
```

If the login ticket you provide expires in less than a year, you will receive a warning. If you receive this warning, ask your Helix Core Server administrator to make sure the **swarm** user is in a Helix Core Server group that has the **Timeout** field set to a year or more. See [p4 group](#) in the *Helix Core Command-Line (P4) Reference*.

```
Checking Swarm user credentials...
```

```
-checking if user [swarm] exists in [myp4host:1666]...
```

```
-user exists
```

```
Obtaining Helix login ticket for [swarm] in [myp4host:1666]...
```

```
-login ticket obtained
```

```
Checking user [swarm]'s ticket against [myp4host:1666]...
```

```
-login ticket is good
```

```
Checking user [swarm] has at least access level [admin]...
```

```
-user has maximum access level [admin]
```

```
-user meets minimum access level [admin]
```

c. Specify the hostname for the Swarm UI.

swarm needs a distinct hostname that users can enter into their browsers to access Swarm. Ideally, this is a fully-qualified domain name,

```
for example  
'swarm.company.com', but it can be just a hostname, for  
example 'swarm'.
```

Whatever hostname you provide should be Swarm-specific and not shared with any other web service on this host.

Note that the hostname you specify typically requires configuration in your network's DNS service. If you are merely testing Swarm, you can add a hostname->IP mapping entry to your computer's hosts configuration.

Hostname for this Swarm server [myhost]:

Note

The default is the current hostname. The configuration script does not verify that the hostname actually works (DNS configuration may not exist yet).

- d. Specify a mail relay host or leave empty to use your local mail handler.

```
Swarm can use a mail relay host to send email notifications.  
Leave empty if  
you want to use the local mail handler (for example Sendmail,  
Postfix etc),  
or enter a hostname (for example mx.yourdomain.com) to use a  
relay host.
```

Mail relay host:

Note

The configuration script does not verify that the mail relay host you provide actually accepts SMTP connections.

- e. To enable file attachments on comments, Swarm must have a depot location to save the attachment files in. Swarm can create this depot and set the depot protections for you automatically or you can manually set it up later.

Important

You must be a user with *super* user permissions to create the **.swarm** depot and set protections.

```
- checking depot
```

```
Swarm has the ability to store attachments against review
comments. To
do this, it needs to have a depot where they are stored. By
default,
this is //.swarm. We can create the depot for you
automatically, and
set the protections on it to the following:
```

```
list user * * -//.swarm/...
admin user swarm * //.swarm/...
super user super * //.swarm/...
```

```
If you want to enable attachments, and for the depot and
protections
to be set up for you, then say yes here. If you say no, then
you can
still do this manually later.
Do you want to create a .swarm depot and set protections?
(y/n) [n]
```

When prompted to automatically create the depot, select one of the following:

- Type **y** to automatically create the **.swarm** depot and set protections.
- Type **n** skip depot creation. You can create the depot and set the protections manually later if you want to, see ["Comment attachments" on page 580](#).

If you choose to automatically create the depot, the script will create the depot, set the protections, and report when completed.

```
Depot // .swarm successfully created. Protections have been set
so
that only the 'swarm' and 'super' have permissions to
access it directly.
```

- f. Swarm needs to know about the exclusive locks to work with exclusively opened files. To enable exclusive locks, set the `filetype.bypasslock` to 1. For more information about exclusive locks, see ["Handling Exclusive Locks" on page 238](#).

Tip

If this setting is not enabled in Helix Core Server, Swarm will report exceptions when working with exclusively opened files similar to **Cannot unshelve review (x). One or more files are exclusively open**, and that you must have the `filetype.bypasslock` configurable enabled.

- g. Swarm needs to know about some Helix Core Server events to operate correctly. Use Helix Core Server Extensions (recommended) or Helix Core Server Triggers to notify Swarm about these events. Helix Core Server Extensions are easier to install and maintain than triggers.

Important

You must be a user with *super* user permissions to install and configure Helix Core Server Extensions.

```
Do you want to use Swarm's Helix Core server extension?
Configuring Server extensions requires super user access to
the Helix Server.
```

```
If you install the Swarm server extension, do not install the
Swarm triggers.
```

```
Server extensions are supported for:
```

- * Linux: Helix server 19.2 and later.
- * Windows: Helix server 21.2 and later.

```
Use server extensions?
```

When prompted to **Use server Extensions**, choose one of the following:

- **Recommended:** Type **y** to use Helix Core Server Extensions. The configuration script will try and install and configure the Swarm Helix Core Server extension.

- Type **n** to use Helix Core Server Triggers. Triggers must be installed manually, Swarm will prompt you to do this when the Swarm configuration script completes.

Note

If you choose to install the Swarm Helix Core Server Extensions, the script will:

- check your Helix Core Server supports Helix Core Server Extensions
- check if Helix Core Server Extensions are installed and configured on your Helix Core Server. If they are, Swarm does not need to do anything
- install and configure the Swarm Helix Core Server extension on your Helix Core Server

If any of these checks fail, Swarm will not install the Swarm Helix Core Server extension and will report the issues on the configuration summary screen.

```

.....
-restarting Apache...
-Apache restarted
configure-swarm.sh: Thu Aug 26 11:31:36 PDT 2021: completed
configuration of Helix Swarm

::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::
:::::::::
::
::  Swarm is now configured and available at:
::
::      http://myhost/
::
::  You may login as the Swarm user [swarm] using the password
::  you specified.
::
::  Ensure that you have configured the Swarm hostname in your
::  network's DNS, or have added an IP address-to-hostname
::  mapping to your computer's hosts configuration so that you
::  can access Swarm.
::
::  Server side extensions are installed and configured
::  on your P4D server.
::
::  Documentation for optional post-install configuration,
such as
::  configuring Swarm to use HTTPS, operate in a sub-folder,
or on a
::  custom port, is available:
::

```



```

::
https://www.perforce.com/perforce/doc.current/manuals/swarm/se
tup.post.html
::

.....
:::

```

Note

If you have installed Swarm on a host that does not provide other web services, you may wish to disable Apache's default site configuration. Doing so means that regardless of the hostname a user might use to reach the web server hosting Swarm, Swarm would be presented.

Be aware that disabling Apache's default site configuration could disable existing web services or content.

To disable Apache's default configuration, manually adjust the Apache configuration. Such changes require familiarity with Apache configuration; for more details, see: <https://httpd.apache.org/docs/2.4/configuring.html>

To add Swarm as an HTML tab in the Helix Core Visual Client (P4V), see [HTML Tabs](#) section in the *P4VJS Developer Guide*.

3. Configure SELinux for Swarm, see ["SELinux configuration"](#) on the facing page.
4. The basic Swarm configuration is now complete.

Important

If your Helix Core Server is configured for Helix Authentication Service, you can force all of your users to authenticate via your Identity Provider (IdP) by disabling fall-back to passwords. To disable fall-back to passwords on the Helix Core Server, run the following command:

```
p4 configure set auth.sso.allow.passwd=0
```

5. Do one of the following:
 - **If the Swarm Helix Core Server extension was installed by the configuration script:** Review the post-install configuration options to customize your Swarm installation, see ["Post-install configuration options"](#) on page 224.
 - **If the Swarm Helix Core Server extension was not installed, you must install triggers:** Configure Helix Core Server for Swarm, see ["Installing Triggers"](#) on page 204.

SELinux configuration

Swarm supports SELinux which is an advanced access control mechanism that improves security for Linux distributions.

SELinux operates in one of three modes:

- **enforcing**: this mode blocks and logs any actions that do not match the defined security policy. This is the default mode for SELinux.
- **permissive**: this mode logs actions that do not match the defined security policy but these actions are not blocked.
- **disabled**: in this mode SELinux is off, actions are not blocked and are not logged.

Tip

To check the mode SELinux is operating in, view the `/etc/selinux/config` file with `vi` or a similar editor:

```
root $ vi /etc/selinux/config
```

SELinux must be configured to enable it to work correctly with Swarm, these configuration steps are shown below.

Note

You must complete the Helix Swarm package "Installation" on page 153 steps, and the "Post-installation configuration" on page 156 steps before configuring SELinux.

Configure SELinux to enforcing mode

Run the following commands as root:

1. Install the `policycoreutils-python` package that contains the `semanage` configuration tool, this is used to configure SELinux


```
root $ yum install policycoreutils-python
```
2. Check the current SELinux mode:


```
root $ getenforce
```
3. SELinux will report its mode as; `enforcing`, `permissive`, or `disabled`.
 - a. If the mode is not set correctly edit the `/etc/selinux/config` file with `vi` or a similar editor.


```
root $ vi /etc/selinux/config
```
 - b. Edit the config file so that `SELinux=` is set to `enforcing`.
 - c. Save the config file.
 - d. Reboot the server to complete the SELinux mode change.

4. Allow content in `/opt/perforce/swarm` to be read and written by the `httpd` process:

```
root $ semanage fcontext -a -t httpd_sys_rw_content_t
"/opt/perforce/swarm(/.*)?"
root $ restorecon -R /opt/perforce/swarm
```

5. Allow the `httpd` process to connect to other networked services, for example P4D and Redis:

```
root $ setsebool -P httpd_can_network_connect 1
```

6. Allow comment attachment thumbnails to be created:

```
root $ setsebool -P httpd_tmp_exec 1
```

7. Allow the files in `p4-bin` to be executed by the `httpd` process:

```
root $ semanage fcontext -a -t httpd_sys_script_exec_t
'/opt/perforce/swarm/p4-bin(/.*)?'
root $ restorecon -R -v /opt/perforce/swarm/p4-bin
```

8. Remove the executable constraints on Redis, allowing it to be started by `systemd` at boot time:

```
root $ semanage fcontext -a -t bin_t
/opt/perforce/swarm/sbin/redis-server-swarm
root $ restorecon -v /opt/perforce/swarm/sbin/redis-server-swarm
```

9. Restart the system:

```
root $ systemctl restart httpd
```

10. Check that you can log in to Swarm.

11. **Only if required:** Relabel your filesystem, see note before relabeling:

Important

Relabeling your file system can be a time consuming process, it is recommended that you only do this if you need to. This depends entirely on your SELinux setup, Perforce cannot give you advice on this.

```
root $ touch /.autorelabel
```

12. Reboot the server.

13. Check that you can log in to Swarm.

14. SELinux is now configured for Swarm.

Note

If you can not log in to Swarm it is possible that SELinux is blocking Swarm because its configuration is incorrect. You will need to troubleshoot the SELinux configuration to find any issues.

Install the `setroubleshoot` package, this contains `sealert` which is used when troubleshooting SELinux:

```
root $ yum install setroubleshoot
```

`sealert` helps you to interpret the contents of the `audit.log`. Run the following command:

```
root $ sealert -a /var/log/audit/audit.log
```

Error message: If you see an error message with a title similar to the message below, it may be because you are running on a Virtual Machine (VM).

```
root $ SELinux is preventing /usr/sbin/ldconfig from write access
on the directory etc.
```

Install `open-vm-tools` on the VM and reboot the VM.

```
root $ yum install open-vm-tools
```

Configure SELinux permissive or disabled mode

Run the following as root:

1. Check the current SELinux mode:

```
root $ getenforce
```

2. SELinux will report its mode as; `enforcing`, `permissive`, or `disabled`.

- a. If the mode is not set correctly edit the `/etc/selinux/config` file with vi or a similar editor.

```
root $ vi /etc/selinux/config
```

- b. Edit the config file so that `SELinux=` is set to `permissive` or `disabled` as required.
- c. Save the config file.
- d. Reboot the server to complete the SELinux mode change.

3. Check that you can log in to Swarm.
4. SELinux is now configured for Swarm.

Uninstall

1. Depending on how you have configured your Helix Core Server events, do one of the following:
 - Uninstall the Swarm Helix Core Server extension. To do this, run the following command on your Helix Core Server:

```
p4 extension --delete Perforce::helix-swarm --yes
```
 - Uninstall the Swarm triggers. As a *super* user, run the `p4 triggers` command from your Helix Core Server and manually remove all the Swarm trigger code lines.
2. Uninstall the Swarm main package, the Swarm triggers package, and the Swarm optional package. The following example assumes that the packages are all installed on the same server (run the command as root):

```
yum remove helix-swarm helix-swarm-triggers helix-swarm-optional
```

3. Remove the Swarm trigger scripts from the server hosting your Helix Core Server.

Important

If you manually installed the trigger script, perhaps because the server hosting your Helix Core Server cannot use packages (for example Windows), manually remove the script.

Run Swarm using a Docker container

Note

- This section assumes that you have basic knowledge of how to use a Docker container. There are numerous ways of configuring a Docker environment, and the examples outlined in this section are only to get you started. For more information on all of the options for Docker configurations, see <https://www.docker.com/>.
- You must be a *super* user if you are configuring Helix Swarm for the first time. If you are re-using an existing `config.php` file, this is not necessary.
- You need access to the internet to download the docker image. You can download an image and transfer it to a machine with no internet access. To do this, use the `docker save` and `docker load` commands. For more information on how to use these commands, see <https://www.docker.com/>.

This section details how to use a Docker container to run Helix Swarm either in a test environment or a production environment.

- "Prerequisites" on the facing page
- "Helix Swarm images" on the facing page
- "Running the Docker container" on page 171

- "Persisting Docker containers for a production environment" on page 174
- "Restarting the Docker container" on page 174
- "Migrating Swarm to a Docker container" on page 176
- "Advanced configuration options " on page 177
- "Example of running Swarm using docker-compose" on page 179

Here is an overview of setting up a Docker container to run Helix Swarm in a production environment:

1. Start Helix Swarm using a Docker container with the `.env` configuration file and a mounted Swarm or data directory.
This stores all the `.env` configuration options into the data directory in the correct format.
2. Stop the Docker container.
3. Start a Docker container using the mounted Swarm or data directory without the `.env` configuration options. The stored configurations in the `config.php` file are used.

Prerequisites

To run Helix Swarm using a Docker container, you must have the following:

- An environment capable of running Docker images.
For more information about Docker images, see <https://www.docker.com>.
- A running Helix Core Server instance.
For more information about the Helix Core Server, see <https://www.perforce.com/manuals/p4sag/Content/P4SAG/Home-p4sag.html>.

Helix Swarm images

Important

- There is a tag for each major version of Helix Swarm, for example, 2022.2. These images are updated for every patch release.
- The **latest** tag points to the very latest released version of Helix Swarm. The major version number of Helix Swarm used by the latest tag will change when new major versions are released, this includes the patches for the latest version.

You can access the Helix Swarm images on Docker Hub here:

<https://hub.docker.com/r/perforce/helix-swarm>

The following Helix Swarm images are available on the Docker Hub:

- `perforce/helix-swarm:latest`
- `perforce/helix-swarm:2022.2`

Helix Swarm image source

You can access the source code for all Helix Swarm images here:

<https://github.com/perforce/helix-swarm-docker>

Running the Docker container

Important

When running Swarm in a Docker container for production systems, your configuration must persist data when it is shutdown. For more information about persisting data and configuration in a production environment, see "[Persisting Docker containers for a production environment](#)" on [page 174](#).

This section details how you can configure a fresh installation of Helix Swarm against a Helix Core Server to get a test instance running. It does not specify how to persist data on an external volume, so should not be used for a production environment. For more information about persisting data and configuration in a production environment, see "[Persisting Docker containers for a production environment](#)" on [page 174](#).

Create a configuration file

Note

- In this section, we create an example `.env` configuration file which is then used to configure a Docker container.
- The `P4D_SUPER_PASSWD` and `SWARM_PASSWD` configurables in the `.env` configuration file must be a password, not ticket. The `SWARM_PASSWD` is converted into a host locked ticket before it is stored.
- If the `/opt/perforce/swarm/data` directory is mapped outside the Docker container, after the initial configuration run, you can remove the configuration options from the `.env` configuration file as they are stored in `config.php` file.

Create a `.env` configuration file in the current directory as follows:

```
P4D_PORT=ssl:myperforce:1666
P4D_SUPER=super
P4D_SUPER_PASSWD=HelixDockerBay94
SWARM_USER=swarm
SWARM_PASSWD=HelixDockerBay94
```

```
SWARM_HOST=mymachine
```

```
SWARM_REDIS=helix-redis
```

```
SWARM_REDIS_PORT=7379
```

```
# If set to 'y', then extensions will be installed even if they
already
```

```
# exist, overwriting existing configuration.
```

```
SWARM_FORCE_EXT=y
```

- `P4D_PORT`

The `P4D_PORT` value must be set to the Helix Core Server instance.

Important

Ensure that the Docker image is able to communicate with the Helix Core Server instance.

- `P4D_SUPER` and `P4D_SUPER_PASSWD`

The `P4D_SUPER` user must be an existing Helix Core Server user with super privileges. The `P4D_SUPER_PASSWD` must be a password and cannot be a ticket.

- `SWARM_USER` and `SWARM_PASSWD`

The `SWARM_USER` is created if it does not already exist. If it already exists, the `SWARM_USER` must have at least admin privileges. The `SWARM_PASSWD` must be set to the existing password if the `SWARM_USER` already exists or you can use it to set a new password. The `SWARM_PASSWD` must obey the normal rules for a Helix Core Server password.

- `SWARM_HOST`

The hostname set by `SWARM_HOST` should be reachable by the Helix Core Server. In this example setup, we assume that the hostname is the name of the host running the Docker image, rather than the Docker image itself. When the Docker image is started, port 80 and, optionally, port 443 are mapped to the host's port. This way anything that can reach the host can reach Helix Swarm.

Note

- With this example setup, the IP address or hostname of the actual Docker image does not matter.
- To run Swarm on a different port than port 80, you must add the port number to the `SWARM_HOST` parameter.

- `SWARM_REDIS` and `SWARM_REDIS_PORT`

This is the hostname and port for the Redis server. This example setup assumes that Redis server is running in a Docker container as setup in ["Start a single Docker container" below](#).

- `SWARM_FORCE_EXT`

When set to `y`, `SWARM_FORCE_EXT` will force the extensions to be installed or reinstalled. The existing SwarmHelix Core Serverextension configuration is overwritten.

Start a single Docker container

Running the following command will start a single instance of Helix Swarm running that uses the `.env` configuration file. It assumes that there is a Redis server and Helix Core Server to connect to. The Redis configurables in the `.env` file in ["Create a configuration file" on page 171](#) section will not work in this example. You must ensure that Helix Swarm points to an existing Redis instance.

```
docker run -d --name helix-swarm --network-alias helix-swarm \
  --env-file .env -p 80:80 perforce/helix-swarm
```

You can shut down the Helix Swarm instance as follows:

```
docker stop helix-swarm
```

You can delete the Helix Swarm instance as follows:

```
docker rm helix-swarm
```

Setup a local private network

A simplified option is to run Redis server as a Docker container. You can setup a local private network to run both, the Redis server and Helix Swarm. This makes the Redis server private to Helix Swarm.

Use the following docker commands to setup a local private network:

Note

The Redis configurables in the `.env` file created in the ["Create a configuration file" on page 171](#) section is used in the example code below:

```
docker network create helix
docker run -d --name helix-redis --network helix --network-alias
helix-redis \
redis redis-server --protected-mode no --port 7379

docker run -d --name helix-swarm --network helix --network-alias
helix-swarm \
--env-file .env -p 80:80 perforce/helix-swarm
```

Persisting Docker containers for a production environment

For a production environment, the Helix Swarm and Redis server cache data should persist even after the Docker container has stopped and deleted. You can do this by using **bind mount** or **volume mount** when starting the Docker containers. This mounts the contents of **/opt/perforce/swarm/data** directory outside the Docker container.

To preserve the Helix Swarm configuration and Redis server cache data outside the Docker container, use the following:

```
docker network create helix

mkdir -p storage/redis-data
docker run -d --name helix-redis --network helix --network-alias
helix-redis \
-v $PWD/storage/redis-data:/data redis \
redis-server --protected-mode no --port 7379 --appendonly yes

mkdir -p storage/swarm-data
docker run -d --name helix-swarm --network helix --network-alias
helix-swarm \
-p 80:80 -v $PWD/storage/swarm-data:/opt/perforce/swarm/data \
--env-file .env perforce/helix-swarm
```

Note

- The **--appendonly yes** option for the Redis server will cause it to write its cache data to disc. The **/data** directory can then be mounted externally to persist it beyond the lifetime of the Docker container.
- Helix Swarm log files, the worker queue, tokens, and workspaces are also preserved because they are part of the **/opt/perforce/swarm/data** directory.

Restarting the Docker container

Note

- If there is no **config.php** file in the **/opt/perforce/swarm/data/** directory, the **.env** configuration file is used.

- When you make any changes to the **config.php** file, you must reload the configuration cache for Swarm to use the updated configurations. If you restart a Docker container, Helix Swarm will not use the updated configuration until the configuration cache has been reloaded. For more information on how to restart configuration cache, see ["Reload Configuration button" on page 732](#) section.
- Once you have restarted the Docker container, you must be a *root* user to access files in the **/opt/perforce/swarm/data/** directory.

When you restart a Docker container, the **/opt/perforce/swarm/data/config.php** file is used to configure the system. You can restart the Docker container as follows:

```
docker run -d --name helix-swarm --network helix --network-alias
helix-swarm \
    -p 80:80 -v $PWD/storage/swarm-data:/opt/perforce/swarm/data \
    perforce/helix-swarm
```

Docker directory

When you restart the Docker container, a Docker directory is created **/opt/perforce/swarm/data/docker** for your Swarm installation. The Docker directory includes the Apache site configuration files, Swarm configuration files, Swarm version file, and the Docker container version. All the files in the Docker directory are linked to the original files in their original locations. This allows you to edit files outside the Docker container. Any change to any file in the Docker directory is preserved between restarts.

For example,

- **/etc/apache2/sites-available -> docker/sites-available**
The entire **/etc/apache2/sites-available** directory is soft linked to the **/opt/perforce/swarm/data/docker/sites-available** directory.
All the configuration files have **a2ensite** run against them. This allows you to automatically run multiple Apache site configurations.
- **/opt/perforce/swarm/public/custom -> docker/custom**
The entire **/opt/perforce/swarm/public/custom** directory is soft linked to the **docker/custom** directory.
- **/opt/perforce/etc/swarm-cron-hosts.conf -> docker/swarm-cron-hosts.conf**
The entire **/opt/perforce/etc/swarm-cron-hosts.conf** file is soft linked to the **docker/swarm-cron-hosts.conf** file.

Note

- If a **docker/sites-available** file exists instead of the directory, then no soft link is created. This allows you to mount the Apache directory outside the Docker container.
- To enable HTTPS, you must setup a new Apache Virtual Host.

Migrating Swarm to a Docker container

Important

When migrating a Swarm installation to a Docker installation on the same system, you must do the following:

- Use a host unlocked ticket.
- Stop the host cron jobs and use the cron jobs available in the Docker container.

Note

- If your Helix Swarm installation uses the default Redis server, then you must configure a new Redis instance for the dockerised Helix Swarm. To configure a new Redis instance, see ["Create a configuration file" on page 171](#).
- If your Helix Swarm installation uses a remote Redis server, you can continue using the remote Redis server or you can configure a new Redis instance. The Redis server configurations are stored in the **config.php** file. To configure a new Redis instance, see ["Create a configuration file" on page 171](#).

You can migrate an existing Swarm server setup into a Docker container. To do this, you must copy the **/opt/perforce/swarm/data** directory to a location that is accessible for the Docker container as follows:

1. Create a **./storage** directory.
2. Copy contents of the **/opt/perforce/swarm/data/** directory to **./storage/swarm-data** directory.
3. Ensure that the network connections configuration are defined correctly in the **config.php** file.
4. Ensure that there is an Apache site configuration file at **swarm-data/etc/perforce-swarm-site.conf**. This Apache site configuration file is copied into the Docker when it starts.

```
docker network create helix
```

```
mkdir -p storage/redis-data
```

```
docker run -d --name helix-redis --network helix --network-alias
```

```

helix-redis \
  -v $PWD/storage/redis-data:/data redis \
  redis-server --protected-mode no --port 7379 --appendonly yes

mkdir -p storage/swarm-data
docker run -d --name helix-swarm --network helix --network-alias
helix-swarm \
  -p 80:80 -v $PWD/storage/swarm-data:/opt/perforce/swarm/data \
  perforce/helix-swarm

```

A **docker** directory is created in the **/opt/perforce/swarm/data/** location. During the initial configuration, the **docker** directory includes the following files:

- Apache site configuration file
perforce-swarm-site.conf
- List of cron worker configuration
swarm-cron-hosts.conf
- Directory which contains any custom configurations for Swarm
custom

Advanced configuration options

Important

- When a **config.php** file exists, all configurations in the **.env** configuration file are ignored.
- When no Apache site configuration file is available, the **SWARM_HOST** parameter in the **.env** configuration file is used to configure the Apache server.
- We recommend that the **/opt/perforce/swarm/data/docker** directory is mounted outside the Docker container. So changes to any files in the Docker directory are preserved between restarts.

To modify the configuration of the Helix Swarm docker image, you can edit the following parameters in the **.env** configuration file:

Configuration	Default	Description
P4D_SUPER	super	Name of the super user.

Configuration	Default	Description
P4D_SUPER_PASSWD	<i>no default</i>	Password of the super user.
P4D_PORT	ssl:perforce:1666	Port for the Helix Core Server (P4D).
SWARM_USER	swarm	Name of the Swarm user.
SWARM_PASSWD	<i>no default</i>	Password of the Swarm user.
SWARM_MAILHOST	localhost	Mail server address
SWARM_HOST	helix-swarm	Hostname of the Swarm server.
SWARM_FORCE_EXT	n	To overwrite Swarm Extensions, set SWARM_FORCE_EXT = 'y'.
SWARM_REDIS	helix-redis	Hostname of the Redis server.
SWARM_REDIS_PORT	7379	Port number of the Redis server.
SWARM_REDIS_PASSWD		Password for the Redis server.
SWARM_REDIS_NAMESPACE		Namespace for the Redis server.

After the Swarm setup is complete, you can edit the **config.php** file to include support for Email configuration, JIRA integration, and for Multiple Helix Core (P4D) instances.

Apache server configuration

By default, the **mod_rewrite** and **mod_ssl** modules are enabled for Apache server. The SSL module is required to run HTTPS from within the docker container.

When using SSL, ensure that you have defined a location for the SSL key and the SSL certificate file. Do this in one of the following ways:

- Save the SSL key and certificate file in the **/opt/perforce/swarm/data/docker** directory and add a reference to them in the Apache site configuration file **perforce-swarm-site.conf**.

- Save the SSL key and certificate file in a different location and mount that location externally. This ensures that the private key file is outside Swarm.
- Save the SSL key and certificate file externally to the docker container. Use a web server as a proxy for the docker container. Enable the proxy web server to provide any HTTPS functionality.

Apache proxy

In cases where the Docker container is behind another web server, you may run into issues with Helix tickets unless they are host unlocked. Generally, the Apache server forwards the IP address of the client to Helix Swarm. This allows a ticket that is obtained on a client machine to work with Helix Swarm, even though Helix Swarm is on a different machine.

If Helix Swarm is behind a web proxy (such as a second Apache server, running SSL), the IP address of the client is lost. In this scenario, host unlocked tickets must be used that can be obtained using `p4 login -ap`.

Alternatively, you can use the `remoteip_module` to allow the container Apache to trust the proxy that enables forwarding the client IP addresses. To do this, add the following code into the `performce-swarm-site.conf` file, at the same level as the `DocumentRoot`.

Note

The `RemoteIPInternalProxy` should be the IP address of the host, as seen by the container.

```
RemoteIPHeader X-Forwarded-For
RemoteIPInternalProxy 172.19.0.1/24
```

Example of running Swarm using docker-compose

This section details an example usage of the docker-compose command. In the following example:

- The `docker-compose.yml` file is setup in a directory from where the commands are run.
- Creates a network for the two containers to run in.
- Sets up a Redis container to run Redis in, with external storage in the `storage` directory.
- Sets up a Swarm container with external storage for the `data` directory and the `/var/www` directory.

```
#
=====
=== #
# This compose file runs a migrated Swarm instance. Redis and Swarm
are in #
# containers, p4d is on the host machine.
#
```

```
#
=====

=== #
version: '2.2'

networks:
  default:
    driver: bridge
    ipam:
      driver: default
      config:
        - subnet: 172.19.0.0/16
          gateway: 172.19.0.1
    name: swarm-net

services:
  helix.swarm:
    image: perforce/helix-swarm:latest
    hostname: helix-swarm
    domainname: helix
    volumes:
      - ./storage/swarm-data:/opt/perforce/swarm/data
      - ./storage/www:/var/www
    ports:
      - 127.0.0.1:80:80
    working_dir: /opt/perforce/swarm
    depends_on:
      - helix.redis
    tty: false

  helix.redis:
    image: "redis"
    hostname: helix.redis
    domainname: helix
```



```
command: redis-server --protected-mode no --port 7379 --appendonly
yes

volumes:
  - ./storage/redis-data:/data
```

Install and configure Swarm manually from a Tarball

Important

- Review the runtime dependencies before you install Swarm, see ["Runtime dependencies" on page 98](#).
- **CentOS/RHEL 7:** Use the Remi repository configuration package (**remi-release-7.rpm**) to give Swarm access to PHP 7/PHP 8 and to LibreOffice which is part of the optional package install. Use the **epel-release-latest-7.noarch.rpm** repository configuration package to give Swarm access to EPEL packages.
- **RHEL 8:** Use the Remi repository configuration package (**remi-release-8.rpm**) to give Swarm access to PHP 7/PHP 8 and to LibreOffice which is part of the optional package install. Use the **epel-release-latest-8.noarch.rpm** repository configuration package to give Swarm access to EPEL packages.
- **RHEL 9:** Use the Remi repository configuration packages (**remi-release-8.rpm** and **remi-release-9.rpm**) to give Swarm access to PHP 7/PHP 8 and to LibreOffice which is part of the optional package install. Use the **epel-release-latest-8.noarch.rpm** and **epel-release-latest-9.noarch.rpm** repository configuration package to give Swarm access to EPEL packages.

1. Download the [Swarm tarball](#).
2. Expand the Swarm package (a *compressed tarball*).

From the command line, expand it via the tar command:

```
tar -zxvf swarm.tgz
```

The contents of the Swarm package are expanded into a top-level folder named **swarm-version**, where **version** corresponds to the version downloaded.

Tip

Many graphical file manager applications (Nautilus on Linux, etc.) can automatically expand the tarball package by simply double-clicking it.

3. Move the contents of the Swarm package to the correct location.

Identify a location for the Swarm files; this should correspond to a location associated to the virtual host configured under Apache (see ["Apache configuration" on page 186](#)).

```
mv /path/to/swarm-version /path/to/vhosts/swarm
```

4. Assign correct ownership and permission for the Swarm files.

The data top-level folder in the Swarm distribution needs to be writeable by the web server. To achieve this effect, simply change ownership of the data folder to the web user:

```
sudo chown -R www /path/to/vhosts/swarm/data
```

The **www** user above is an example of what the web server user name might be. Depending on your distribution, this could be **_www**, **web**, **nobody** or something else entirely.

If your web server is already running, you can discover the user with:

```
ps aux | grep -E 'apache|httpd'
root      3592  0.0  0.5 405240 20708 ?        Ss   May03   4:32
/usr/sbin/apache2 -k start
www       20016  0.0  0.2 405264  9796 ?        S    07:45   0:00
/usr/sbin/apache2 -k start
```

In this example, **www** is the user Apache is running as.

From a security perspective, we recommend that the minimum file permissions should be granted to the user/group under which the web server runs against the Swarm distribution.

5. Configure Redis, see ["Redis configuration" below](#).

Redis configuration

Swarm requires Redis to manage its caches and by default Swarm uses its own Redis server on the Swarm machine. Swarm caches data from the Helix Core Server to improve the performance of common searches in Swarm and to reduce the load on the Helix Core Server. Redis is included in the Swarm Tarball installation.

This section describes how to configure the Swarm Redis service on the Swarm machine. Do not change the default values of the following configuration files if you are using the Swarm Redis server.

Tip

If required, you can use your own Redis server instead of the Swarm Redis server. For instructions on how to configure Swarm to use your own Redis server, see ["Use your own Redis server" on page 235](#).

Swarm has two Redis binaries in **SWARM_ROOT/p4-bin/bin.linux26x86_64/**:

- **redis-server-swarm**
- **redis-cli-swarm**

1. Configure the Redis cache database, enter the following details in the **redis-server.conf** file in **/opt/perforce/etc/**:

```
bind 127.0.0.1
port 7379
supervised auto
save ""
dir /opt/perforce/swarm/redis
```

Default values:

Tip

- The default settings are shown below, the **redis-server.conf** file contains more detailed information about the Redis configuration for Swarm.
- On Swarm systems with a large number of users, groups, and projects, start-up time can be improved by persisting the memory cache. You can persist the memory cache by disabling background saves and enabling append saves, see the **redis-server.conf** file comments for detailed information.

- **bind 127.0.0.1** - Redis server IP address (loopback interface)
- **port 7379** - Redis server port number
- **supervised auto** - detects the use of **upstart** or **systemd** automatically to signal that the process is ready to use the supervisors
- **save ""** - background saves disabled, recommended.
- **dir /opt/perforce/swarm/redis** - the directory the Redis cache database is stored in.

Tip

To fine-tune your Redis server settings, make your changes in the Laminas Redis adapter. For information about the Laminas Redis adapter, see the [Laminas Redis Adapter documentation](#).

2. The Redis cache database must be running before Swarm starts so we recommend setting it up as a service so that it starts automatically at boot time.

The following example script will:

- run the Redis service as the Perforce user
- use the configuration file in **/opt/perforce/etc/redis-server.conf**
- assume the database server is installed in **/opt/perforce/swarm/p4-bin/bin.linux26x86_64/**

To configure Redis as a service, add a script with the following content in **/etc/systemd/system**

redis-server-swarm.service

```
[Unit]
Description=Redis Server for Swarm
After=network.target

[Service]
ExecStart=/opt/perforce/swarm/p4-bin/bin.linux26x86_64/redis-server-swarm /opt/perforce/etc/redis-server.conf
ExecStop=/opt/perforce/swarm/p4-bin/bin.linux26x86_64/redis-cli-swarm shutdown
Restart=always
RestartSec=10
StandardOutput=syslog
StandardError=syslog
SyslogIdentifier=redis
User=perforce
    Group=perforce

[Install]
WantedBy=multi-user.target
```

3. Create the **`SWARM_ROOT/data/config.php`** file if it does not already exist. The **redis** block of the **`SWARM_ROOT/data/config.php`** file contains the **password**, **namespace**, **host** and **port** details of the Swarm Redis server:

```
<?php
// this block should be a peer of 'p4'
'redis' => array(
    'options' => array(
        'password' => null, // Defaults to null
        'namespace' => 'Swarm',
        'server' => array(
            'host' => 'localhost', // Defaults to 'localhost'
            or enter your Redis server hostname
            'port' => '7379', // Defaults to '7379 or enter
```

```

your Redis server port
    ),
    ),
    'items_batch_size' => 100000,
    'check_integrity' => '03:00', // Defaults to '03:00' Use
one of the following two formats:
                                // 1) The time of day that
the integrity check starts each day. Set in 24 hour format with
leading zeros and a : separator
                                // 2) The number of seconds
between each integrity check. Set as a positive integer. Specify
'0' to disable the integrity check.
    'population_lock_timeout' => 300, // Timeout for initial
cache population. Defaults to 300 seconds.
    ),

```

Configurables:

- **password**: Redis server password. Defaults to **null** and should be left at default if using the Swarm Redis server.
- **namespace**: the prefix used for key values in the Redis cache. Defaults to **Swarm** and should be left at default if using the Swarm Redis server.

Note

If you have multiple Swarm instances running against a single Redis server, each Swarm server must use a different Redis **namespace**. This enables the cache data for the individual Swarm instances to be identified. The **namespace** is limited to ≤ 128 characters.

If one or more of your Swarm instances is connected to multiple Helix Core Servers, the Redis **namespace** includes the **server label** and the character limit is reduced to ≤ 127 characters, see ["Multiple Helix Core Server instances" on page 646](#).

- **host**: Redis server hostname. Defaults to **localhost** and should be left at default if using the Swarm Redis server.
- **port**: Redis server port number. Defaults to **7379**. Swarm uses port **7379** as its default to avoid clashing with other Redis servers that might be on your network. It should be left at default if using the Swarm Redis server. The default port for a non-Swarm Redis server is **6379**.

- **items_batch_size**: Maximum number of key/value pairs allowed in an **mset** call to Redis. Sets exceeding this will be batched according to this maximum for efficiency. Defaults to **100000**.

Note

The default value of **100000** was chosen to strike a balance between efficiency and project data complexity. This value should not normally need to be changed, contact support before making a change to this value.

- **check_integrity**: In some circumstances, such as when changes are made in the Helix Core Server when Swarm is down or if errors occur during updates, the Redis cache can get out of sync with the Helix Core Server. Swarm can run a regular integrity check to make sure that the Redis caches and Helix Core Server are in sync. If an integrity check finds an out of sync cache file, Swarm automatically updates the data in that cache.

The **check_integrity** configurable specifies when the Redis cache integrity check is run. Set as a specific time of day (24 hour format with leading zeros) or a number of seconds (positive integer) between checks. Disable the integrity check with **'0'**. Defaults to **'03:00'**.

- **population_lock_timeout**: specifies the timeout, in seconds, for initial cache population. If you have a large Swarm system, increase this time if the initial cache population times out. Defaults to **300** seconds.

Configure Apache, see ["Apache configuration" below](#).

Apache configuration

The configuration of the Apache HTTP Server (Apache) can vary between OS distributions; see the documentation specific to your installation of Apache.

Important

The following notes are applicable for CentOS/RHEL7, RHEL 8, and RHEL 9:

- **Swarm 2019.1 to 2020.1**: these versions of Swarm used PHP packages from the SCL repositories for CentOS/RHEL 7. This was to provide access to more recent versions of PHP than are available as standard on CentOS/RHEL. This required the use of the **httpd24-httpd** package for Apache. These were all installed into **/opt/rh**

Swarm 2020.2 and later: these versions of Swarm uses the Remi repository for CentOS/RHEL 7, RHEL 8, and RHEL 9. This provides PHP 7.4 installed in the standard file system structure. This means that the old **httpd24-httpd** version of Apache is no longer needed, and the standard system version of Apache is being used again.

The SCL Apache site configuration file was installed at this location for Swarm 2019.1 to 2020.1:

```
/opt/rh/httpd24/root/etc/httpd/conf.d/perforce-swarm-site.conf
```

If this exists when Swarm is upgraded to 2020.2 and later, this file is copied to `/etc/httpd/conf.d/perforce-swarm-site.conf` if there is no file at the destination. It is also re-written to change references from `/var/log/httpd24` to `/var/log/httpd`

If a site configuration file for Swarm already exists in `/etc/httpd`, the copy and re-write is not performed.

After upgrade, `httpd24-httpd` is disabled.

- To avoid seeing the Apache HTTP server Linux test page when you start the Apache server, comment out the content of the `welcome.conf` file located in the `/etc/httpd/conf.d/` directory.
- To avoid loading the Apache HTTP server example configuration instead of the Swarm configuration when the Apache server starts, rename the `autoindex.conf` file located in the `/etc/httpd/conf.d/` directory to `z-autoindex.conf` or similar. This is required because Apache runs the first conf file it finds in the `/etc/httpd/conf.d/` directory (alphabetical order) and that must be the `perforce-swarm-site.conf` file.

1. Locate your system's Apache configuration.

Common configuration directories include:

- `/etc/httpd/conf/`
- `/etc/apache2/`
- `/usr/local/apache2/conf/`
- `/Applications/XAMPP/etc/`

Within the configuration path, the main Apache configuration file is usually named one of the following:

- `httpd.conf`
- `apache2.conf`

Tip

A longer discussion on the possible locations and names of Apache configuration files is available here : <https://cwiki.apache.org/confluence/display/httpd/DistrosDefaultLayout>

2. Set up an Apache virtual host (vhost) for your installation.

If your Apache configuration directory contains the directories `sites-available` and `sites-enabled`:

- a. Copy the appropriate virtual host definition below into the file **sites-available/swarm**.
- b. Enable the Swarm virtual host definition.

```
sudo a2ensite swarm
```

Otherwise, copy the virtual host definition below into the bottom of the main Apache configuration file, **httpd.conf** or **apache2.conf**.

Virtual host definition example for Apache 2.4:

```
<VirtualHost *:80>
    ServerName myswarm.host
    AllowEncodedSlashes NoDecode
    ServerAlias myswarm
    ErrorLog "/path/to/apache/logs/myswarm.error_log"
    CustomLog "/path/to/apache/logs/myswarm.access_log" common
    DocumentRoot "/path/to/swarm/public"
    <Directory "/path/to/swarm/public">
        AllowOverride All
        Require all granted
    </Directory>
</VirtualHost>
```

Note

See Apache's virtual host documentation for details:
<https://httpd.apache.org/docs/2.4/vhosts/>

3. Customize the virtual host definition.
 - a. Replace **myswarm.host** with the hostname for Swarm on your network. This may require adjusting the DNS configuration on your network.
 - b. Replace **myswarm** with the name of the subdomain hosting Swarm. Many administrators choose swarm.

Note

The string **myswarm** in the log file paths: this should match the subdomain name and prefix for the log files, to help coordinate the active host with the log files for that host. Doing this is particularly useful when your Apache server hosts multiple instances of Swarm.

- c. Replace `/path/to/apache/logs` with the path where your Apache stores its log files. Apache's log files are typically named `access_log` and `error_log`.
 - d. Replace `/path/to/swarm` with the path to the Swarm directory.
4. Verify that the correct Apache modules are enabled.
 - To query whether the PHP and Rewrite modules are active, use the `apachectl` utility to list all of the active modules (this may be named `apache2ctl` on your system):


```
apachectl -t -D DUMP_MODULES
```
 - Simply look for `php7_module` and `rewrite_module` in the output. If you see them, skip ahead to step 5.
 - If the Apache utility `a2enmod` is installed, use it to enable the PHP and Rewrite modules:


```
sudo a2enmod php7 rewrite
```
 - Without the `a2enmod` utility, edit the Apache configuration file by hand. Locate your Apache configuration file for modules and either uncomment or add the following lines:


```
LoadModule  php7_module      libexec/apache2/libphp7.so
LoadModule  rewrite_module   libexec/apache2/mod_rewrite.so
```
 - Note that your Apache installation may have different paths for the location of its modules (the `.so` files).
 5. Restart your web server.
 - To ensure that the Apache configuration changes you made become active, restart the web server.


```
sudo apachectl restart
```
 - Query Apache's active virtual hosts and modules to confirm your changes are in effect:


```
apachectl -t -D DUMP_VHOSTS
apachectl -t -D DUMP_MODULES
```

Important

Apache must be configured to use the prefork MPM because P4PHP does not support threaded operation.

The prefork **MPM** is the default for Linux Apache installations, so you may not have to do anything.

For more information on Apache MPMs and configuration, see:
<https://httpd.apache.org/docs/2.4/mpm.html>

6. Configure PHP, see "PHP configuration" on the facing page.

PHP configuration

PHP can vary between OS distributions; see the documentation specific to your installation of PHP.

1. First determine which PHP `.ini` file is in use by the PHP Apache module.

Note

It may not necessarily be the same `.ini` file that is in use when calling PHP from the command line (running `php --ini` from the command line reports this).

If you are having trouble determining which `.ini` file the PHP Apache module is using, create a PHP file that can be served through Apache with the following contents:

```
<?php phpinfo();?>
```

Point your browser to this file and look for this table row in the resulting table:

Loaded Configuration File

2. Ensure that `date.timezone` is set correctly for your system.

Some distributions do not make a default timezone available to PHP, so the best practice to set the timezone for PHP explicitly. See the [list of supported timezones](#).

An example `date.timezone` setting in `php.ini`:

```
date.timezone = America/Vancouver
```

3. Verify that the `iconv`, `json`, and `session` extensions are present.

They are usually enabled by default, although you may have to install packages for them through your OS distribution. Verify they are present by searching for their respective names in the `phpinfo` output above.

4. Enable P4PHP, the Perforce extension for PHP:

For Swarm to communicate with the Helix Core Server, it needs the P4PHP extension. P4PHP variants are available for each [PHP 7 version supported](#) by Swarm and each variant is available as either `openssl 1.0.2` or `openssl 1.1.1`. All of the P4PHP variants are supplied with the Swarm package and tarball installations.

Note

- For Linux, the default variants are compiled with `glibc 2.11`.
- P4PHP 2019.1 or later is required for Swarm 2019.1 and later.

Swarm package and tarball installations: two versions of P4PHP are supplied for each PHP 7 and PHP 8 version supported by Swarm. They are located in the `p4-bin/bin.linux26x86_64` directory.

- `perforce-php7x.so` and `perforce-php8x.so` compatible with systems using `SSL 1.0.2`

- **perforce-php7x-ssl1.1.1.so** and **perforce-php8x-ssl1.1.1.so** compatible with systems using SSL 1.1.1 (by default Ubuntu 20.04 uses SSL 1.1.1)
- **perforce-php8x-ssl3.so** compatible with systems using SSL 3.0.0 (by default Ubuntu 22.04 uses SSL 3.0.0)

Where **x** is the version of PHP 7 or PHP 8.

If the **perforce.ini** file is not pointing at the correct version of P4PHP and you connect to an SSL enabled Helix Core Server:

- The Swarm web-page will not load and you might see a **Connection Reset** error.
- There might be an **undefined symbol: SSLey** message in the Apache error log

To enable P4PHP, edit the web server's **php.ini** file and add the following line:

```
extension=/path/to/swarm/p4-bin/bin.<platform>/perforce-
<variant>.so
```

Example 1: for a 64-bit Linux system running PHP 7.2 and SSL 1.0.2:

```
extension=/path/to/swarm/p4-bin/bin.linux26x86_64/perforce-
php72.so
```

Example 2: for a 64-bit Linux system running PHP 7.4 and SSL 1.0.2:

```
extension=/path/to/swarm/p4-bin/bin.linux26x86_64/perforce-
php74.so
```

Example 3: for a 64-bit Linux system running PHP 8.0 and SSL 1.0.2:

```
extension=/path/to/swarm/p4-bin/bin.linux26x86_64/perforce-
php80.so
```

Example 4: for a 64-bit Linux system running PHP 8.1 and SSL 1.0.2:

```
extension=/path/to/swarm/p4-bin/bin.linux26x86_64/perforce-
php81.so
```

Example 5: for a 64-bit Linux system running PHP 8.2 and SSL 1.0.2:

```
extension=/path/to/swarm/p4-bin/bin.linux26x86_64/perforce-
php82.so
```

Example 6: for a 64-bit Linux system running PHP 7.2 and SSL 1.1.1:

```
extension=/path/to/swarm/p4-bin/bin.linux26x86_64/perforce-
php72-ssl1.1.1.so
```

Example 7: for a 64-bit Linux system running PHP 7.4 and SSL 1.1.1:

```
extension=/path/to/swarm/p4-bin/bin.linux26x86_64/perforce-  
php74-ssl1.1.1.so
```

Example 8: for a 64-bit Linux system running PHP 8.0 and SSL 1.1.1:

```
extension=/path/to/swarm/p4-bin/bin.linux26x86_64/perforce-  
php80-ssl1.1.1.so
```

Example 9: for a 64-bit Linux system running PHP 8.1 and SSL 1.1.1:

```
extension=/path/to/swarm/p4-bin/bin.linux26x86_64/perforce-  
php81-ssl1.1.1.so
```

Example 10: for a 64-bit Linux system running PHP 8.2 and SSL 1.1.1:

```
extension=/path/to/swarm/p4-bin/bin.linux26x86_64/perforce-  
php82-ssl1.1.1.so
```

Example 11: for a 64-bit Linux system running PHP 8.2 and SSL 3.0.0:

```
extension=/path/to/swarm/p4-bin/bin.linux26x86_64/perforce-  
php82-ssl3.so
```

Alternatively, copy the extension file to the default location for PHP extensions, and then just add this line instead:

```
extension=perforce-<variant>.so
```

5. Restart Apache for the changes to become active.
6. To verify that P4PHP is active, navigate to the `phpinfo` file you [created above](#). You should then see a perforce section (search for "Perforce Module"). It should report that the module is enabled and display the version information.

Note

Be aware that any operating system upgrades on the machine hosting Swarm may involve updates to PHP. If this occurs, the PHP `.ini` file needs to be updated to point to the correct **variant** of P4PHP to match the version of PHP that the upgraded operating system is using.

7. In addition, Swarm greatly benefits from the following optional extension: ImageMagick extension for PHP to improve Swarm's ability to preview graphic formats that web browsers typically cannot display, see "[ImageMagick \(imagick\) extension for PHP](#)" below.
8. Configure Swarm, see "[Swarm configuration](#)" on the next page.

ImageMagick (imagick) extension for PHP

Imagick is a PHP extension that integrates the ImageMagick graphics library's API for the creation and manipulation of images. Enabling Imagick improves Swarm's ability to preview graphics formats that web browsers typically cannot display and to create comment attachment thumbnails.

For more information, see:

- <https://secure.php.net/imagick>
 - <https://pecl.php.net/package/imagick>
1. We recommend that you install Imagick from your OS distribution, via **apt-get**, **yum**, etc. If your distribution does not offer the imagick package for PHP, install it via PECL (although you may have to resolve system dependencies):


```
sudo pecl install imagick
```
 2. Verify that imagick is enabled in your PHP Apache module's PHP **.ini** file (as determined in the section above for P4PHP). You may need to add the following line:


```
extension=imagick.so
```
 3. Configure the ImageMagick **policy.xml** file to allow Swarm to create thumbnails if it is not already configured:
 - a. Open the ImageMagick **policy.xml** file for editing:


```
/etc/ImageMagick-6/policy.xml
```
 - b. Disable lines with the pattern below for the following **<filetype>**s: **PS**, **PS2**, **EPS**, **PDF**, or **XPS**:


```
<policy domain="coder" rights="none" pattern="<filetype>" />
```
 - c. Save the **policy.xml** file.
 4. Restart Apache for the changes to become active.
 5. To verify that imagick is active, navigate to the **phpinfo** file you [created earlier](#). You should then see an imagick section. It should report its version information and a table for its directives, supported image file formats, and more.

Warning

Once you have completed installing and enabling P4PHP and imagick, we recommend that you remove the **phpinfo** file you created to avoid disclosing information about your installation.

Swarm configuration

Now that Swarm is ready for use, you need to configure it to work in your environment.

Important

Swarm does not support Helix Core Servers configured to use **P4AUTH**, see [Centralized authentication server \(P4AUTH\)](#) in the *Helix Core Server Administrator Guide*.

Note

- Swarm can be connected to Helix Core Servers (P4D) and commit servers.
To configure Swarm to connect to more than one Helix Core Server (P4D), see ["Multiple Helix Core Server instances" on page 646](#).
To configure Swarm to connect to a Helix Core Server configured to use *commit-edge architecture*, see ["Commit-edge deployment" on page 587](#).
- Swarm must not be connected to Helix Broker, Helix Proxy, Helix Edge, forwarding replica, or read-only replica servers.

Swarm configuration file

Important

- If your Helix Core Server is configured for Helix Authentication Service, the Helix Core Server must be temporarily configured to allow fall-back to passwords while you establish a connection to the Helix Core Server. Run the following command on the Helix Core Server to enable fall-back to passwords:

```
p4 configure set auth.sso.allow.passwd=1
```

Edit the [SWARM_ROOT/data/config.php](#) file so that it contains the following configuration block:

```
<?php
return array(
    'p4' => array(
        'port'      => 'my-helix-core-server:1666',
        'user'      => 'admin_userid',
        'password'  => 'admin user ticket or password',
        'sso'       => 'disabled', //
        'disabled'|'optional'|'enabled'
                                // default value is 'disabled'
    ),
    'log' => array(
        'priority' => 3, // 7 for max, defaults to 3
    ),
    'mail' => array(
        'transport' => array(
            'host' => 'my.mx.host',
        ),
    ),
);
```

```
    ),
  );
```

port

For the port value, replace *my-helix-core-server:1666* with the P4PORT value used to connect to your Helix Core Server.

Important

If your Helix Core Server is deployed using the commit-edge architecture, ensure that Swarm's port value points to the commit server.

For more information, see [Commit-edge](#) in the *Helix Core Server Administrator Guide*.

Warning

If the port points to a Helix Broker, ensure that the broker does not delegate commands to different replicas, edge servers, or proxies. Such delegation can cause odd problems or outright failures in Swarm.

Swarm needs to have a consistent, current view of the state of Helix Core Server, and works best when it connects to a central/commit server.

user

For the user value, replace **admin_userid** with a normal Helix Core Server userid that has **admin**-level access to Helix Core Server.

Note

- This user is used by Swarm to communicate with the Helix Core Server. It should not be used to perform everyday Swarm reviewing tasks.
- Helix Core Server 2020.1 and later, permissions have changed for viewing and [editing stream spec files](#) in Swarm. To view and edit stream spec files in Swarm, the Swarm user must have *admin* permissions for the entire depot `/ / . . .`

password

For the password value, while a plain-text password works, we recommend that you use a ticket value instead. Obtain the ticket value for the **admin_userid** during login with this command:

```
p4 -p my-helix-core-server:1666 -u admin_userid login -p
```

Important

- If your Helix Core Server authentication is configured in the one of the following ways, ticket-based authentication is required:
 - Authentication configured with security level 3.
 - Authentication configured for LDAP.
 - Authentication configured for Helix Authentication Service.
- You must use a long-lived login ticket for the Swarm user.

You can determine when the admin `userid` ticket will expire with:

```
p4 -p my-helix-core-server:1666 -u admin_userid -P ticket_value
login -s
```

For more information about tickets, see the section [Ticket-based authentication](#) in the *Helix Core Server Administrator Guide*.

SSO

Swarm can be configured for the Helix Authentication Service when the Helix Core Server is also configured for the Helix Authentication Service.

Set `sso` to one of the following:

- **enabled** all users must use Helix Authentication Service to log in to Swarm.

The login page displays an input box to enter the email or username. Once the input box is filled, the **Log in with SSO** button is activated. See ["Log in with SSO" on page 406](#).

If [SAML](#) is configured when the `sso` configurable is enabled, the user can only log in to Swarm using the [Log in with SSO](#) button.

- **optional** Helix Authentication Service is available for users to log in to Swarm but is not enforced.

The login page displays an input box to enter the email or username. Once the input box is filled, the **Log in with SSO** button is enabled. To log in with SSO, see ["Log in with SSO" on page 406](#).

The user can also manually log in to Swarm using the **Log in with credentials** button. To log in manually, see ["Log in with a password" on page 405](#).

If you switch to **Log in with credentials** and then decide to use single sign-on, click **Log in with SSO**.

If [SAML](#) is configured when the `sso` configurable is set to optional, the user can either log in to Swarm using the [Log in with SSO](#) button or the [Log in with credentials](#) button.

- **disabled** Helix Authentication Service is not available to Swarm. This is the default value.

The login page displays two input boxes to enter the email or username, and to enter the password. Once the input boxes are filled, the **Log in with credentials** button is activated. See ["Log in with a password" on page 405](#).

Important

From Swarm 2021.1, the `sso_enabled` configurable is deprecated but remains supported. It is replaced with the more flexible `sso` configurable. If the `sso_enabled` configurable and `sso` configurable are both present in the `p4` configuration block, Swarm uses the `sso` configurable value.

host

For the host value, replace `my.mx.host` with the hostname of the mail exchange service that Swarm should use to send its email notifications.

Note

Since this configuration file contains the credentials for a Helix Core Server *admin*-level user, we recommend that this file's ownership and permissions be adjusted such that only the web server user can read the file, and that no user can write the file.

Next step

Configure the Helix Core Server to notify Swarm about Helix Core Server events, see ["Configuring Helix Core Server event notification" below](#).

Configuring Helix Core Server event notification

Tip

Triggers are still supported, but we recommend you use Helix Core Server Extensions. Helix Core Server Extensions are easier to install and maintain than Triggers.

Swarm needs to know about a number of Helix Core Server events to operate correctly, this can be done by using Helix Core Server Extensions (recommended) or Helix Core Server Triggers. Swarm installs include the Swarm Helix Core Server extension file and trigger scripts required for Swarm to get the events it needs from your Helix Core Server.

You must install and configure Helix Core Server Extensions or Helix Core Server Triggers to complete the Swarm installation.

Do one of the following so that Swarm is notified about events on the Helix Core Server:

- Install and configure the Swarm Helix Core Server extension, see ["Installing the Swarm Helix Core Server extension \(recommended\)" on the facing page](#)
- Install and configure Swarm Helix Core Server Triggers, see ["Installing Triggers" on page 204](#)

Installing the Swarm Helix Core Server extension (recommended)

Important

- If you are using the Swarm Helix Core Server extension, Swarm Helix Core Server Triggers must not be installed.
- You must be a user with *super* user permissions to install and configure Helix Core Server Extensions.

Swarm needs to know about a number of Helix Core Server events to operate correctly, this is done by using Helix Core Server Extensions. Helix Core Server Extensions must be installed and configured on your Helix Core Server to complete your Swarm installation. The Swarm Helix Core Server extension is included with the Swarm product download.

Prerequisites

To install the Swarm Helix Core Server extension you need:

A compatible version of Helix Core Server for Helix Core Server Extensions:

- **Linux:** Helix Core Server 2021.2 and later. If you are using an earlier version of Helix Core Server, you must use triggers.
- **Windows:** Helix Core Server 2021.2 and later. If you are using an earlier version of Helix Core Server, you must use triggers.

You will also need:

- Helix Core Server Extensions installed and configured on your Helix Core Server.
- A user with *super* permissions to install and configure the Swarm Helix Core Server extension.

Swarm Helix Core Server extension installation

Tip

The following commands can be run on the Swarm server or the Helix Core Server. You must be a Helix Core Server *super* user to install and configure Helix Core Server Extensions.

The Swarm Helix Core Server extension file is included for all of the Swarm installation types. The **helix-swarm.p4-extension** file is located in "[swarm_root](#)" on [page 724/p4-bin/extensions](#)

Run the following commands as a Helix Core Server user with *super* permissions:

1. Install the Swarm Helix Core Server extension on the Helix Core Server with:

```
p4 extension --yes --install helix-swarm.p4-extension
```

Example response:

```
Extension 'Perforce::helix-swarm' version '2022.1.20221215'
installed successfully.
Perform the following steps to turn on the Extension:

# Create a global configuration if one doesn't already exist.
p4 extension --configure Perforce::helix-swarm

# Create an instance configuration to enable the Extension.
p4 extension --configure Perforce::helix-swarm --name
Perforce::helix-swarm-instanceName

For more information, visit:
https://www.perforce.com/manuals/extensions/Content/Extensions/Home-extensions.html
```

2. Create a global configuration if one doesn't already exist:

```
p4 extension --configure Perforce::helix-swarm
```

The spec file opens in your text editor, for example:

```
ExtName:          helix-swarm
ExtDescription:
    Helix Swarm Extension

ExtVersion:       2022.1.DEV.20220120

ExtUUID:          4532BC59-7BC8-478F-ADF6-0A563C42563D

ExtRev: 1

ExtMaxScriptTime:      unset

ExtMaxScriptMem:       unset

ExtAllowedGroups:

ExtEnabled:           true
```

```
ExtP4USER:      sampleExtensionsUser

Name:   helix-swarm

Owner:   super

Update: 2022/01/21 10:43:46

Description:
    The description of your config.

ExtConfig:
    Debug:
        2
    Swarm-Secure:
        true
    Swarm-Token:
        ... SWARM-TOKEN
    Swarm-URL:
        http://localhost/
```

3. Edit the spec file to match your system:

- **ExtP4USER:** Set to an existing user with *super* permissions.
- **ExtConfig:** Check and edit the values in this block:

- **Debug:**

Debug levels **0** to **3** control the amount of debug information sent to the Helix Core Server Extensions log.

Important

A debug level of **10** and higher sends all debug information to the client. This is useful for debugging, but should not be run in a production environment.

- **SSL-CA-File:** (optional) Set the absolute path of the [Certificate Authority \(CA\) Privacy-Enhanced Mail \(PEM\)](#) file to validate the SwarmHelix Core Server's certificate. Ensure that this file is accessible to the Helix Core Server process owner.

If a suitable certificate is not found in the local SSL certificate store maintained by the operating system or if there are verification issues from using a self-signed certificate then the default CA file is used.

By default, the local certificate stored in the default CA file is used if the absolute path for the [Certificate Authority \(CA\)](#) is not specified.

- **Swarm-Secure:**
 - `true` will only accept secure SSL certificates.
 - `false` will accept insecure SSL certificates.
- **Swarm-Token:** Set to the Swarm trigger token value.

Important

To set multiple Helix Core Server token settings for Extensions, repeat the following steps for each Helix Core Server. The process to obtain an extension token of your Swarm instance is similar to obtaining a trigger token of your Swarm instance.

To obtain the trigger token of your Swarm instance:

- Log in to Swarm as a *super* user.
- Click your *userid*, found at the right of the main toolbar.
- Select **About Swarm**.
- The **About Swarm** dialog is displayed and Swarm generates an API token if it doesn't already exist.
- Copy the trigger token value displayed at the bottom of the dialog and paste the token into the spec file.

- **Swarm-URL:** Set to the URL of your Swarm instance.

- When you have finished editing the spec file, save it in your text editor.
- Create an instance configuration to enable the extension:

```
p4 extension --configure Perforce::helix-swarm --name swarm
```

The spec file opens in your text editor, for example:

```
ExtName:          helix-swarm

ExtDescription:
    Helix Swarm Extension

ExtVersion:       2022.1.DEV.20220120

ExtUUID:          4532BC59-7BC8-478F-ADF6-0A563C42563D

ExtRev: 1

ExtMaxScriptTime:      unset

ExtMaxScriptMem:       unset

ExtEnabled:            true

ExtDebug:              none

Name:    swarm

Owner:    super

Update: 2022/01/21 10:58:41

Description:
    The description of your config.

ExtConfig:
    depot-path:
        //...
    enableStrict:
        true
```

```

enableWorkflow:
    true
httpTimeout:
    30
ignoreErrors:
    false

```

We recommend that these settings are left at their default values, but the configuration needs to be run to set the default values.

- **enableStrict:** Enable strict checking of reviews for workflow. Verifies that the file content in a commit matches the file content of its associated approved review. If one or more files in a commit do not match the content of the file in its associated review, the commit is rejected.
- **enableWorkflow:** Enable workflow.
- **httpTimeout:** Timeout (in seconds) when communicating with the Swarm server.
- **ignoreErrors:** If the server returns an error (timeout, not there), then default to allowing a request. It means workflow rules can be skipped, but if the Swarm server is down it will not block submits.

6. Save the spec file in your text editor.
7. Confirm your Swarm Helix Core Server extension is installed and configured by running the following command on the Helix Core Server:

```
p4 extension --run swarm ping
```

- If the Swarm Helix Core Server extension is working, the Helix Core Server responds with **OK**
- If the Swarm Helix Core Server extension is not working, the Helix Core Server responds with an error message. For example, **BAD (Cannot reach https://swarm-example.com/)**

For more information about Helix Core Server Extensions, see:

- [Extension overview](#) in the *Helix Core Extensions Developer Guide*.
- [p4 extension](#) in the *Helix Core Command-Line (P4) Reference*.

8. Depending on your installation type, do one of the following:
 - **Swarm package installation:** Review the post-install configuration options to customize your Swarm installation, see "[Post-install configuration options](#)" on page 224.
 - **Swarm tarball installation:** Set up a recurring task to spawn workers, see "[Set up a recurring task to spawn workers](#)" on page 220.

Add Swarm as an HTML tab in the Helix Core Visual Client (P4V)

To add Swarm as an HTML tab in the Helix Core Visual Client (P4V), see [HTML Tabs](#) section in the *P4VJS Developer Guide*.

Installing Triggers

Important

- If you are using Swarm Helix Core Server Triggers, the Swarm Helix Core Server extension must not be installed.
- Swarm does not support Helix Core Servers configured to use **P4AUTH**, see [Centralized authentication server \(P4AUTH\)](#) in the *Helix Core Server Administrator Guide*.

Note

- Swarm can be connected to Helix Core Servers (P4D) and commit servers.
To configure Swarm to connect to more than one Helix Core Server (P4D), see ["Multiple Helix Core Server instances"](#) on page 646.
To configure Swarm to connect to a Helix Core Server configured to use *commit-edge architecture*, see ["Commit-edge deployment"](#) on page 587.
- Swarm must not be connected to Helix Broker, Helix Proxy, Helix Edge, forwarding replica, or read-only replica servers.

Tip

Triggers are still supported, but we recommend you use Helix Core Server Extensions. Helix Core Server Extensions are easier to install and maintain than Triggers. For instructions on installing Extensions, see ["Installing the Swarm Helix Core Server extension \(recommended\)"](#) on page 198.

Now that you have a configured instance of Swarm, the last piece is to configure your Helix Core Server to tell Swarm about interesting events. This can be accomplished through the use of triggers.

For more information about Helix Core Server triggers, see [Triggers](#) in *Helix Core Server Administrator Guide*.

Using triggers to push events to Swarm

Helix Core Server provides a facility called *triggers* to customize the operation of the server, or to invoke additional processing for specific kinds of versioning operations. Swarm provides a trigger script written in Perl that notifies Swarm about activity within the Helix Core Server.

See ["Trigger options"](#) on page 735 for more information on configuring the Perl trigger.

Set up Swarm triggers with a Windows or Linux-hosted Helix Core Server

Tip

Many of the steps in this procedure are common to both Windows and Linux-hosted Helix Core Servers. When there is a difference between the Windows and Linux operating systems, this is indicated in the procedure.

1. Ensure that the required ["Trigger dependencies" on page 108](#) have been installed on the machine hosting the Helix Core Server.
2. Copy the Swarm trigger to the Helix Core Server
 - If your Helix Core Server is version 2014.1 (or later), we recommend submitting the trigger script, `p4-bin/scripts/swarm-trigger.pl`, to Helix Core Server and running it from the depot. The recommended depot location is `/.swarm/triggers/swarm-trigger.pl`, especially if you have already setup ["Comment attachments" on page 580](#).
 - If your Helix Core Server is older than version 2014.1, or you prefer that the trigger file exist in the filesystem, you must copy the `p4-bin/scripts/swarm-trigger.pl` script to the server hosting Helix Core Server. If your Helix Core Server deployment uses the commit-edge architecture, the script must also be copied to all edge servers, and it must exist in the same path on all servers.
3. Trigger requests to Swarm without a valid API token are ignored. Obtain your API token:
 - a. [Log in](#) to Swarm as a *super* user.
 - b. Click your *userid*, found at the right of the main toolbar.
 - c. Select **About Swarm**.

The **About Swarm** dialog is displayed and Swarm generates an API token if it doesn't already exist.
 - d. Make a note of the trigger token value displayed at the bottom of the dialog, you will need the trigger token when you configure Helix Core Server for Swarm.

Tip

Click on the trigger token to select it and then copy it to your clipboard. You can paste it into the `swarm-trigger.conf` file when you configure Helix Core Server for Swarm.

4. Configure the trigger.

You need to use the API token established in the previous step.

`swarm-trigger.pl` can be configured directly, but the preferred approach is to create a configuration file called `swarm-trigger.conf`. Using the configuration file greatly simplifies upgrades.

- If you are using the Swarm triggers package described in ["Install and configure Swarm from a package \(recommended\)" on page 115](#), the configuration file is available at `/opt/perforce/etc/swarm-trigger.conf`
 - If you submitted the trigger script to the depot in the previous step, you should similarly submit the configuration file to the depot. The recommended depot location is `/.swarm/triggers/swarm-trigger.conf`.
 - If you copied the trigger script to the commit server and all edge servers in the previous step, also copy the configuration file to the commit server and all edge servers, making sure that they exist in the same path on all servers.
- If you are not using the Swarm trigger package, create `swarm-trigger.conf` in the same directory as `swarm-trigger.pl`.
 - If you submitted the trigger script to the depot in the previous step, you should similarly submit the configuration file to the depot. The recommended depot location is `/.swarm/triggers/swarm-trigger.conf`.
 - If you copied the trigger script to the commit server and all edge servers in the previous step, also copy the configuration file to the commit server and all edge servers, making sure that they exist in the same path on all servers.

The following is a sample of what your `swarm-trigger.conf` should contain:

```
# SWARM_HOST (required)
# Hostname of your Swarm instance, with leading "http://" or
# "https://".
SWARM_HOST="http://my-swarm-host"

# SWARM_TOKEN (required)
# The token used when talking to Swarm to offer some security. To
# obtain the
# value, log in to Swarm as a super user and select 'About Swarm'
# to see the
# token value.
SWARM_TOKEN="MY-UUID-STYLE-TOKEN"

# ADMIN_USER (optional) Do not use if the Workflow feature is
# enabled (default)
# For enforcing reviewed changes, optionally specify the normal
# Perforce user
# with admin privileges (to read keys); if not set, will use
```

```

whatever Perforce
# user is set in environment.
ADMIN_USER=

# ADMIN_TICKET_FILE (optional) Do not use if the Workflow feature
is enabled (default)
# For enforcing reviewed changes, optionally specify the location
of the
# p4tickets file if different from the default
($HOME/.p4tickets).
# Ensure this user is a member of a group with an 'unlimited' or
very long
# timeout; then, manually login as this user from the Perforce
server machine to
# set the ticket.
ADMIN_TICKET_FILE=

# VERIFY_SSL (optional)
# If HTTPS is being used on the Swarm web server, then this
controls whether
# the SSL certificate is validated or not. By default this is set
to 1, which
# means any SSL certificates must be valid. If the web server is
using a self
# signed certificate, then this must be set to 0.
VERIFY_SSL=1

```

Modify the **swarm-trigger.conf** configuration file to set the **SWARM_HOST** and the **SWARM_TOKEN** variables appropriately.

Note
Windows only:

You may need to edit the trigger script to specify the full path to **curl.exe** if the **HTTP::Tiny** module is not installed.

Note**Linux only:**

swarm-trigger.pl looks for configuration in the following files. Variables defined in the later files will override the earlier defined variables of the same name:

- Variables set inside the **swarm-trigger.pl** script itself
- **/etc/perforce/swarm-trigger.conf**
- **/opt/perforce/etc/swarm-trigger.conf**
- The **swarm-trigger.conf** file stored in the same directory as **swarm-trigger.pl**.
- Any file passed to the **swarm-trigger.pl** script using the **-c** command line argument

Important**Windows and Linux:**

If you specify **ADMIN_USER**, the ticket contained in **%USERPROFILE%/p4tickets.txt** (Windows), **\$HOME/.p4tickets** (Linux), or the optional ticket file specified with the **ADMIN_TICKET_FILE** must use the port that was used to start Helix Core Server. For example, if **p4d** is started with:

```
C:\> p4d -p my-helix-core-server:1666 ...
```

then the ticket for the admin user specified with **ADMIN_USER** should be established with:

```
C:\> p4 -p my-helix-core-server:1666 -u admin_userid login
```

If the ticket was established using the wrong port, the error message you encounter includes the port that the trigger is attempting to use:

```
'swarm.strict.1' validation failed: Invalid login credentials
to [port] within this trigger script; please contact your
administrator.
```

5. **Linux only:** Ensure that the script has execute permissions.

Important

Skip this step if you have committed the script to the Helix Core Server.

```
$ chmod +x /path/to/swarm-trigger.pl
```

6. Verify that the trigger script executes correctly.

- **For Windows:**

Run:

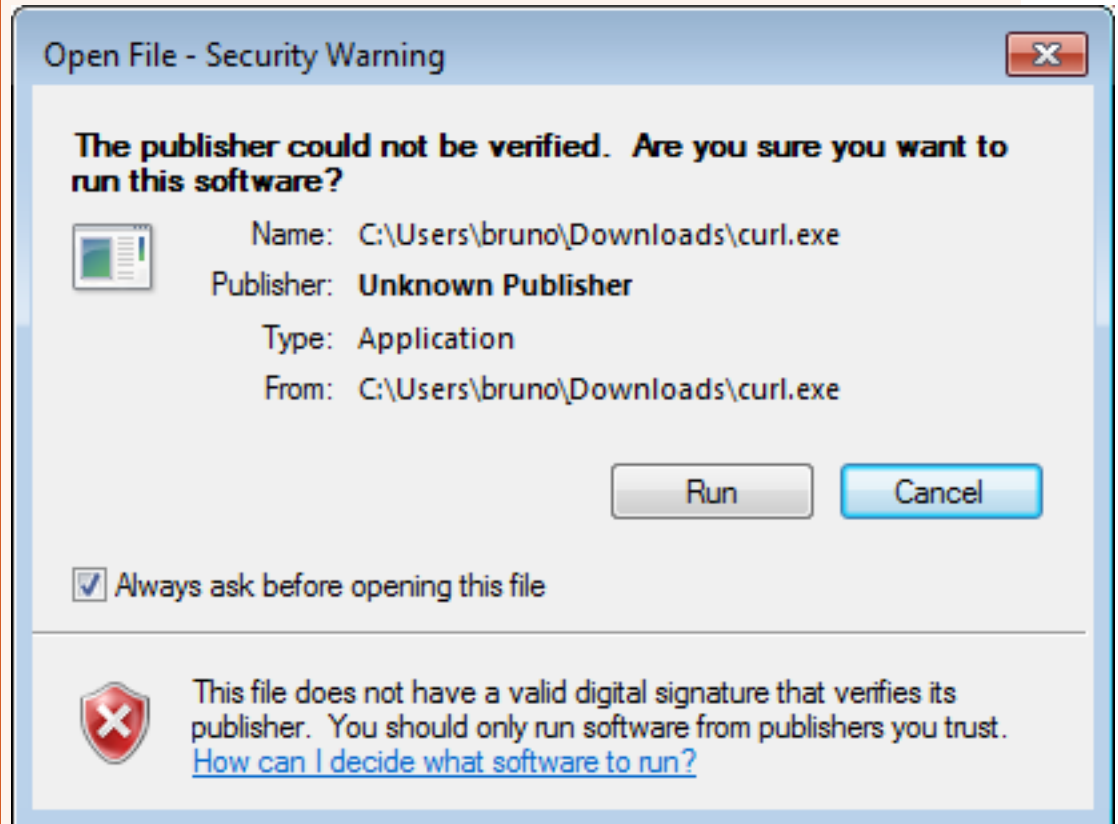
```
C:\> perl "C:\path\to\swarm-trigger.pl" -t ping -v 0
```

Use the full path to **perl** if it is not available in your command path.

You should expect to see no output. If the trigger is misconfigured, such as using an invalid [trigger token](#), you would see an error.

Warning**Only if using Curl:**

Installation of the triggers may cause a security warning dialog to appear when `curl.exe` executes:



If this occurs, the triggers hang, creating zombie Perl processes. Due to the way triggers are invoked by Helix Core Server, the dialog is normally not visible even though Windows is waiting on interaction.

To resolve this:

- a. Clear the **Always ask before opening this file** check box and click **Run**.
- b. Right-click `curl.exe`, select **Properties**, and click **Unblock**.

■ For Linux:**Important**

Skip this step if you have committed the script to the Helix Core Server.

```
$ /path/to/swarm-trigger.pl -t ping -v 0
```

You should expect to see no output. If the trigger is misconfigured, such as using an invalid [trigger token](#), you would see an error.

Note

Run the trigger script without any arguments to see additional usage information.

7. Update the Helix Core Server triggers table to run the trigger script.

Warning

- The `swarm.shelvedel shelve-delete` trigger line was added to Swarm in version 2018.1 and updated in version 2020.1.
 - **New installation or upgrading from Swarm 2017.4 and earlier:** add the `swarm.shelvedel shelve-delete` trigger line to the Helix Core Server trigger table if it is not already present.
 - **Upgrading from Swarm 2018.x and 2019.x:** replace the existing `swarm.shelvedel shelve-delete` trigger line in the Helix Core Server trigger table with the one supplied in the Swarm version you are upgrading to.
- **Workflow feature:**

The [Workflow feature](#) is enabled by default in Swarm 2019.2 and later. The trigger lines required when workflow is enabled are different to those required when workflow is disabled:

- **Workflow feature [enabled](#) (default):**
 - Comment out the `swarm.enforce.1`, `swarm.enforce.2`, `swarm.strict.1`, and `swarm.strict.2` trigger lines in the Helix Core Server trigger table if they are present.
 - Add the `swarm.enforce change-submit`, `swarm.strict change-content`, and `swarm.shelvesub shelve-submit` trigger lines to the Helix Core Server trigger table if they are not already present.
- **Workflow feature [disabled](#):**
 - Comment out the `swarm.enforce change-submit`, `swarm.strict change-content`, and `swarm.shelvesub shelve-submit` trigger lines in the Helix Core Server trigger table if they are present.
 - The `swarm.enforce.1`, `swarm.enforce.2`, `swarm.strict.1`, and `swarm.strict.2` trigger lines are commented out because they are optional. They require that the `DEPOT_PATH1` and `DEPOT_PATH2` values are configured appropriately. Support for these triggers will be dropped in a later release.
 - The first two lines configure the **enforce** feature, which rejects any submitted changes that are not tied to an approved review.
 - The second two lines configure the **strict** feature, which rejects any submitted changes when the contents of the changelist do not match the contents of its associated approved review.
 - If you need to apply **enforce** or **strict** to more depot paths, copy the lines and tweak the depot paths as necessary.

Tip

The trigger script can provide the list of trigger lines that should work, with little to no adjustment, by executing it with the `-o` option:

Windows

```
C:\> perl "C:\path\to\swarm-trigger.pl" -o
```

Linux

```
$ /path/to/swarm-trigger.pl -o
```


Tip

- The `%quote%c:\path\to\perl.exe%quote%` entry is only required for Windows-hosted systems.
- An initial tabbed indent is required for each trigger line.
- Commented out trigger lines are not stored.
- If you disable the workflow feature in the Swarm `config.php` file, workflow will not be processed by Swarm but a small overhead is still incurred by the Helix Core Server each time it runs a workflow trigger script. This overhead can be eliminated by commenting out the `swarm.enforce change-submit`, `swarm.strict change-content`, and `swarm.shelvesub shelve-submit` workflow triggers.

As a Helix Core Server user with *super* privileges, edit the Helix Core Server trigger table by running the `p4 triggers` command and add the following lines (including the initial tab character). Update the `perl.exe`, trigger script, and configuration file paths in each line below to reflect the actual paths on your Helix Core Server:

- If you have committed both the trigger script and the configuration file to the Helix Core Server:

```

    swarm.job          form-commit    job
"%quote%c:\path\to\perl.exe%quote% %//.swarm/triggers/swarm-
trigger.pl% -c %//.swarm/triggers/swarm-trigger.conf% -t job
-v %formname%"
    swarm.user         form-commit    user
"%quote%c:\path\to\perl.exe%quote% %//.swarm/triggers/swarm-
trigger.pl% -c %//.swarm/triggers/swarm-trigger.conf% -t user
-v %formname%"
    swarm.userdel      form-delete    user
"%quote%c:\path\to\perl.exe%quote% %//.swarm/triggers/swarm-
trigger.pl% -c %//.swarm/triggers/swarm-trigger.conf% -t
userdel      -v %formname%"
    swarm.group        form-commit    group
"%quote%c:\path\to\perl.exe%quote% %//.swarm/triggers/swarm-
trigger.pl% -c %//.swarm/triggers/swarm-trigger.conf% -t group
-v %formname%"
    swarm.groupdel     form-delete    group

```

```

"%quote%c:\path\to\perl.exe%quote% %//.swarm/triggers/swarm-
trigger.pl% -c %//.swarm/triggers/swarm-trigger.conf% -t
groupdel -v %formname%"

    swarm.changesave form-save change
"%quote%c:\path\to\perl.exe%quote% %//.swarm/triggers/swarm-
trigger.pl% -c %//.swarm/triggers/swarm-trigger.conf% -t
changesave -v %formname%"

    swarm.shelve shelve-commit //...
"%quote%c:\path\to\perl.exe%quote% %//.swarm/triggers/swarm-
trigger.pl% -c %//.swarm/triggers/swarm-trigger.conf% -t
shelve -v %change%"

    swarm.commit change-commit //...
"%quote%c:\path\to\perl.exe%quote% %//.swarm/triggers/swarm-
trigger.pl% -c %//.swarm/triggers/swarm-trigger.conf% -t
commit -v %change%"

    swarm.shelvedel shelve-delete //...
"%quote%c:\path\to\perl.exe%quote% %//.swarm/triggers/swarm-
trigger.pl% -c %//.swarm/triggers/swarm-trigger.conf% -t
shelvedel -v %change% -w %client% -u %user% -d
%quote%%clientcwd%^^^%quote% -a %quote%%argsQuoted%%quote% -s
%quote%%serverVersion%%quote%"
# The following three triggers are used by workflow. If
workflow is disabled in the Swarm
# configuration then they should be disabled.
    swarm.enforce change-submit //...
"%quote%c:\path\to\perl.exe%quote% %//.swarm/triggers/swarm-
trigger.pl% -c %//.swarm/triggers/swarm-trigger.conf% -t
checkenforced -v %change% -u %user%"

    swarm.strict change-content //...
"%quote%c:\path\to\perl.exe%quote% %//.swarm/triggers/swarm-
trigger.pl% -c %//.swarm/triggers/swarm-trigger.conf% -t
checkstrict -v %change% -u %user%"

    swarm.shelvesub shelve-submit //...

```

```

"%quote%c:\path\to\perl.exe%quote% %//.swarm/triggers/swarm-
trigger.pl% -c %//.swarm/triggers/swarm-trigger.conf% -t
checkshelve -v %change% -u %user%"
#      The following triggers are only used to prevent a commit
without an approved review.
#      They predate the workflow functionality and should only be
used if workflow is disabled.
#      Support for these will be dropped in a later release.
#      swarm.enforce.1 change-submit //DEPOT_PATH1/...
"%quote%c:\path\to\perl.exe%quote% %//.swarm/triggers/swarm-
trigger.pl% -c %//.swarm/triggers/swarm-trigger.conf% -t
enforce -v %change% -p %serverport%"
#      swarm.enforce.2 change-submit //DEPOT_PATH2/...
"%quote%c:\path\to\perl.exe%quote% %//.swarm/triggers/swarm-
trigger.pl% -c %//.swarm/triggers/swarm-trigger.conf% -t
enforce -v %change% -p %serverport%"
#      swarm.strict.1 change-content //DEPOT_PATH1/...
"%quote%c:\path\to\perl.exe%quote% %//.swarm/triggers/swarm-
trigger.pl% -c %//.swarm/triggers/swarm-trigger.conf% -t
strict -v %change% -p %serverport%"
#      swarm.strict.2 change-content //DEPOT_PATH2/...
"%quote%c:\path\to\perl.exe%quote% %//.swarm/triggers/swarm-
trigger.pl% -c %//.swarm/triggers/swarm-trigger.conf% -t
strict -v %change% -p %serverport%"

```

- b. If you have copied the trigger script and configuration file to common paths on all servers:

```

        swarm.job          form-commit    job
"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t job          -v %formname%"

        swarm.user        form-commit    user
"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t user         -v %formname%"

        swarm.userdel     form-delete    user
"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t userdel      -v %formname%"

        swarm.group       form-commit    group
"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t group        -v %formname%"

        swarm.groupdel    form-delete    group
"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t groupdel     -v %formname%"

        swarm.changesave  form-save      change
"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t changesave   -v %formname%"

        swarm.shelve      shelve-commit  //...
"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t shelve       -v %change%"

        swarm.commit      change-commit  //...
"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t commit       -v %change%"

        swarm.shelvedel   shelve-delete  //...

```

```

"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t shelvedel -v %change% -w %client% -u %user% -d
%quote%%clientcwd%^^^%quote% -a %quote%%argsQuoted%%quote% -s
%quote%%serverVersion%%quote%"
#       The following three triggers are used by workflow. If
workflow is disabled in the Swarm
#       configuration then they should be disabled.
        swarm.enforce      change-submit  //...
"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t checkenforced -v %change% -u %user%"
        swarm.strict      change-content  //...
"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t checkstrict  -v %change% -u %user%"
        swarm.shelvesub  shelve-submit  //...
"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t checkshelve  -v %change% -u %user%"
#       The following triggers are only used to prevent a commit
without an approved review.
#       They predate the workflow functionality and should only be
used if workflow is disabled.
#       Support for these will be dropped in a later release.
#       swarm.enforce.1 change-submit  //DEPOT_PATH1/...
"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t enforce -v %change% -p %serverport%"
#       swarm.enforce.2 change-submit  //DEPOT_PATH2/...
"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t enforce -v %change% -p %serverport%"

```

```
#      swarm.strict.1  change-content //DEPOT_PATH1/...
"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t strict -v %change% -p %serverport%"
#      swarm.strict.2  change-content //DEPOT_PATH2/...
"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t strict -v %change% -p %serverport%"
```

Warning
Windows and Linux:

The use of `%quote%` is not supported on 2010.2 servers (it is harmless though); if you are using this version, ensure that you do not have any spaces in the path to `perl.exe` or the script's path.

Important
Workflow feature disabled:

If your Helix Core Server has SSL enabled and is older than the 2014.1 release, the `%serverport%` trigger variable does not include the necessary transport indicator, which can cause the `enforce` and `strict` triggers to fail.

To solve this problem, add `ssl:` immediately before `%serverport%` in the trigger lines. For example:

```
#      The following triggers are only used to prevent a commit
#      without an approved review.
#      They predate the workflow functionality and should only be
#      used if workflow is disabled.
#      Support for these will be dropped in a later release.
#      swarm.enforce.1 change-submit //DEPOT_PATH1/...
"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t enforce -v %change% -p ssl:%serverport%"
#      swarm.enforce.2 change-submit //DEPOT_PATH2/...
"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t enforce -v %change% -p ssl:%serverport%"
#      swarm.strict.1  change-content //DEPOT_PATH1/...
"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t strict -v %change% -p ssl:%serverport%"
#      swarm.strict.2  change-content //DEPOT_PATH2/...
"%quote%C:\path\to\perl.exe%quote% %quote%/path/to/swarm-
trigger.pl%quote% -c %quote%/path/to/swarm-trigger.conf%quote%
-t strict -v %change% -p ssl:%serverport%"
```

8. Configure the Helix Core Server to promote all shelved changes.

When this configurable is set, Swarm has access to all shelved changelists, which is a requirement for pre-commit reviews. When it is not set, users connected to an edge server must remember to use the `-p` option when shelving files to promote their shelves to the commit server when initiating a pre-commit review.

```
p4 configure set dm.shelve.promote=1
```

9. Optionally forward logins to the commit server.

If you intend to use P4V and its Swarm integration, you should consider forwarding logins to the commit server. See ["P4V Authentication" on page 588](#) for details.

10. Depending on your installation type, do one of the following:

- **Swarm package installation:** Review the post-install configuration options to customize your Swarm installation, see ["Post-install configuration options" on page 224](#).
- **Swarm tarball installation:** Set up a recurring task to spawn workers, see ["Set up a recurring task to spawn workers" below](#).

Set up a recurring task to spawn workers

To ensure that incoming Helix Core Server events are automatically processed by Swarm, it is important to set up a cron job to do this. We recommend that the cron job is installed on the Swarm host machine.

Install curl or wget

You must install either **curl** or **wget** on the Swarm server machine to run the recurring task that spawns the Swarm workers:

- To install **curl**, see ["Install curl" below](#).
- To install **wget**, see ["Install wget" on the next page](#).

Install curl

1. Download **curl** from: <https://curl.haxx.se/download.html>
2. Install **curl**.
3. Verify that **curl** is installed, see ["Verify that curl or wget is installed" on the next page](#).

Warning

If **curl** cannot execute as expected, trigger execution may block or fail. Prior to configuring the [triggers](#), verify that **curl** executes, run:

```
$ curl -h
```


The start of the output should be similar to:

```
Usage: curl [options...] <url>
Options: (H) means HTTP/HTTPS only, (F) means FTP only
  --anyauth          Pick "any" authentication method (H)
  -a, --append       Append to target file when uploading (F/SFTP)
  --cacert FILE      CA certificate to verify peer against (SSL)
  --capath DIR       CA directory to verify peer against (SSL)
...[truncated for brevity]...
```

For a more thorough test that actually fetches content over a network, try the following test:

```
$ curl https://www.perforce.com/
```

The output should look like HTML.

Install wget

1. Download **wget** from: <https://ftp.gnu.org/gnu/wget/>
2. Install **wget**.
3. Verify that **wget** is installed, see "[Verify that curl or wget is installed](#)" below.

Verify that curl or wget is installed

Warning

curl or **wget** must be installed or workers do not spawn and Swarm cannot process any events.

Verify that curl or wget is installed:

1. Use the **which** command.

For example to verify that **curl** is installed:

```
$ which curl
```

2. If you see any output, the referenced command is installed.

Create a recurring task to spawn workers

Note

Swarm package installation:

- The **helix-swarm**, **swarm-cron-hosts.conf**, and **swarm-cron.sh** files are automatically created with the correct content including the correct URL for Swarm. They do not need to be edited.
- If you want to run Swarm in a sub-folder of an existing website, or with a custom port, the **swarm-cron-host.conf** file must be edited. For instructions on how to configure Swarm to run in a sub-folder, or with a custom port, see [Run Swarm in a sub-folder of an existing web site](#), or [Run Swarm's virtual host on a custom port](#).

helix-swarm cron file

The **helix-swarm** file is located in the **/etc/cron.d** and points to the **swarm-cron.sh** script file.

1. Create a file named **helix-swarm** in **/etc/cron.d** if it does not already exist.
2. Edit the **helix-swarm** file so that it has the following content:

```
#
# Cron job to start Swarm workers every minute
#
* * * * * nobody [ -x /opt/perforce/swarm/p4-bin/scripts/swarm-
cron.sh ] && /opt/perforce/swarm/p4-bin/scripts/swarm-cron.sh
```

3. Save the **helix-swarm** file.

swarm-cron.sh script

The **swarm-cron.sh** file in **/opt/perforce/swarm/p4-bin/scripts** ensures that a worker is fired up every minute.

By default, the **DEFAULT_CONFIG_FILE=** configuration line is set to **/opt/perforce/etc/swarm-cron-hosts.conf**. The **swarm-cron.sh** script takes the Swarm hostname from the **swarm-cron-hosts.conf** file.

The **swarm-cron.sh** script contains the correct content for Swarm installations using **wget**, or **curl**.

Tip

You must leave the **wget**, and **curl** configuration lines in the script. The script is written so that it can be used with Swarm installations that use **wget**, and Swarm installations that use **curl**. The **wget** configuration line is tried first, if that fails the **curl** configuration line is tried.

- **wget** configuration line for HTTP, and HTTPS:

```
wget --quiet --no-check-certificate --output-document /dev/null -
-timeout 5 "${SWARM_HOST}/queue/worker"
```

- **curl** configuration line for HTTP, and HTTPS:

```
curl --silent --insecure --output /dev/null --max-time 5
"${SWARM_HOST}/queue/worker"
```

In the **wget** and **curl** configuration lines above, where you see **--timeout 5** or **--max-time 5**, the **5** is the number of seconds that the cron task will wait for a response from the Swarm host. When the cron task is installed on the Swarm host, that value could be reduced to **1** second (e.g. **--timeout 1** or **--max-time 1**).

Note

If you configure Swarm to use HTTPS, and you install a self-signed certificate, the cron jobs avoid the certificate validity test which could cause silent failures to process events. The command to avoid the certificate validity test is included by default for **wget**, and **curl**:

- wget command: **--no-check-certificate**
- curl command: **--insecure**

If you have configured an HTTP connection to Swarm, the avoid validity check command is ignored.

swarm-cron-hosts.conf configuration file

The **swarm-cron-hosts.conf** file in **/opt/perforce/etc** specifies the connection type (HTTP or HTTPS), hostname, and port number for Swarm cron jobs.

1. Edit the **swarm-cron-hosts.conf** file so that it contains the actual Swarm hostname, and port you have configured for Swarm (which may include a [sub-folder](#) or a [custom port](#)). The following format is used:

```
[http[s]://]<swarm-host>[:<port>]
```

Default if value not specified:

- **[http[s]://]** http
 - **<swarm-host>** must be specified
 - **:<port>** 80
2. Save the edited file.

Tip

To check or modify Swarm worker configuration, see ["Workers" on page 749](#).

Note

If the recurring task is disabled, or stops functioning for any reason, logged-in users see the following error message when Swarm detects that no workers are running:

Hmm... no queue workers? Ask your administrator to check the [worker setup](#).

3. The basic Swarm configuration is now complete.

Important

If your Helix Core Server is configured for Helix Authentication Service, you can force all of your users to authenticate via your Identity Provider (IdP) by disabling fall-back to passwords. To disable fall-back to passwords on the Helix Core Server, run the following command:

```
p4 configure set auth.sso.allow.passwd=0
```

4. Review the post-install configuration options to customize your Swarm installation, see "[Post-install configuration options](#)" below.

Common post-installation options

This chapter covers the post-installation options that are common for all installation processes.

Post-install configuration options

There are a few options for customizing your Swarm installation's operation. This section covers the options that are officially supported.

This section contains:

- "No customization required" on the next page
- "Back up your Swarm virtual host first" on the next page
- "HTTPS" on the next page
- "Run Swarm in a sub-folder of an existing web site" on page 230
- "Run the Swarm virtual host on a custom port" on page 233
- "Use your own Redis server" on page 235
- "Hiding Swarm storage from regular users" on page 238
- "Handling Exclusive Locks" on page 238

No customization required

If you do not need to customize your Swarm installation, your Swarm installation is complete but you must check that it is working correctly before using it in production. For instructions on how to check your Swarm installation, see ["Validate your Swarm installation" on page 239](#).

Back up your Swarm virtual host first

Before undertaking any of the following customization options, ensure that you have backed up your Swarm virtual host configuration. Choose the most appropriate option:

- If your Apache configuration directory contains the directories **sites-available** and **sites-enabled**:

```
$ cd /path/to/apache/configuration/..
$ cp -a sites-available sites-available.bak
```

Important

If the **sites-enabled** directory contains files, and not just symbolic links, you need to backup this folder as well:

```
$ cd /path/to/apache/configuration/..
$ cp -a sites-enabled sites-enabled.bak
```

- For CentOS/RHEL systems, if you used the [Swarm packages](#) to install Swarm:

```
$ cd /path/to/apache/configuration/..
$ cp -a conf.d conf.d.bak
```

- Otherwise, back up your Apache configuration.

HTTPS

Important

Back up your Swarm virtual host configuration before customizing your Swarm installation, see ["Back up your Swarm virtual host first" above](#).

This section describes how to make your Swarm installation more secure by using HTTPS.

Before you begin the following procedure, locate your system's Apache configuration. Common configuration directories include:

- **/etc/httpd/conf/**
- **/etc/apache2/**

- `/usr/local/apache2/conf/`
- `/Applications/XAMPP/etc/`

Within the Apache configuration path, the main Apache configuration file is usually named one of the following:

- `httpd.conf`
- `apache2.conf`

A longer discussion on the possible locations and names of Apache configuration files is available here: <https://cwiki.apache.org/confluence/display/httpd/DistrosDefaultLayout>

1. Enable SSL in Apache.

If the Apache utility `a2enmod` is installed:

```
$ sudo a2enmod ssl
```

Without the `a2enmod` utility, edit the Apache configuration file by hand. Locate your Apache configuration file for modules and either uncomment or add the following lines:

```
LoadModule ssl_module libexec/apache2/mod_ssl.so
```

2. Create a directory to store certificates.

```
$ sudo mkdir -p /etc/apache2/ssl
```

3. Create a certificate/key pair.

```
$ cd /etc/apache2/ssl
```

```
$ sudo openssl req -x509 -nodes -days 365 -newkey rsa:2048 -
keyout apache.key -out apache.crt
```

This command generates a private key and a certificate. To form the certificate, `openssl` prompts you for several details:

```
Generating a 2048 bit RSA private key
.....+++
.....+++
writing new private key to 'apache.key'
-----

You are about to be asked to enter information that will be
incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished
Name or a DN.
```

```
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [AU]:CA
State or Province Name (full name) [Some-State]:British Columbia
Locality Name (eg, city) []:Victoria
Organization Name (eg, company) [Internet Widgits Pty
Ltd]:Perforce Software
Organizational Unit Name (eg, section) []:Swarm development team
Common Name (e.g. server FQDN or YOUR name) []:myswarm.host
Email Address []:admin@myswarm.host
```

The output above includes some example details. You should replace anything in italics with your own details. Since the certificate request details that can help users determine whether your certificate is valid, enter legitimate information whenever possible.

Important

The Common Name field must match the hostname for your Swarm installation exactly.

4. Secure the certificate directory.

```
$ sudo chmod 600 /etc/apache2/ssl
```

5. Edit the virtual host configuration.

Note

The virtual host configuration should be in the file you [backed up initially](#).

Edit the virtual host configuration to match:

```
<VirtualHost *:80>
    ServerName myswarm
    AllowEncodedSlashes NoDecode
    ServerAlias myswarm.host
    ErrorLog "/path/to/apache/logs/myswarm.error_log"
    CustomLog "/path/to/apache/logs/myswarm.access_log" common
    DocumentRoot "/path/to/swarm/public"
    <Directory "/path/to/swarm/public">
        AllowOverride All
        Require all granted
    </Directory>

    Redirect / https://myswarm.host/
</VirtualHost>

<VirtualHost *:443>
    SSLEngine on
    SSLCertificateFile /etc/apache2/ssl/apache.crt
    SSLCertificateKeyFile /etc/apache2/ssl/apache.key

    ServerName myswarm.host
    AllowEncodedSlashes NoDecode
    ServerAlias myswarm
    ErrorLog "/path/to/apache/logs/myswarm.error_log"
    CustomLog "/path/to/apache/logs/myswarm.access_log" common
    DocumentRoot "/path/to/swarm/public"
    <Directory "/path/to/swarm/public">
        AllowOverride All
        Require all granted
    </Directory>
</VirtualHost>
```

See Apache's virtual host documentation for details: <https://httpd.apache.org/docs/2.4/vhosts/>

6. Customize the virtual host definition.

- a. Replace `myswarm.host` with the hostname for Swarm on your network.
- b. Replace `myswarm` with the name of the subdomain hosting Swarm. Many administrators choose `swarm`.

Note

The string `myswarm` in the log file paths: this should match the subdomain name and prefix for the log files, to help coordinate the active host with the log files for that host. Doing this is particularly useful when your Apache server hosts multiple instances of Swarm.

- c. Replace `/path/to/apache/logs` with the path where your Apache store its log files. Apache's log files are typically named `access_log` and `error_log`.
 - d. Replace `/path/to/swarm` with the path to the Swarm directory.
7. Restart your web server.

```
$ sudo apachectl restart
```

8. Adjust your firewall configuration to allow connections to the standard SSL port for web servers.

- For CentOS/RHEL:

```
$ sudo firewall-cmd --zone=public --add-port=443/tcp --  
permanent  
$ sudo systemctl reload firewalld
```

- For other distributions, consult with your network administrator or operating system documentation to determine how to adjust your firewall configuration.

9. Test your HTTPS URL from a web browser.

Important

If the `myswarm.host` value in the virtual host configuration and the certificate do not match, the P4V integration with Swarm fails with the message **SSL handshake failed**.

Also, when a reverse DNS lookup is performed, `myswarm.host` should be the answer when querying for the Swarm server's IP address.

10. Validate that Swarm is working correctly before using it in production, see ["Validate your Swarm installation" on page 239](#).

Run Swarm in a sub-folder of an existing web site

Warning

Multiple Helix Core Server instances on a single Swarm instance: You cannot run Swarm in a sub-folder if Swarm is connected to multiple Helix Core Servers.

Issue: Swarm will lose connection to all of the Helix Core Servers if you edit the `base_url` configurable value in the `environment` block of `<swarm_root>/data/config.php`. This will stop your system working.

Fix: Remove the `base_url` configurable from the `environment` block of `<swarm_root>/data/config.php`.

Important

Back up your Swarm virtual host configuration before customizing your Swarm installation, see ["Back up your Swarm virtual host first" on page 225](#).

If you cannot run Swarm in its own virtual host, which might be necessary when you do not control the hostname to be used with Swarm, installing Swarm in a sub-folder of an existing virtual host configuration can be a good solution.

Installing Swarm in a sub-folder requires modification of the previous installation steps covered in this chapter:

- The ["Apache configuration" on page 186](#) is entirely different; instead of establishing a new virtual host, you need to [modify an existing virtual host configuration](#). Often, this would be Apache's default site.
- [Swarm's configuration file](#) requires an [extra item](#).

The following sections cover the specifics of sub-folder installation.

See ["base_url" on page 611](#) for more details.

Important

If you used the [Swarm packages](#) to install Swarm, you can adjust Swarm's configuration using the package configuration script `/opt/perforce/swarm/sbin/configure-swarm.sh`.

`configure-swarm.sh` does not read any existing Swarm configuration; you must provide all of the configuration details each time you execute `configure-swarm.sh`:

```
$ sudo /opt/perforce/swarm/sbin/configure-swarm.sh -n -p  
myp4host:1666 -u swarm -w password -e mx.example.com -H myhost -B  
/swarm
```

In the example above, the `-B` option is used to specify the name of the sub-folder.

If you use `configure-swarm.sh` to adjust the Swarm configuration, you only need to follow the "Apache configuration" below steps described below; all of the changes listed in the "Swarm configuration" on the facing page section below have been completed by `configure-swarm.sh`.

Apache configuration

1. Ensure that the `SWARM_ROOT` is not within the document root of the intended virtual host.

This step ensures that Swarm's source code and configuration is impossible to browse, preventing access to important details such as stored credentials, and active sessions and workspaces.

2. Adjust the virtual host configuration that you are already using.

Note

Depending on the method used to install Swarm, the filename for virtual host configuration you need to edit is:

- For [Swarm package](#) installations, edit `perforce-swarm-site.conf`.
- For manual installations following Swarm's recommended "Apache configuration" on [page 186](#) edit `swarm`.
- For other installations, you may have to edit `httpd.conf` or nearby files.

Add the following lines to the virtual host definition:

```
Alias "/swarm" "SWARM_ROOT/public"

DocumentRoot "/var/www/html"

<Directory "SWARM_ROOT/public">
    AllowOverride All
    Require all granted
</Directory>
```

The **Alias** line configures Apache to respond to requests to `https://myhost/swarm` with content from Swarm's **public** folder. You can change the `/swarm` portion of the Alias line to anything you want.

The **DocumentRoot** line configures the Apache DocumentRoot directory location.

The **<Directory>** block grants access to everything within Swarm's **public** folder. Replace `SWARM_ROOT` with the actual path to Swarm.

3. Restart your web server.

```
$ sudo apachectl restart
```

Swarm configuration

To successfully operate within a sub-folder, the `"swarm_root"` on [page 724/data/config.php](#) file needs to be adjusted to contain the following lines (as a peer of the p4 item):

```
'environment' => array(
    'base_url' => '/swarm'
),
```

Ensure that `/swarm` matches the first item in the `Alias` line in the virtual host configuration.

See ["Environment" on page 609](#) for more details.

Swarm trigger configuration

The `swarm-trigger.conf` file must be updated to include the `base_url` in the `SWARM_HOST` path.

Tip

The location of the `swarm-trigger.conf` depends on your installation. For information on the location of the `swarm-trigger.conf` file, see ["Set up Swarm triggers with a Windows or Linux-hosted Helix Core Server" on page 205](#).

1. Edit the `swarm-trigger.conf` file.
2. Replace:

```
SWARM_HOST="http://my-swarm-host"
```

With:

```
SWARM_HOST="http://website-hostname/swarm"
```

3. Save the file.

Cron configuration

Swarm's [recurring task](#) configuration must be updated to reflect the sub-folder that you have configured in [Apache's](#) and [Swarm's](#) configurations. This is configured in the `swarm-cron-hosts.conf` file.

1. Edit `/opt/perforce/etc/swarm-cron-hosts.conf`.
2. Replace:

```
https://swarm-host
```

with:

```
https://swarm-host/swarm/
```

Where **swarm-host** is the hostname of your Swarm installation, and **swarm** is the sub-folder you want to use.

3. Save the edited file.

New workers should be started at the start of the next minute.

4. Validate that Swarm is working correctly before using it in production, see ["Validate your Swarm installation" on page 239](#).

Run the Swarm virtual host on a custom port

Important

Back up your Swarm virtual host configuration before customizing your Swarm installation, see ["Back up your Swarm virtual host first" on page 225](#).

If you cannot run Swarm on port 80 (or port 443 for HTTPS), perhaps because you do not have root access, it is possible to run Swarm on a custom port.

Installing Swarm to use a custom port requires modification of the previous installation steps covered in this chapter: The [Apache configuration](#) is slightly different, requiring [modification of Swarm's virtual host definition](#).

The following section covers the specifics of the custom port configuration.

Note

In addition to the following instructions, you may also need to apply the **external_url** item described in the ["Environment" on page 609](#) section if your Swarm is behind a proxy, or you have multiple Swarm instances connected to Helix Core Server.

Important

If you used the [Swarm packages](#) to install Swarm, you can adjust Swarm's configuration using the package configuration script `/opt/perforce/swarm/sbin/configure-swarm.sh`.

configure-swarm.sh does not read any existing Swarm configuration; you must provide all of the configuration details each time you execute **configure-swarm.sh**:

```
$ sudo /opt/perforce/swarm/sbin/configure-swarm.sh -n -p
myp4host:1666 -u swarm -w password -e mx.example.com -H myhost -P
8080
```

In the example above, the **-P** option is used to specify the custom port that Swarm should use.

If you use `configure-swarm.sh` to adjust Swarm's configuration, follow the additional steps that it describes. Once those steps are complete, do not perform any of the steps described below.

Apache configuration

1. Edit the virtual host configuration.

Note

Depending on the method used to install Swarm, the filename for virtual host configuration you need to edit is:

- For [Swarm package](#) installations, edit `perforce-swarm-site.conf`.
- For manual installations following Swarm's [recommended Apache configuration](#), edit `swarm`.
- For other installations, you may have to edit `httpd.conf` or nearby files.

- a. Add the following line *outside* of the `<VirtualHost>` block:

```
Listen 8080
```

- b. Edit the `<VirtualHost *:80>` line to read:

```
<VirtualHost *:8080>
```

For both lines, replace `8080` with the custom port you wish to use.

Important

If you choose a port that is already in use, Apache refuses to start.

2. Restart your web server.

```
$ sudo apachectl restart
```

3. Adjust your firewall configuration to allow connections to the custom port.

- For CentOS/RHEL:

```
$ sudo firewall-cmd --zone=public --add-port=8080/tcp --
permanent
$ sudo systemctl reload firewalld
```

Replace `8080` with the custom port you wish to use.

- For other distributions, consult with your network administrator or operating system documentation to determine how to adjust your firewall configuration.

Cron configuration

Swarm's [recurring task](#) configuration must be updated to reflect the custom port that you have configured in [Apache's](#) configuration. This is configured in the `swarm-cron-hosts.conf` file.

1. Edit `/opt/perforce/etc/swarm-cron-hosts.conf`.
2. Replace:

```
https://swarm-host:80
```

with:

```
https://swarm-host:8080
```

Where `swarm-host` is the hostname of your Swarm installation, and `8080` is the custom port you want to use.

3. Save the edited file.
New workers should be started at the start of the next minute.
4. Validate that Swarm is working correctly before using it in production, see ["Validate your Swarm installation" on page 239](#).

Use your own Redis server

By default Swarm uses its own Redis server on the Swarm machine.

This section describes how to connect Swarm to your own Redis server.

To use your own Redis server:

1. Add the `redis` block to the `$SWARM_ROOT/data/config.php` file, it contains the `password`, `namespace`, `host`, and `port` details of your Redis server

```
<?php
// this block should be a peer of 'p4'
'redis' => array(
    'options' => array(
        'password' => null, // Defaults to null
        'namespace' => 'Swarm',
        'server' => array(
            'host' => 'MyRedisServerHostname', // Defaults to
'localhost' or enter your Redis server hostname
            'port' => 'MyRedisServerPortNumber', // Defaults
to '7379 or enter your Redis server port
        ),
    ),
```

```

        'items_batch_size' => 100000,
        'check_integrity' => '03:00', // Defaults to '03:00' Use
one of the following two formats:
                                // 1) The time of day that
the integrity check starts each day. Set in 24 hour format with
leading zeros and a : separator
                                // 2) The number of seconds
between each integrity check. Set as a positive integer. Specify
'0' to disable the integrity check.
        'population_lock_timeout' => 300, // Timeout for initial
cache population. Defaults to 300 seconds.
    ),

```

- **password**: set to the password of your Redis server, defaults to null
- **namespace**: the prefix used for key values in the Redis cache. Defaults to **Swarm**.

Note

If you have multiple Swarm instances running against a single Redis server, each Swarm server must use a different Redis **namespace**. This enables the cache data for the individual Swarm instances to be identified. The **namespace** is limited to ≤ 128 characters.

If one or more of your Swarm instances is connected to multiple Helix Core Servers, the Redis **namespace** includes the **server label** and the character limit is reduced to ≤ 127 characters, see ["Multiple Helix Core Server instances" on page 646](#).

- **host** *MyRedisServerHostname*: Redis server host name
- **port** *MyRedisServerPortNumber*: Redis server port number
- **items_batch_size**: Maximum number of key/value pairs allowed in an **mset** call to Redis. Sets exceeding this will be batched according to this maximum for efficiency. Defaults to **100000**.

Note

The default value of **100000** was chosen to strike a balance between efficiency and project data complexity. This value should not normally need to be changed, contact support before making a change to this value.

- **check_integrity**: In some circumstances, such as when changes are made in the Helix Core Server when Swarm is down or if errors occur during updates, the Redis cache can get out of sync with the Helix Core Server. Swarm can run a regular integrity

check to make sure that the Redis caches and Helix Core Server are in sync. If an integrity check finds an out of sync cache file, Swarm automatically updates the data in that cache.

The `check_integrity` configurable specifies when the Redis cache integrity check is run. Set as a specific time of day (24 hour format with leading zeros) or a number of seconds (positive integer) between checks. Disable the integrity check with `'0'`. Defaults to `'03:00'`.

- `population_lock_timeout`: specifies the timeout, in seconds, for initial cache population. If you have a large Swarm system, increase this time if the initial cache population times out. Defaults to `300` seconds.
2. Edit the Redis server configuration in the `/opt/perforce/etc/redis-server.conf` file:

```
bind MyRedisServerHostname
port MyRedisServerPortNumber
supervised auto
save ""
dir /MyRedis/CacheDatabase/Directory/Location
```

- `bind MyRedisServerHostname` - Redis server host name
- `port MyRedisServerPortNumber` - Redis server port number
- `supervised auto` - detects the use of `upstart` or `systemd` automatically to signal that the process is ready to use the supervisors
- `save ""` - background saves disabled, recommended.
- `dir /MyRedis/CacheDatabase/Directory/Location` - the directory the Redis cache database is stored in

Tip

- The `redis-server.conf` file contains more detailed information about the Redis configuration for Swarm.
- On Swarm systems with a large number of users, groups, and projects, start-up time can be improved by persisting the memory cache. You can persist the memory cache by disabling background saves and enabling append saves, see the `redis-server.conf` file comments for detailed information.

Swarm is now connected to your Redis server and your Swarm installation is complete.

3. Validate that Swarm is working correctly before using it in production, see ["Validate your Swarm installation" on page 239](#).

Hiding Swarm storage from regular users

Swarm information storage uses Helix Core Server's **keys** facility. By default, users with *list*-level access can search keys and potentially obtain information they would not otherwise have access to, and users with *review*-level access can write or modify keys potentially corrupting or destroying data.

We recommend that you set the **dm.keys.hide** configurable to 2 to require *admin*-level access for searching and modifying keys. Note that **dm.keys.hide** is available in Helix Core Server versions 2013.1 and newer.

When **dm.keys.hide** is set to 2, both the **p4 keys** and **p4 key** commands require *admin*-level access in the Helix Core Server. When **dm.keys.hide** is set to 1, only the **p4 keys** command requires *admin*-level access in the Helix Core Server. When **dm.keys.hide** is set to 1, or is not set, users who know (or can deduce) key names can read values (if they have *list*-level access) or write values (if they have *review*-level access) with the **p4 key** command.

To set **dm.keys.hide**:

```
p4 configure set dm.keys.hide=2
```

To confirm the current value of **dm.keys.hide**:

```
p4 configure show dm.keys.hide
```

To unset **dm.keys.hide**:

```
p4 configure unset dm.keys.hide
```

Next step

Validate that Swarm is working correctly before using it in production, see ["Validate your Swarm installation" on the next page](#).

Handling Exclusive Locks

Tip

If this setting is not enabled in Helix Core Server, Swarm will report exceptions when working with exclusively opened files similar to **Cannot unshelve review (x). One or more files are exclusively open**, and that you must have the **filetype.bypasslock** configurable enabled.

Swarm takes copies of files when it is creating reviews. Some of the files managed by Helix Core Server may be limited to 'exclusive open' by having the filetype modifier '+' set. This file-level setting ensures only one user at a time can open the file for editing.

To allow Swarm to work with these 'exclusive open' files, you must enable **filetype.bypasslock** in the Helix Core Server configuration.

To set **filetype.bypasslock**:

```
p4 configure set filetype.bypasslock=1
```

To confirm the current value of **filetype.bypasslock**:

```
p4 configure show filetype.bypasslock
```

To unset `filetype.bypasslock`:

```
p4 configure unset filetype.bypasslock
```

Next step

Validate that Swarm is working correctly before using it in production, see ["Validate your Swarm installation" below](#).

Validate your Swarm installation

Tip

When Swarm starts it verifies the Redis cache, during this time you cannot log in to Swarm. The time taken to verify the Redis cache depends on the number of users, groups, and projects Swarm has. Start-up time can be improved by persisting the memory cache. You can persist the memory cache by disabling background saves and enabling append saves in the `redis-server.conf` file, see ["Redis server configuration file" on page 676](#).

Now that Swarm is fully installed the final step is to check that your Swarm installation is working correctly by doing the following:

1. Create a new changelist that:
 - a. Contains at least one modified file
 - b. Contains the `#review` keyword in the changelist description
2. Right click on the new changelist in P4V and click **Shelve Files...**

Important

Do not select **Request New Swarm Review...** because this method uses the API and will not fully test the Swarm Helix Core Server extension.

This is also true if you are using Swarm Helix Core Server Triggers instead of the Swarm Helix Core Server extension.

3. Check that a new Swarm review is created for the changelist.
 - If a Swarm review is created, the Swarm Helix Core Server extension is working. If you are using Swarm Helix Core Server Triggers instead of the Swarm Helix Core Server extension and the review is created, the triggers are working.
 - If a Swarm review is not created, see ["Review not created" on the facing page](#).

You are now all set to start using Swarm, enjoy!

Tip

To get started with Swarm, see the ["Quickstart" on page 29](#) chapter.

Optional Swarm integrations

Swarm provides a number of optional integrations that can be enabled if required, see ["Integrations" on page 552](#).

Review not created

If you are using the Swarm Helix Core Server extension

If a new Swarm review is not created when you validate your installation, your Swarm Helix Core Server extension is probably not working correctly on the Helix Core Server.

Check the Swarm Helix Core Server extension is working

To check that the Swarm Helix Core Server extension is working, run the following command on your Helix Core Server:

```
p4 extension --run swarm ping
```

Success message:

```
OK
```

Example failure message:

```
BAD (Cannot reach https://swarm-example.com/)
```

For information about Helix Core Server Extensions on your Helix Core Server, see the [Helix Core Extensions Developer Guide](#).

Check installed versions

To check your Swarm version and your Swarm Helix Core Server extension version, run the following command on your Helix Core Server:

```
p4 extension --run swarm version
```

Example response:

```
Swarm Version: SWARM/2022.1/2226999 (2022/03/20)
Extension Version: 2022.1.20221215
```

If you are using Swarm Helix Core Server Triggers

If a new Swarm review is not created when you validate your installation, your Swarm Helix Core Server Triggers are probably not installed correctly on the Helix Core Server. For instructions on how to install the Swarm Helix Core Server Triggers on your Helix Core Server, see ["Installing Triggers" on page 204](#).

Review cannot be updated when using triggers

This is rare but can sometimes happen if your Swarm server is configured for SSL.

Issue

The `swarm-trigger.pl` tries to use the Perl `HTTP::Tiny` module if it is installed. Sometimes it has a problem with certificates and a 599 error is generated.

To check if this is the issue, examine the operating system log of your Helix Core Server for errors logged by `swarm-trigger.pl`.

- Windows platforms, look in the Windows Event viewer
- Linux platforms:
 - **Ubuntu/Debian:** look in `/var/log/syslog`
 - **CentOS/RedHat:** look in `/var/log/messages`

The error message looks something like: **Error: (599/Internal Exception)**
(probably invalid SSL certificate)

Workaround

If you see this error, modify the Swarm trigger script file on the Helix Core Server so that it uses `Curl` instead of `HTTP::Tiny`.

Modify the trigger script to use Curl:

There are two lines in the `swarm-trigger.pl` script that enable the use of `HTTP::Tiny`, edit these two lines to make the script fallback to `Curl`:

1. Find both instances of the following line:

```
if ($HAVE_TINY) {
```
2. Add `!` to `$HAVE_TINY` for each line so that they look like this:

```
if (!$HAVE_TINY) {
```
3. Save the trigger script file.

Upgrading Swarm

The section describes how to upgrade a Swarm package or tarball installation to a newer release.

Note When upgrading to Helix Swarm 2022.3 or later, ensure that you upgrade Helix Core Visual Client (P4V) to 2021.3 or later.

Tip

If you are not already running Swarm, none of these instructions apply to you. Instead, see the [Swarm installation instructions](#).

Choose your upgrade process:

- "Upgrading a Ubuntu package installation" below.
- "Upgrading a CentOS/RHEL package installation" on page 254.
- "Upgrading a Docker container" on page 274.
- "Upgrading a tarball installation" on page 275.

Upgrading a Ubuntu package installation

The section describes how to upgrade a Swarm Ubuntu package installation to a newer release.

Important

- Swarm runtime dependencies change between releases, you must check that your system satisfies the Swarm runtime dependencies before starting the upgrade, see ["Runtime dependencies" on page 98](#).
- Review the PHP requirements before you upgrade Swarm, see ["PHP" on page 102](#).
- Review the Helix Core Server requirements before you upgrade Swarm, see ["Helix Core Server requirements" on page 106](#).
- Helix Core Server 2020.1 and later, permissions have changed for viewing and [editing stream spec files](#) in Swarm. To view and edit stream spec files in Swarm, the Swarm user must have *admin* permissions for the entire depot `// . . .`
- If you are upgrading from Swarm 2020.2 or earlier and have *userids* that contain the forward slash (/) character, add `AllowEncodedSlashes NoDecode` to the `VirtualHost` block of your `/etc/apache2/sites-enabled/perforce-swarm-site.conf` file. For more information about the `VirtualHost` block, see ["Apache configuration" on page 186](#).

Upgrade Swarm

Tip

From Swarm 2021.1, the Swarm package upgrade installs logrotate to manage your Swarm log rotation. If the package upgrade finds an existing custom logrotate configuration file for Swarm, the upgrade will notify you and give you details on how to disable the new logrotate configuration.

For information about the logrotate configuration, see ["Logrotate" on page 633](#).

The following process attempts to minimize downtime, but a short period of downtime for Swarm users is unavoidable. There should be no downtime for your Helix Core Server. After a successful upgrade, all Swarm users are logged out.

If you are using Swarm in a production environment, we encourage you to test this upgrade process in a non-production environment first.

1. Run the following commands:

```
sudo apt-get update
sudo apt-get install helix-swarm helix-swarm-triggers helix-
swarm-optional
```

2. Swarm generally has several major updates each year, and may occasionally have a patch update between major updates. To determine whether a Swarm update is available, run the following commands:

```
sudo apt-get update
sudo apt-get -s upgrade | grep swarm
```

Tip

Swarm uses a Redis server to manage its caches. This is installed and configured on the Swarm machine during the upgrade. If you prefer to use your own Redis server, see ["Use your own Redis server" on page 235](#).

Configuring Helix Core Server event notification

Swarm needs to know about a number of Helix Core Server events to operate correctly, this can be done by using Helix Core Server Extensions or Helix Core Server Triggers. Swarm installs include the Swarm Helix Core Server extension file and trigger scripts required for Swarm to get the events it needs from your Helix Core Server.

You must install the Swarm Helix Core Server extension or update your Triggers to complete the Swarm upgrade.

Do one of the following so that Swarm is notified about events on the Helix Core Server:

- Install and configure the Swarm Helix Core Server extension, see ["Installing the Swarm Helix Core Server extension" below](#)
- Update your Swarm Triggers, see ["Updating your triggers" on page 249](#)

Installing the Swarm Helix Core Server extension

Important

- If you are using the Swarm Helix Core Server extension, Swarm Helix Core Server Triggers must not be installed.
- You must be a user with *super* user permissions to install and configure Helix Core Server Extensions.

Prerequisites

To install the Swarm Helix Core Server extension you need:

A compatible version of Helix Core Server for Helix Core Server Extensions:

- **Linux:** Helix Core Server 2021.2 and later. If you are using an earlier version of Helix Core Server, you must use triggers.
- **Windows:** Helix Core Server 2021.2 and later. If you are using an earlier version of Helix Core Server, you must use triggers.

You will also need:

- Helix Core Server Extensions installed and configured on your Helix Core Server.
- Your Swarm user password
- A user with *super* permissions to install and configure Helix Core Server Extensions.

To install the Swarm Helix Core Server extension

There are two ways to install the Swarm Helix Core Server extension:

- Run the interactive post-installation configuration script. See [Install process using the interactive post-installation configuration script](#).

Note

The Swarm post-installation configuration script can be used in a few different ways. The steps below outline the most straightforward configuration using an interactive install, but you can review the options by running:

```
sudo /opt/perforce/swarm/sbin/configure-swarm.sh -h
```


- Alternatively, use the `swarm-extctl.sh` script file located in `/opt/perforce/swarm/sbin` directory. The `swarm-extctl.sh` script file specifically handles the extension upgrade. See [Install process using the swarm-extctl.sh script](#).

Install process using the interactive post-installation configuration script

1. Run the interactive post-installation configuration script to install the Swarm Helix Core Server extension:

```
sudo /opt/perforce/swarm/sbin/configure-swarm.sh
```

The configuration script displays a summary for your Swarm instance:

```
-----
configure-swarm.sh: Thu Feb 17 14:20:49 PDT 2021: commencing
configuration of Swarm
Summary of arguments passed:
Interactive?      [yes]
Force?           [no]
P4PORT           [(not specified, will prompt)]
Swarm user       [(not specified, will prompt, will suggest
swarm)]
Swarm password   [(not specified, will prompt)]
Email host       [(not specified, will prompt)]
Use Extensions?  [(not specified, will prompt)]
Swarm host       [(not specified, will prompt, will suggest
myhost)]
Swarm port       [(default (80))]
Swarm base URL   [(default (empty))]
Create Swarm user? [yes]
Super user       [(not specified, will prompt)]
Super password   [(not specified, will prompt)]
```

2. The configuration script prompts you for your Swarm configuration details.

Enter your Swarm user password when prompted, all of the other details are pre-filled with your existing Swarm configuration. Press the **[Enter]** key to accept each of them.

3. The script prompts you to use the Swarm Helix Core Server extension.

```
Do you want to use Swarm's Helix Core server extension?
Configuring Server extensions requires super user access to the
Helix Server.
```

If you install the Swarm server extension, do not install the Swarm triggers.

Server extensions are supported for:

- * Linux: Helix server 19.2 and later.
- * Windows: Helix server 21.2 and later.

Use server extensions? (y/n) [n]

Type **y** to use Helix Core Server Extensions.

4. Sign in as a user with *super* permissions when prompted.
5. The script checks that Helix Core Server Extensions can be installed.

Note

The script will:

- check your Helix Core Server supports Helix Core Server Extensions
- check if Helix Core Server Extensions are installed and configured on your Helix Core Server. If they are, Swarm does not need to do anything
- check if Swarm Helix Core Server Triggers are installed on the server

If any of these checks fail, Swarm will not install the Swarm Helix Core Server extension and will report the issues on the configuration summary screen.

6. If Swarm Triggers are installed, you are prompted to remove them.

Swarm triggers are installed on this server and must be removed before installing extensions. You can either have this script automatically remove them, or you can quit and do it manually.

Do you want to automatically remove the triggers? (y/n) [n]

- **Recommended:** Type **y** to get the script to automatically remove the Swarm Helix Core Server Triggers.

The Swarm Helix Core Server Triggers are not deleted immediately, the triggers to be removed are listed so you can review them before removing them.

- Type **n** to manually remove the Swarm Helix Core Server Triggers from the trigger spec file.

7. The script requests confirmation that it is okay to remove the triggers:

Are you sure it is okay to remove these triggers? (y/n) [n]

- **Recommended:** Type **y** to automatically remove the triggers if the list of triggers looks correct.

Note

The script will:

- save a timestamped copy of your old trigger spec to `/opt/perforce/swarm/triggers.saved.yyyymmdd-hhmmss`
- save a timestamped copy of the new trigger spec to `/opt/perforce/swarm/triggers.noswarm.yyyymmdd-hhmmss`
- remove the Swarm Helix Core Server Triggers from the trigger spec and save it

If either of the trigger specs contain sensitive information, move them to a secure location.

- Type **n** if you see something wrong with the Swarm Helix Core Server Triggers marked for removal or if you want to remove the triggers from the trigger spec manually. You are offered the option of automatically opening your trigger spec file in your default text editor or opening the file manually in a text editor. Save your changes and rerun the post-installation configuration script to complete the installation of the Swarm Helix Core Serverextension.
8. The script will now install and configure the Swarm Helix Core Server extension for you. When it has completed the configuration, the configuration summary screen is displayed, for example:

```

.....
- installing instance configuration
Extension config swarm saved
configure-swarm.sh: Thu Feb 17 14:31:36 PDT 2021: completed
configuration of Helix Swarm

::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::
:::::
::
::  Swarm is now configured and available at:
::
::      http://myhost/
::
::  Ensure that you have configured the Swarm hostname in your
::  network's DNS, or have added an IP address-to-hostname
::  mapping to your computer's hosts configuration so that you
::  can access Swarm.
::
::  You may login as the Swarm user [swarm] using the password
::  you specified.
::
::  Server side extensions are installed and configured
::  on your P4D server.
::
::  Documentation for optional post-install configuration, such
as
::  configuring Swarm to use HTTPS, operate in a sub-folder, or
on a
::  custom port, is available:
::
::
https://www.perforce.com/perforce/doc.current/manuals/swarm/setup
.post.html

```

```

::
.....
::::

```

9. The Swarm upgrade is now complete.

Install process using the `swarm-extctl.sh` script

1. Navigate to the the `swarm-extctl.sh` script file located in `/opt/perforce/swarm/sbin` directory.
2. Run the following commands as root:
 - a. Backup the current configuration. Ensure that the current configuration is correctly backed up before continuing.


```
sudo ./swarm-extctl.sh -p <masterserverIp:port> -u <username> save
```
 - b. Remove the existing extension from the master.


```
sudo ./swarm-extctl.sh -p <masterserverIp:port> -u <username> delete
```
 - c. Install the latest version of Swarm Helix Core Server extension located on the Swarm server on the master. The following command will check for all saved configurations and apply them during install.


```
sudo ./swarm-extctl.sh -p <masterserverIp:port> -u <username> install
```
3. The Swarm upgrade is now complete.

Next step

Now "Validate your upgrade" on page 253.

Updating your triggers

Important

If you are using Swarm Helix Core Server Triggers, the Swarm Helix Core Server extension must not be installed.

1. Copy the new Swarm trigger script to your Helix Core Server machine. The trigger script is `SWARM_ROOT/p4-bin/scripts/swarm-trigger.pl`, and requires installation of Perl 5.08+ (use the latest available) on the Helix Core Server machine. If Swarm is using SSL, then the triggers also require the `IO::Socket::SSL` Perl module.

Warning

Do not overwrite any existing trigger script at this time. Give the script a new name, for example: **swarm-trigger-new.pl**.

2. Configure the Swarm trigger script by creating, in the same directory on the Helix Core Server machine, **swarm-trigger.conf**. It should contain:

Note

If you already have a **swarm-trigger.conf** file, no additional configuration is required.

```
# SWARM_HOST (required)
# Hostname of your Swarm instance, with leading "http://" or
# "https://".
SWARM_HOST="http://my-swarm-host"

# SWARM_TOKEN (required)
# The token used when talking to Swarm to offer some security. To
# obtain the
# value, log in to Swarm as a super user and select 'About Swarm'
# to see the
# token value.
SWARM_TOKEN="MY-UUID-STYLE-TOKEN"

# ADMIN_USER (optional) Do not use if the Workflow feature is
# enabled (default)
# For enforcing reviewed changes, optionally specify the normal
# Perforce user
# with admin privileges (to read keys); if not set, will use
# whatever Perforce
# user is set in environment.
ADMIN_USER=

# ADMIN_TICKET_FILE (optional) Do not use if the Workflow feature
# is enabled (default)
# For enforcing reviewed changes, optionally specify the location
# of the
```

```
# p4tickets file if different from the default
($HOME/.p4tickets).
# Ensure this user is a member of a group with an 'unlimited' or
very long
# timeout; then, manually login as this user from the Perforce
server machine to
# set the ticket.
ADMIN_TICKET_FILE=

# VERIFY_SSL (optional)
# If HTTPS is being used on the Swarm web server, then this
controls whether
# the SSL certificate is validated or not. By default this is set
to 1, which
# means any SSL certificates must be valid. If the web server is
using a self
# signed certificate, then this must be set to 0.
# set the ticket.
VERIFY_SSL=1
```

Fill in the required **SWARM_HOST** and **SWARM_TOKEN** variables with the configuration from any previous Swarm trigger script, typically **swarm-trigger.pl**.

Tip

The **ADMIN_USER** and **ADMIN_TICKET** variables were used by the '**enforce triggers**' in Swarm 2019.1 and earlier. They can be removed unless you are explicitly [disabling workflow](#) and using the deprecated '**enforce triggers**'.

Note

Swarm 2015.4 and earlier: Swarm trigger script files were available as shell scripts in these earlier Swarm versions, typically **swarm-trigger.sh**.

Swarm must now use a Perl trigger script file, typically **swarm-trigger.pl**.

3. **On Linux:** ensure that the script is executable:

```
sudo chmod +x swarm-trigger-new.pl
```

4. Rename the new trigger script:

On Linux:

```
mv swarm-trigger-new.pl swarm-trigger.pl
```

On Windows:

```
ren swarm-trigger-new.pl swarm-trigger.pl
```

5. Update the triggers in your Helix Core Server.

Warning

- The **swarm.shelvedel shelve-delete** trigger line was added to Swarm in version 2018.1 and updated in version 2020.1.
 - **Upgrading from Swarm 2017.4 and earlier:** add the **swarm.shelvedel shelve-delete** trigger line to the Helix Core Server trigger table if it is not already present, see "[Update the Helix Core Server triggers table to run the trigger script.](#)" on page 211
 - **Upgrading from Swarm 2018.x and 2019.x:** replace the existing **swarm.shelvedel shelve-delete** trigger line in the Helix Core Server trigger table with the one supplied in the Swarm version you are upgrading to.
- **Workflow feature:**

The [Workflow feature](#) is enabled by default in Swarm 2019.2 and later. The trigger lines required when workflow is enabled are different to those required when workflow is disabled:

 - **Workflow feature [enabled](#) (default):**
 - Comment out the **swarm.enforce.1**, **swarm.enforce.2**, **swarm.strict.1**, and **swarm.strict.2** trigger lines in the Helix Core Server trigger table if they are present, see "[Update the Helix Core Server triggers table to run the trigger script.](#)" on page 211
 - Add the **swarm.enforce change-submit**, **swarm.strict change-content**, and **swarm.shelvesub shelve-submit** trigger lines to the Helix Core Server trigger table if they are not already present, see "[Update the Helix Core Server triggers table to run the trigger script.](#)" on page 211
 - **Workflow feature [disabled](#):**

Comment out the **swarm.enforce change-submit**, **swarm.strict change-content**, and **swarm.shelvesub shelve-submit** trigger lines in the Helix Core Server trigger table if they are present, see "[Update the Helix Core Server triggers table to run the trigger script.](#)" on page 211

- a. Run the Swarm trigger script to capture (using **Ctrl+C** on Windows and Linux) the trigger lines that should be included in the Perforce trigger table:

On Linux:

```
./swarm-trigger.pl -o
```


On Windows:

```
path/to/perl swarm-trigger.pl -o
```

- b. As a Perforce user with *super* privileges, update the Perforce trigger table by running **p4 triggers** command and replacing any **swarm.*** lines with the previously captured trigger line output (using **Ctrl+V** on Windows and Linux).

Important

If you previously customized the Swarm trigger lines, perhaps to apply various ["Trigger options" on page 735](#), be sure to repeat those customizations within the updated trigger lines.

Next step

Now ["Validate your upgrade" below](#).

Validate your upgrade

Tip

When Swarm starts it verifies the Redis cache, during this time you cannot log in to Swarm. The time taken to verify the Redis cache depends on the number of users, groups, and projects Swarm has. Start-up time can be improved by persisting the memory cache. You can persist the memory cache by disabling background saves and enabling append saves in the **redis-server.conf** file, see ["Redis server configuration file" on page 676](#).

Check that your upgraded Swarm instance is working correctly by doing the following:

1. Create a new changelist that:
 - a. Contains at least one modified file
 - b. Contains the **#review** keyword in the changelist description
2. Right click on the new changelist in P4V and click **Shelve Files...**

Important

Do not select **Request New Swarm Review...** because this method uses the API and will not fully test the Swarm Helix Core Server extension.

This is also true if you are using Swarm Helix Core Server Triggers instead of the Swarm Helix Core Server extension.

3. Check that a new Swarm review is created for the changelist.
 - If a Swarm review is created, the Swarm Helix Core Server extension is working. If you are using Swarm Helix Core Server Triggers instead of the Swarm Helix Core Server

extension and the review is created, the triggers are working.

- If a Swarm review is not created, see "[Review not created](#)" on page 240.

Swarm index upgrade

If you are upgrading from Swarm 2017.2 or earlier you should run the index upgrade, this ensures that the review activity history is displayed in the correct order on the [Dashboard](#), and [Reviews list](#) pages.

Note

If you are upgrading from Swarm version 2017.3 or later, the index upgrade step is not required.

The index upgrade process can be configured to suit your Swarm system specifications. See [Upgrade index](#) for details.

Run the upgrade as an *Admin* user by visiting the following URL:

```
http://SWARM-HOST/upgrade
```

Note

After the Swarm upgrade, on the first visit to some Swarm pages, users might see a message to perform a browser hard refresh. This happens because we are updating the UI and content of some Swarm pages, so the user's page cache is no longer valid and requires a hard refresh to load the upgraded page. For example, for Chrome on Windows/Linux **[CTRL]+[F5]**.

All done!

Upgrading a CentOS/RHEL package installation

The section describes how to upgrade a Swarm CentOS/RHEL package installation to a newer release.

Important

- Swarm runtime dependencies change between releases, you must check that your system satisfies the Swarm runtime dependencies before starting the upgrade, see "[Runtime dependencies](#)" on page 98.
- Review the PHP requirements before you upgrade Swarm, see "[PHP](#)" on page 102.
- Review the Helix Core Server requirements before you upgrade Swarm, see "[Helix Core Server requirements](#)" on page 106.
- Helix Core Server 2020.1 and later, permissions have changed for viewing and [editing stream spec files](#) in Swarm. To view and edit stream spec files in Swarm, the Swarm user must have *admin* permissions for the entire depot `// . . .`.

Note

- If you are upgrading from Swarm 2017.2 or earlier, run the Swarm index upgrade after you have validated your upgrade. This is the last step of the upgrade and ensures that the review activity history is displayed in the correct order on the [Dashboard](#), and [Reviews list](#) pages.
- If you are upgrading from Swarm 2020.2 or earlier and have `userids` that contain the forward slash (/) character, add `AllowEncodedSlashes NoDecode` to the `VirtualHost` block of your `/etc/apache2/sites-enabled/perforce-swarm-site.conf` file. For more information about the `VirtualHost` block, see "[Apache configuration](#)" on page 186.

Upgrade Swarm

Important

The following notes are applicable for CentOS/RHEL7, RHEL 8, and RHEL 9:

- As part of the PHP upgrade process your Apache 2.2 server will be stopped and disabled, if you are currently using the Apache 2.2 server for any other applications they will stop working. You will either need to upgrade these applications to work with PHP 7 or PHP 8 and Apache 2.4 or move them to another server.
- **Swarm 2019.1 to 2020.1:** these versions of Swarm used PHP packages from the SCL repositories for CentOS/RHEL 7. This was to provide access to more recent versions of PHP than are available as standard on CentOS/RHEL. This required the use of the `httpd24-httpd` package for Apache. These were all installed into `/opt/rh`

Swarm 2020.2 and later: these versions of Swarm uses the Remi repository for CentOS/RHEL 7, RHEL 8, and RHEL 9. This provides PHP 7.4 installed in the standard file system structure. This means that the old `httpd24-httpd` version of Apache is no longer needed, and the standard system version of Apache is being used again.

The SCL Apache site configuration file was installed at this location for Swarm 2019.1 to 2020.1:

```
/opt/rh/httpd24/root/etc/httpd/conf.d/perforce-swarm-site.conf
```

If this exists when Swarm is upgraded to 2020.2 and later, this file is copied to `/etc/httpd/conf.d/perforce-swarm-site.conf` if there is no file at the destination. It is also re-written to change references from `/var/log/httpd24` to `/var/log/httpd`

If a site configuration file for Swarm already exists in `/etc/httpd`, the copy and re-write is not performed.

After upgrade, `httpd24-httpd` is disabled.

- To avoid seeing the Apache HTTP server Linux test page when you start the Apache server, comment out the content of the `welcome.conf` file located in the `/etc/httpd/conf.d/` directory.
- To avoid loading the Apache HTTP server example configuration instead of the Swarm configuration when the Apache server starts, rename the `autoindex.conf` file located in the `/etc/httpd/conf.d/` directory to `z-autoindex.conf` or similar. This is required because Apache runs the first conf file it finds in the `/etc/httpd/conf.d/` directory (alphabetical order) and that must be the `perforce-swarm-site.conf` file.

Tip

From Swarm 2021.1, the Swarm package upgrade installs logrotate to manage your Swarm log rotation. If the package upgrade finds an existing custom logrotate configuration file for Swarm, the upgrade will notify you and give you details on how to disable the new logrotate configuration.

For information about the logrotate configuration, see ["Logrotate" on page 633](#).

The following process attempts to minimize downtime, but a short period of downtime for Swarm users is unavoidable. There should be no downtime for your Helix Core Server. After a successful upgrade, all Swarm users are logged out.

If you are using Swarm in a production environment, we encourage you to test this upgrade process in a non-production environment first.

1. Install the main Swarm package on the server to host Swarm.

Follow the instructions for your OS distribution, see [CentOS/REHL 7](#), [RHEL 8](#), or [RHEL 9](#) below:

- **CentOS/RHEL 7** (run these commands as root):

The full PHP and Apache upgrade process described below is only required the first time you upgrade to PHP 7.x. For future Swarm upgrades just run the Swarm upgrade steps, the ImageMagick steps if you want to install ImageMagick, and check that SELinux is working correctly.

The upgrade process will deploy `epel-release-latest-7.noarch.rpm` to give Swarm access to the EPEL repository, deploy `remi-release-7.rpm` to give Swarm access to PHP 7.x and Apache 2.4, upgrade Swarm, stop and disable the Apache 2.2 server, copy and edit some Swarm files so they work with Apache 2.4, and finally start and enable the Apache 2.4 server.

- a. Deploy the `epel-release-latest-7.noarch.rpm` repository configuration package to give Swarm access to EPEL packages:

```
yum install https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm
```

- b. Deploy the Remi repository configuration package to give Swarm access to PHP 7.x (only required the first time you upgrade to PHP 7.x):

```
yum install https://rpms.remirepo.net/enterprise/remi-release-7.rpm
```

Tip

If you don't deploy the Remi repository, you will see dependency errors when you do the next steps.

- c. Install the `yum-utils` package to give access to the `yum-config-manager` command:

```
yum install yum-utils
```

- d. **RHEL only:** Enable the optional channel for some dependencies:

```
subscription-manager repos --enable=rhel-7-server-optional-rpms
```

- e. Install the **Default/Single version** of PHP:

- i. Disable `remi-php*`:

```
yum-config-manager --disable 'remi-php*'
```

- ii. Enable PHP 7.4:

```
yum-config-manager --enable remi-php74
```

- iii. Run an upgrade for PHP, this will also upgrade the Swarm packages:

```
yum update
```

- f. If you didn't run `yum update` in the previous step, run the Swarm package upgrade now:

```
yum install helix-swarm helix-swarm-triggers helix-swarm-optional
```

Important

If you are upgrading from Swarm 2019.3 to Swarm 2021.1 or later, remove the Swarm PHP 7.1 files:

```
yum remove $(yum list installed | grep php71 | awk '{print $1}')
```

- g. If you are upgrading from Apache 2.2, disable your Apache 2.2 HTTP server so that it does not automatically start:

```
systemctl disable httpd
```

- h. If you are upgrading from Apache 2.2, stop your Apache 2.2 HTTP server:

```
systemctl stop httpd
```

- i. If you have any special `php.ini` configuration options set, copy the `php.ini` file to the `/etc/php.d/php.ini/` directory, for example:

```
cp /etc/opt/rh/rh-php72/php.d/php.ini /etc/php.d/php.ini
```

- j. Copy the `perforce-swarm-site.conf` file to the `/etc/httpd/conf.d/` directory if it is not already in there, for example:

```
cp /opt/rh/httpd24/root/etc/httpd/conf.d/perforce-swarm-site.conf /etc/httpd/conf.d/perforce-swarm-site.conf
```

- k. Replace the `/log/httpd24` filepath with `/log/httpd/` in the `perforce-swarm-site.conf` file using the `sed` command:

```
sed -i "s#/log/httpd24/#/log/httpd/#g" /etc/httpd/conf.d/perforce-swarm-site.conf
```

- l. Enable your Apache 2.4 HTTP server so that it automatically starts:

```
systemctl enable httpd
```

- m. Start the Apache 2.4 HTTP server:

```
systemctl start httpd
```

- n. **Optional, ImageMagick:** when ImageMagick is enabled, Swarm can display the following image formats: BMP, EPS, PSD, TGA, TIFF.

- i. Install ImageMagick:

```
yum install ImageMagick
```

- ii. Restart the web server so that ImageMagick is picked up.

```
systemctl restart httpd
```

- o. If you are upgrading from Swarm 2020.1, restart your Redis server:

```
systemctl restart redis-server-swarm
```

- p. Check that SELinux is working correctly for your system. For instructions on configuring SELinux on CentOS/RHEL, see ["SELinux configuration" on page 149](#).

■ **RHEL 8** (run these commands as root):

The full PHP and Apache upgrade process described below is only required the first time you upgrade to PHP 7.x. For future Swarm upgrades just run the Swarm upgrade steps, the ImageMagick steps if you want to install ImageMagick, and check that SELinux is working correctly.

The upgrade process will deploy `epel-release-latest-8.noarch.rpm` to give Swarm access to the EPEL repository, deploy `remi-release-8.rpm` to give Swarm access to PHP 7.x and Apache 2.4, upgrade Swarm, stop and disable the Apache 2.2 server, copy and edit some Swarm files so they work with Apache 2.4, and finally start and enable the Apache 2.4 server.

- a. Deploy the `epel-release-latest-8.noarch.rpm` repository configuration package to give Swarm access to EPEL packages:

```
dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-8.noarch.rpm
```

- b. Deploy the Remi repository configuration package to give Swarm access to PHP 7.x (only required the first time you upgrade to PHP 7.x):

```
dnf install https://rpms.remirepo.net/enterprise/remi-release-8.rpm
```

Tip

If you don't deploy the Remi repository, you will see dependency errors when you do the next steps.

- c. Install the `yum-utils` package to give access to the `yum-config-manager` command:

```
dnf install yum-utils
```

- d. Enable the optional channel for some dependencies:


```
subscription-manager repos --enable=rhel-8-server-optional-rpms
```

- e. Install the **Default/Single version** of PHP:

- i. Enable the module stream for PHP 7.4:

```
dnf module reset php
```

- ii. Install PHP 7.4:

```
dnf module install php:remi-7.4
```

- iii. Run an upgrade for PHP, this will also upgrade the Swarm packages:

```
dnf update
```

- f. If you didn't run `dnf update` in the previous step, run the Swarm package upgrade now:

```
yum install helix-swarm helix-swarm-triggers helix-swarm-optional
```

Important

If you are upgrading from Swarm 2019.3 to Swarm 2021.1 or later, remove the Swarm PHP 7.1 files:

```
yum remove $(yum list installed | grep php71 | awk '{print $1}')
```

- g. If you are upgrading from Apache 2.2, disable your Apache 2.2 HTTP server so that it does not automatically start:

```
systemctl disable httpd
```

- h. If you are upgrading from Apache 2.2, stop your Apache 2.2 HTTP server:

```
systemctl stop httpd
```

- i. If you have any special `php.ini` configuration options set, copy the `php.ini` file to the `/etc/php.d/php.ini/` directory, for example:

```
cp /etc/opt/rh/rh-php72/php.d/php.ini /etc/php.d/php.ini
```

- j. Copy the `perforce-swarm-site.conf` file to the `/etc/httpd/conf.d/` directory if it is not already in there, for example:

```
cp /opt/rh/httpd24/root/etc/httpd/conf.d/perforce-swarm-site.conf /etc/httpd/conf.d/perforce-swarm-site.conf
```

- k. Replace the `/log/httpd24` filepath with `/log/httpd/` in the `perforce-swarm-site.conf` file using the `sed` command:

```
sed -i "s#/log/httpd24/#!/log/httpd/#g"
/etc/httpd/conf.d/perforce-swarm-site.conf
```

- l. Enable your Apache 2.4 HTTP server so that it automatically starts:

```
systemctl enable httpd
```

- m. Start the Apache 2.4 HTTP server:

```
systemctl start httpd
```

- n. Enable your Apache 2.4 HTTP server so that it automatically starts:

```
systemctl enable httpd
```

- o. Start the Apache 2.4 HTTP server:

```
systemctl start httpd
```

- p. **Optional, ImageMagick:** when ImageMagick is enabled, Swarm can display the following image formats: BMP, EPS, PSD, TGA, TIFF.

- i. Install ImageMagick:

```
yum install ImageMagick
```

- ii. Restart the web server so that ImageMagick is picked up.

```
systemctl restart httpd
```

- q. If you are upgrading from Swarm 2020.1, restart your Redis server:

```
systemctl restart redis-server-swarm
```

- r. Check that SELinux is working correctly for your system. For instructions on configuring SELinux on CentOS/RHEL, see ["SELinux configuration" on page 149](#).

■ RHEL 9 (run these commands as root):

- a. Run an upgrade for PHP, this will also upgrade the Swarm packages:

```
dnf update
```

- b. If you didn't run `dnf update` in the previous step, run the Swarm package upgrade now:

```
yum install helix-swarm helix-swarm-triggers helix-swarm-optional
```

Important

If you are upgrading from Swarm 2019.3 to Swarm 2021.1 or later, remove the Swarm PHP 7.1 files:

```
yum remove $(yum list installed | grep php71 | awk '{print $1}')
```

- c. **Optional, ImageMagick:** when ImageMagick is enabled, Swarm can display the following image formats: BMP, EPS, PSD, TGA, TIFF.
 - i. Install ImageMagick:


```
yum install ImageMagick
```
 - ii. Restart the web server so that ImageMagick is picked up.


```
systemctl restart httpd
```
 - d. Check that SELinux is working correctly for your system. For instructions on configuring SELinux on CentOS/RHEL, see ["SELinux configuration" on page 149](#).
2. Swarm generally has several major updates each year, and may occasionally have a patch update between major updates. To determine whether a Swarm update is available.
- Run the following command as root:

```
yum list updates | grep swarm
```

Tip

Swarm uses a Redis server to manage its caches. This is installed and configured on the Swarm machine during the upgrade. If you prefer to use your own Redis server, see ["Use your own Redis server" on page 235](#).

Configuring Helix Core Server event notification

Swarm needs to know about a number of Helix Core Server events to operate correctly, this can be done by using Helix Core Server Extensions or Helix Core Server Triggers. Swarm installs include the Swarm Helix Core Server extension file and trigger scripts required for Swarm to get the events it needs from your Helix Core Server.

You must install the Swarm Helix Core Server extension or update your Triggers to complete the Swarm upgrade.

Do one of the following so that Swarm is notified about events on the Helix Core Server:

- Install and configure the Swarm Helix Core Server extension, see ["Installing the Swarm Helix Core Server extension" below](#)
- Update your Swarm Triggers, see ["Updating your triggers" on page 269](#)

Installing the Swarm Helix Core Server extension

Important

- If you are using the Swarm Helix Core Server extension, Swarm Helix Core Server Triggers must not be installed.

- You must be a user with *super* user permissions to install and configure Helix Core Server Extensions.

Prerequisites

To install the Swarm Helix Core Server extension you need:

A compatible version of Helix Core Server for Helix Core Server Extensions:

- **Linux:** Helix Core Server 2021.2 and later. If you are using an earlier version of Helix Core Server, you must use triggers.
- **Windows:** Helix Core Server 2021.2 and later. If you are using an earlier version of Helix Core Server, you must use triggers.

You will also need:

- Helix Core Server Extensions installed and configured on your Helix Core Server.
- Your Swarm user password
- A user with *super* permissions to install and configure Helix Core Server Extensions.

To install the Swarm Helix Core Server extension

There are two ways to install the Swarm Helix Core Server extension:

- Run the interactive post-installation configuration script. See [Install process using the interactive post-installation configuration script](#).

Note

The Swarm post-installation configuration script can be used in a few different ways. The steps below outline the most straightforward configuration using an interactive install, but you can review the options by running:

```
sudo /opt/perforce/swarm/sbin/configure-swarm.sh -h
```

- Alternatively, use the `swarm-extctl.sh` script file located in `/opt/perforce/swarm/sbin` directory. The `swarm-extctl.sh` script file specifically handles the extension upgrade. See [Install process using the swarm-extctl.sh script](#).

Install process using the interactive post-installation configuration script

1. Run the interactive post-installation configuration script to install the Swarm Helix Core Server extension:

```
sudo /opt/perforce/swarm/sbin/configure-swarm.sh
```

The configuration script displays a summary for your Swarm instance:

```

-----
configure-swarm.sh: Thu Feb 17 14:20:49 PDT 2021: commencing
configuration of Swarm
Summary of arguments passed:
Interactive?      [yes]
Force?           [no]
P4PORT           [(not specified, will prompt)]
Swarm user       [(not specified, will prompt, will suggest
swarm)]
Swarm password   [(not specified, will prompt)]
Email host       [(not specified, will prompt)]
Use Extensions?  [(not specified, will prompt)]
Swarm host       [(not specified, will prompt, will suggest
myhost)]
Swarm port       [(default (80))]
Swarm base URL   [(default (empty))]
Create Swarm user? [yes]
Super user       [(not specified, will prompt)]
Super password   [(not specified, will prompt)]

```

2. The configuration script prompts you for your Swarm configuration details.

Enter your Swarm user password when prompted, all of the other details are pre-filled with your existing Swarm configuration. Press the **[Enter]** key to accept each of them.

3. The script prompts you to use the Swarm Helix Core Server extension.

```

Do you want to use Swarm's Helix Core server extension?
Configuring Server extensions requires super user access to the
Helix Server.
If you install the Swarm server extension, do not install the
Swarm triggers.
Server extensions are supported for:
* Linux: Helix server 19.2 and later.
* Windows: Helix server 21.2 and later.

Use server extensions? (y/n) [n]

```

Type **y** to use Helix Core Server Extensions.

4. Sign in as a user with *super* permissions when prompted.
5. The script checks that Helix Core Server Extensions can be installed.

Note

The script will:

- check your Helix Core Server supports Helix Core Server Extensions
- check if Helix Core Server Extensions are installed and configured on your Helix Core Server. If they are, Swarm does not need to do anything
- check if Swarm Helix Core Server Triggers are installed on the server

If any of these checks fail, Swarm will not install the Swarm Helix Core Server extension and will report the issues on the configuration summary screen.

6. If Swarm Triggers are installed, you are prompted to remove them.

```
Swarm triggers are installed on this server and must be removed
before installing extensions. You can either have this script
automatically remove them, or you can quit and do it manually.
Do you want to automatically remove the triggers? (y/n) [n]
```

- **Recommended:** Type **y** to get the script to automatically remove the Swarm Helix Core Server Triggers.

The Swarm Helix Core Server Triggers are not deleted immediately, the triggers to be removed are listed so you can review them before removing them.

- Type **n** to manually remove the Swarm Helix Core Server Triggers from the trigger spec file.

7. The script requests confirmation that it is okay to remove the triggers:

```
Are you sure it is okay to remove these triggers? (y/n) [n]
```

- **Recommended:** Type **y** to automatically remove the triggers if the list of triggers looks correct.

Note

The script will:

- save a timestamped copy of your old trigger spec to
`/opt/perforce/swarm/triggers.saved.yyyymmdd-hhmmss`
- save a timestamped copy of the new trigger spec to
`/opt/perforce/swarm/triggers.noswarm.yyyymmdd-hhmmss`
- remove the Swarm Helix Core Server Triggers from the trigger spec and save it

If either of the trigger specs contain sensitive information, move them to a secure location.

- Type **n** if you see something wrong with the Swarm Helix Core Server Triggers marked for removal or if you want to remove the triggers from the trigger spec manually. You are offered the option of automatically opening your trigger spec file in your default text editor or opening the file manually in a text editor. Save your changes and rerun the post-installation configuration script to complete the installation of the Swarm Helix Core Serverextension.
8. The script will now install and configure the Swarm Helix Core Server extension for you. When it has completed the configuration, the configuration summary screen is displayed, for example:

```

.....
- installing instance configuration
Extension config swarm saved
configure-swarm.sh: Thu Feb 17 14:31:36 PDT 2021: completed
configuration of Helix Swarm

::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::
:::::
::
::  Swarm is now configured and available at:
::
::      http://myhost/
::
::  Ensure that you have configured the Swarm hostname in your
::  network's DNS, or have added an IP address-to-hostname
::  mapping to your computer's hosts configuration so that you
::  can access Swarm.
::
::  You may login as the Swarm user [swarm] using the password
::  you specified.
::
::  Server side extensions are installed and configured
::  on your P4D server.
::
::  Documentation for optional post-install configuration, such
as
::  configuring Swarm to use HTTPS, operate in a sub-folder, or
on a
::  custom port, is available:
::
::
https://www.perforce.com/perforce/doc.current/manuals/swarm/setup
.post.html

```



```

::
.....
::::

```

9. The Swarm upgrade is now complete.

Install process using the `swarm-extctl.sh` script

1. Navigate to the the `swarm-extctl.sh` script file located in `/opt/perforce/swarm/sbin` directory.
2. Run the following commands as root:
 - a. Backup the current configuration. Ensure that the current configuration is correctly backed up before continuing.


```
sudo ./swarm-extctl.sh -p <masterserverIp:port> -u <username> save
```
 - b. Remove the existing extension from the master.


```
sudo ./swarm-extctl.sh -p <masterserverIp:port> -u <username> delete
```
 - c. Install the latest version of Swarm Helix Core Server extension located on the Swarm server on the master. The following command will check for all saved configurations and apply them during install.


```
sudo ./swarm-extctl.sh -p <masterserverIp:port> -u <username> install
```
3. The Swarm upgrade is now complete.

Next step

Now "Validate your upgrade" on page 273.

Updating your triggers

Important

If you are using Swarm Helix Core Server Triggers, the Swarm Helix Core Server extension must not be installed.

1. Copy the new Swarm trigger script to your Helix Core Server machine. The trigger script is `SWARM_ROOT/p4-bin/scripts/swarm-trigger.pl`, and requires installation of Perl 5.08+ (use the latest available) on the Helix Core Server machine. If Swarm is using SSL, then the triggers also require the `IO::Socket::SSL` Perl module.

Warning

Do not overwrite any existing trigger script at this time. Give the script a new name, for example: **swarm-trigger-new.pl**.

2. Configure the Swarm trigger script by creating, in the same directory on the Helix Core Server machine, **swarm-trigger.conf**. It should contain:

Note

If you already have a **swarm-trigger.conf** file, no additional configuration is required.

```
# SWARM_HOST (required)
# Hostname of your Swarm instance, with leading "http://" or
# "https://".
SWARM_HOST="http://my-swarm-host"

# SWARM_TOKEN (required)
# The token used when talking to Swarm to offer some security. To
# obtain the
# value, log in to Swarm as a super user and select 'About Swarm'
# to see the
# token value.
SWARM_TOKEN="MY-UUID-STYLE-TOKEN"

# ADMIN_USER (optional) Do not use if the Workflow feature is
# enabled (default)
# For enforcing reviewed changes, optionally specify the normal
# Perforce user
# with admin privileges (to read keys); if not set, will use
# whatever Perforce
# user is set in environment.
ADMIN_USER=

# ADMIN_TICKET_FILE (optional) Do not use if the Workflow feature
# is enabled (default)
# For enforcing reviewed changes, optionally specify the location
# of the
```

```
# p4tickets file if different from the default
($HOME/.p4tickets).
# Ensure this user is a member of a group with an 'unlimited' or
very long
# timeout; then, manually login as this user from the Perforce
server machine to
# set the ticket.
ADMIN_TICKET_FILE=

# VERIFY_SSL (optional)
# If HTTPS is being used on the Swarm web server, then this
controls whether
# the SSL certificate is validated or not. By default this is set
to 1, which
# means any SSL certificates must be valid. If the web server is
using a self
# signed certificate, then this must be set to 0.
# set the ticket.
VERIFY_SSL=1
```

Fill in the required **SWARM_HOST** and **SWARM_TOKEN** variables with the configuration from any previous Swarm trigger script, typically **swarm-trigger.pl**.

Tip

The **ADMIN_USER** and **ADMIN_TICKET** variables were used by the **'enforce triggers'** in Swarm 2019.1 and earlier. They can be removed unless you are explicitly [disabling workflow](#) and using the deprecated **'enforce triggers'**.

Note

Swarm 2015.4 and earlier: Swarm trigger script files were available as shell scripts in these earlier Swarm versions, typically **swarm-trigger.sh**.

Swarm must now use a Perl trigger script file, typically **swarm-trigger.pl**.

3. **On Linux:** ensure that the script is executable:

```
sudo chmod +x swarm-trigger-new.pl
```

4. Rename the new trigger script:

On Linux:

```
mv swarm-trigger-new.pl swarm-trigger.pl
```

On Windows:

```
ren swarm-trigger-new.pl swarm-trigger.pl
```

5. Update the triggers in your Helix Core Server.

Warning

- The **swarm.shelvedel shelve-delete** trigger line was added to Swarm in version 2018.1 and updated in version 2020.1.

- **Upgrading from Swarm 2017.4 and earlier:** add the **swarm.shelvedel shelve-delete** trigger line to the Helix Core Server trigger table if it is not already present, see "[Update the Helix Core Server triggers table to run the trigger script.](#)" on page 211
- **Upgrading from Swarm 2018.x and 2019.x:** replace the existing **swarm.shelvedel shelve-delete** trigger line in the Helix Core Server trigger table with the one supplied in the Swarm version you are upgrading to.

■ Workflow feature:

The [Workflow feature](#) is enabled by default in Swarm 2019.2 and later. The trigger lines required when workflow is enabled are different to those required when workflow is disabled:

- **Workflow feature [enabled](#) (default):**
 - Comment out the **swarm.enforce.1**, **swarm.enforce.2**, **swarm.strict.1**, and **swarm.strict.2** trigger lines in the Helix Core Server trigger table if they are present, see "[Update the Helix Core Server triggers table to run the trigger script.](#)" on page 211
 - Add the **swarm.enforce change-submit**, **swarm.strict change-content**, and **swarm.shelvesub shelve-submit** trigger lines to the Helix Core Server trigger table if they are not already present, see "[Update the Helix Core Server triggers table to run the trigger script.](#)" on page 211
- **Workflow feature [disabled](#):**

Comment out the **swarm.enforce change-submit**, **swarm.strict change-content**, and **swarm.shelvesub shelve-submit** trigger lines in the Helix Core Server trigger table if they are present, see "[Update the Helix Core Server triggers table to run the trigger script.](#)" on page 211

- a. Run the Swarm trigger script to capture (using **Ctrl+C** on Windows and Linux) the trigger lines that should be included in the Perforce trigger table:

On Linux:

```
./swarm-trigger.pl -o
```

On Windows:

```
path/to/perl swarm-trigger.pl -o
```

- b. As a Perforce user with *super* privileges, update the Perforce trigger table by running **p4 triggers** command and replacing any **swarm.*** lines with the previously captured trigger line output (using **Ctrl+V** on Windows and Linux).

Important

If you previously customized the Swarm trigger lines, perhaps to apply various ["Trigger options" on page 735](#), be sure to repeat those customizations within the updated trigger lines.

Next step

Now ["Validate your upgrade" below](#).

Validate your upgrade

Tip

When Swarm starts it verifies the Redis cache, during this time you cannot log in to Swarm. The time taken to verify the Redis cache depends on the number of users, groups, and projects Swarm has. Start-up time can be improved by persisting the memory cache. You can persist the memory cache by disabling background saves and enabling append saves in the **redis-server.conf** file, see ["Redis server configuration file" on page 676](#).

Check that your upgraded Swarm instance is working correctly by doing the following:

1. Create a new changelist that:
 - a. Contains at least one modified file
 - b. Contains the **#review** keyword in the changelist description
2. Right click on the new changelist in P4V and click **Shelve Files...**

Important

Do not select **Request New Swarm Review...** because this method uses the API and will not fully test the Swarm Helix Core Server extension.

This is also true if you are using Swarm Helix Core Server Triggers instead of the Swarm Helix Core Server extension.

3. Check that a new Swarm review is created for the changelist.
 - If a Swarm review is created, the Swarm Helix Core Server extension is working. If you are using Swarm Helix Core Server Triggers instead of the Swarm Helix Core Server

extension and the review is created, the triggers are working.

- If a Swarm review is not created, see "[Review not created](#)" on page 240.

Swarm index upgrade

If you are upgrading from Swarm 2017.2 or earlier you should run the index upgrade, this ensures that the review activity history is displayed in the correct order on the [Dashboard](#), and [Reviews list](#) pages.

Note

If you are upgrading from Swarm version 2017.3 or later, the index upgrade step is not required.

The index upgrade process can be configured to suit your Swarm system specifications. See [Upgrade index](#) for details.

Run the upgrade as an *Admin* user by visiting the following URL:

```
http://SWARM-HOST/upgrade
```

Note

After the Swarm upgrade, on the first visit to some Swarm pages, users might see a message to perform a browser hard refresh. This happens because we are updating the UI and content of some Swarm pages, so the user's page cache is no longer valid and requires a hard refresh to load the upgraded page. For example, for Chrome on Windows/Linux **[CTRL]+[F5]**.

All done!

Upgrading a Docker container

This section describes how to upgrade a Docker container that is running Swarm to a newer release.

The following process attempts to minimize downtime, but a short period of downtime for Swarm users is unavoidable. There should be no downtime for your Helix Core Server. After a successful upgrade, all Swarm users are logged out.

If you are using Swarm in a production environment, we encourage you to test this upgrade process in a non-production environment first.

To upgrade a Docker container, run the following commands in your Docker setup directory:

1. Pull down the latest version of the Helix Swarm image. In the below code example, you can change the *latest* tag to the current version of your Helix Swarm installation, for example, 2022.2.

```
docker pull perforce/helix-swarm:latest
```

2. View what is currently running in your Docker container using **docker ps** command.

```

docker ps
CONTAINER ID   IMAGE                                COMMAND
CREATED        STATUS        PORTS
NAMES
0b3baaf10387   perforce/helix-swarm               "/bin/sh -c /opt/per..." 22
seconds ago    Up 20 seconds  0.0.0.0:80->80/tcp,
:::80->80/tcp
helix-swarm
f2030d449679   redis                               "docker-entrypoint.s..." 38
seconds ago    Up 37 seconds  6379/tcp helix-redis

```

3. Stop and delete the running Helix Swarm image.

```

docker stop helix-swarm
docker rm helix-swarm

```

4. Start a new Helix Swarm image using a **docker run** command. For more information on how to run a Helix Swarm image on a Docker container, see ["Run Swarm using a Docker container" on page 169](#).

Upgrading a tarball installation

The section describes how to upgrade a Swarm tarball installation to a newer release.

The following process attempts to minimize downtime, but a short period of downtime for Swarm users is unavoidable. There should be no downtime for your Helix Core Server. After a successful upgrade, all Swarm users are logged out.

If you are using Swarm in a production environment, we encourage you to test this upgrade process in a non-production environment first.

Warning

- Swarm runtime dependencies change between releases, you must check that your system satisfies the Swarm runtime dependencies before starting the upgrade, see ["Runtime dependencies" on page 98](#).
- P4PHP should be upgraded to the version included in the new Swarm release.
 - If you have already configured PHP to use the Swarm-provided P4PHP ([as recommended](#)), this happens automatically.

- If you have manually installed P4PHP in some other fashion, configure Swarm to use the version of P4PHP included in the new Swarm tarball before you perform any of the upgrade steps below. See ["PHP configuration" on page 190](#) for details.

Swarm package and tarball installations: two versions of P4PHP are supplied for each PHP 7 and PHP 8 version supported by Swarm. They are located in the `p4-bin/bin.linux26x86_64` directory.

- `perforce-php7x.so` and `perforce-php8x.so` compatible with systems using SSL 1.0.2
- `perforce-php7x-ssl1.1.1.so` and `perforce-php8x-ssl1.1.1.so` compatible with systems using SSL 1.1.1 (by default Ubuntu 20.04 uses SSL 1.1.1)
- `perforce-php8x-ssl3.so` compatible with systems using SSL 3.0.0 (by default Ubuntu 22.04 uses SSL 3.0.0)

Where `x` is the version of PHP 7 or PHP 8.

If the `perforce.ini` file is not pointing at the correct version of P4PHP and you connect to an SSL enabled Helix Core Server:

- The Swarm web-page will not load and you might see a **Connection Reset** error.
 - There might be an **undefined symbol: SSLey** message in the Apache error log
- Review the PHP requirements before you upgrade Swarm, see ["PHP" on page 102](#).
 - Review the Helix Core Server requirements before you upgrade Swarm, see ["Helix Core Server requirements" on page 106](#).
 - Helix Core Server 2020.1 and later, permissions have changed for viewing and [editing stream spec files](#) in Swarm. To view and edit stream spec files in Swarm, the Swarm user must have *admin* permissions for the entire depot `// . . .`

Note

- If you are upgrading from Swarm 2017.2 or earlier, run the Swarm index upgrade after you have validated your upgrade. This is the last step of the upgrade and ensures that the review activity history is displayed in the correct order on the [Dashboard](#), and [Reviews list](#) pages.
- If you are upgrading from Swarm 2020.2 or earlier and have **userids** that contain the forward slash (/) character, add **AllowEncodedSlashes NoDecode** to the **VirtualHost** block of your `/etc/apache2/sites-enabled/perforce-swarm-site.conf` file. For more information about the VirtualHost block, see ["Apache configuration" on page 186](#).

PHP and Apache 2.4 installation

CentOS and RHEL do not have PHP 7 or PHP 8 and Apache 2.4 by default so you must upgrade your system before you can upgrade Swarm. This process is only required the first time you upgrade to PHP 7 or PHP 8. If you have already upgraded to PHP 7 or PHP 8 and Apache 2.4, see ["Configure Redis on the Swarm machine" on page 284](#).

Carry out the following steps:

Important

The following notes are applicable for CentOS/RHEL7, RHEL 8, and RHEL 9:

- As part of the PHP upgrade process your Apache 2.2 server will be stopped and disabled, if you are currently using the Apache 2.2 server for any other applications they will stop working. You will either need to upgrade these applications to work with PHP 7 or PHP 8 and Apache 2.4 or move them to another server.
- **CentOS/RHEL 7:** Use the Remi repository configuration package (**remi-release-7.rpm**) to give Swarm access to PHP 7/PHP 8 and to LibreOffice which is part of the optional package install. Use the **epel-release-latest-7.noarch.rpm** repository configuration package to give Swarm access to EPEL packages.
- **RHEL 8:** Use the Remi repository configuration package (**remi-release-8.rpm**) to give Swarm access to PHP 7/PHP 8 and to LibreOffice which is part of the optional package install. Use the **epel-release-latest-8.noarch.rpm** repository configuration package to give Swarm access to EPEL packages.
- **RHEL 9:** Use the Remi repository configuration packages (**remi-release-8.rpm** and **remi-release-9.rpm**) to give Swarm access to PHP 7/PHP 8 and to LibreOffice which is part of the optional package install. Use the **epel-release-latest-8.noarch.rpm** and **epel-release-latest-9.noarch.rpm** repository configuration package to give Swarm access to EPEL packages.
- **Swarm 2019.1 to 2020.1:** these versions of Swarm used PHP packages from the SCL repositories for CentOS/RHEL 7. This was to provide access to more recent versions of PHP than are available as standard on CentOS/RHEL. This required the use of the **httpd24-httpd** package for Apache. These were all installed into **/opt/rh**

Swarm 2020.2 and later: these versions of Swarm uses the Remi repository for CentOS/RHEL 7, RHEL 8, and RHEL 9. This provides PHP 7.4 installed in the standard file system structure. This means that the old **httpd24-httpd** version of Apache is no longer needed, and the standard system version of Apache is being used again.

The SCL Apache site configuration file was installed at this location for Swarm 2019.1 to 2020.1:

```
/opt/rh/httpd24/root/etc/httpd/conf.d/perforce-swarm-site.conf
```

If this exists when Swarm is upgraded to 2020.2 and later, this file is copied to `/etc/httpd/conf.d/perforce-swarm-site.conf` if there is no file at the destination. It is also re-written to change references from `/var/log/httpd24` to `/var/log/httpd`

If a site configuration file for Swarm already exists in `/etc/httpd`, the copy and re-write is not performed.

After upgrade, `httpd24-httpd` is disabled.

- To avoid seeing the Apache HTTP server Linux test page when you start the Apache server, comment out the content of the `welcome.conf` file located in the `/etc/httpd/conf.d/` directory.
- To avoid loading the Apache HTTP server example configuration instead of the Swarm configuration when the Apache server starts, rename the `autoindex.conf` file located in the `/etc/httpd/conf.d/` directory to `z-autoindex.conf` or similar. This is required because Apache runs the first conf file it finds in the `/etc/httpd/conf.d/` directory (alphabetical order) and that must be the `perforce-swarm-site.conf` file.

This section describes how to install PHP 7 or PHP 8 and Apache 2.4 on your system using a repository .

1. This process is only required the first time you upgrade to PHP 7 or PHP 8. Install PHP 7 or PHP 8 and migrate the Swarm Apache configuration to use Apache 2.4.

Follow the instructions for your OS distribution, see [CentOS/REHL 7](#), [RHEL 8](#), or [RHEL 9](#) below:

- **CentOS/RHEL 7** (run these commands as root):

- a. Deploy the `epel-release-latest-7.noarch.rpm` repository configuration package to give Swarm access to EPEL packages:

```
yum install https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm
```
- b. Deploy the Remi repository configuration package to give Swarm access to PHP 7 or PHP 8 (only required the first time you upgrade to PHP 7 or PHP 8):

```
yum install https://rpms.remirepo.net/enterprise/remi-release-7.rpm
```
- c. Install the `yum-utils` package to give access to the `yum-config-manager` command:

```
yum install yum-utils
```
- d. **RHEL only:** Enable the optional channel for some dependencies:

```
subscription-manager repos --enable=rhel-7-server-optional-rpms
```
- e. Install the **Default/Single version** of PHP 7 or PHP 8:
 - i. Disable `remi-php*`:

```
yum-config-manager --disable 'remi-php*'
```
 - ii. Enable PHP 7 or PHP 8. Based on the version of PHP that you want to enable, use either of the following commands

```
yum-config-manager --enable remi-php7x
```

```
yum-config-manager --enable remi-php8x
```

Where **x** is the version of PHP 7 or PHP 8.
 - iii. Run an upgrade for PHP:

```
yum update
```

Important

If you are upgrading from Swarm 2019.3 to Swarm 2021.1 or later, remove the Swarm PHP 7.1 files:

```
yum remove $(yum list installed | grep php71 | awk '{print $1}')
```

- f. If you are upgrading from Apache 2.2, disable your Apache 2.2 HTTP server so that it does not automatically start:

```
systemctl disable httpd
```
- g. If you are upgrading from Apache 2.2, stop your Apache 2.2 HTTP server:

```
systemctl stop httpd
```

- h. If you are upgrading from PHP 7.2, copy the `30-perforce.ini` file to `/etc/php.d/` directory, this step also changes the `ini` filename to `perforce.ini`:

```
cp /etc/opt/rh/rh-php72/php.d/30-perforce.ini
/etc/php.d/perforce.ini
```

- i. If you are upgrading from PHP 7.2 to PHP 7.4, 8.0, 8.1, or 8.2 replace the PHP 7.2 references with PHP version number in the `perforce.ini` file using the `sed` command. For example, to upgrade from PHP 7.2 to PHP 7.4 use the following `sed` command:

```
sed -i "s/72/74/g" /etc/php.d/perforce.ini
```

- j. If you have any special `php.ini` configuration options set, copy the `php.ini` file to the `/etc/php.d/` directory, for example:

```
cp /etc/opt/rh/rh-php72/php.d/php.ini
/etc/php.d/php.ini
```

- k. Copy the `perforce-swarm-site.conf` file to the `/etc/httpd/conf.d/` directory if it is not already in there, for example:

```
cp /opt/rh/httpd24/root/etc/httpd/conf.d/perforce-
swarm-site.conf /etc/httpd/conf.d/perforce-swarm-
site.conf
```

- l. Replace the `/log/httpd24` filepath with `/log/httpd/` in the `perforce-swarm-site.conf` file using the `sed` command:

```
sed -i "s#/log/httpd24/#/log/httpd/#g"
/etc/httpd/conf.d/perforce-swarm-site.conf
```

- m. Enable your Apache 2.4 HTTP server so that it automatically starts:

```
systemctl enable httpd
```

- n. Start the Apache 2.4 HTTP server:

```
systemctl start httpd
```

- o. Check that SELinux is working correctly for your system. For instructions on configuring SELinux on CentOS/RHEL, see "[SELinux configuration](#)" on page 149.

■ **RHEL 8** (run these commands as root):

- a. Deploy the `epel-release-latest-8.noarch.rpm` repository configuration package to give Swarm access to EPEL packages:

```
dnf install https://dl.fedoraproject.org/pub/epel/epel-
release-latest-8.noarch.rpm
```

- b. Deploy the Remi repository configuration package to give Swarm access to PHP 7 or PHP 8 (only required the first time you upgrade to PHP 7 or PHP 8):

```
dnf install https://rpms.remirepo.net/enterprise/remi-release-8.rpm
```

- c. Install the `yum-utils` package to give access to the `yum-config-manager` command:

```
dnf install yum-utils
```

- d. Enable the optional channel for some dependencies:

```
subscription-manager repos --enable=rhel-8-server-optional-rpms
```

- e. Install the **Default/Single version** of PHP 7 or PHP 8:

- i. Enable the module stream for PHP:

```
dnf module reset php
```

- ii. Install PHP.

If you want to install PHP 7.4, use the following command:

```
dnf module install php:remi-7.4
```

To install a specific version of PHP, replace the PHP 7.4 reference in the above command.

- iii. Run an upgrade for PHP:

```
dnf update
```

Important

If you are upgrading from Swarm 2019.3 to Swarm 2021.1 or later, remove the Swarm PHP 7.1 files:

```
yum remove $(yum list installed | grep php71 | awk '{print $1}')
```

- f. If you are upgrading from Apache 2.2, disable your Apache 2.2 HTTP server so that it does not automatically start:

```
systemctl disable httpd
```

- g. If you are upgrading from Apache 2.2, stop your Apache 2.2 HTTP server:

```
systemctl stop httpd
```

- h. If you are upgrading from PHP 7.2, copy the `30-perforce.ini` file to `/etc/php.d/` directory, this step also changes the `ini` filename to `perforce.ini`:

```
cp /etc/opt/rh/rh-php72/php.d/30-perforce.ini
/etc/php.d/perforce.ini
```

- i. If you are upgrading from PHP 7.2, replace the PHP 7.2 references with PHP version number in the `perforce.ini` file using the `sed` command. For example, to upgrade from PHP 7.2 to PHP 7.4 use the following `sed` command:

```
sed -i "s/72/74/g" /etc/php.d/perforce.ini
```

- j. If you have any special `php.ini` configuration options set, copy the `php.ini` file to the `/etc/php.d/` directory, for example:

```
cp /etc/opt/rh/rh-php72/php.d/php.ini
/etc/php.d/php.ini
```

- k. Copy the `perforce-swarm-site.conf` file to the `/etc/httpd/conf.d/` directory if it is not already in there, for example:

```
cp /opt/rh/httpd24/root/etc/httpd/conf.d/perforce-
swarm-site.conf /etc/httpd/conf.d/perforce-swarm-
site.conf
```

- l. Replace the `/log/httpd24` filepath with `/log/httpd/` in the `perforce-swarm-site.conf` file using the `sed` command:

```
sed -i "s#/log/httpd24/#/log/httpd/#g"
/etc/httpd/conf.d/perforce-swarm-site.conf
```

- m. Enable your Apache 2.4 HTTP server so that it automatically starts:

```
systemctl enable httpd
```

- n. Start the Apache 2.4 HTTP server:

```
systemctl start httpd
```

- o. Check that SELinux is working correctly for your system. For instructions on configuring SELinux on CentOS/RHEL, see ["SELinux configuration" on page 149](#).

■ **RHEL 9** (run these commands as root):

- a. Run an upgrade for PHP, this will also upgrade the Swarm packages:

```
dnf update
```

- b. If you didn't run `dnf update` in the previous step, run the Swarm package upgrade now:

```
yum install helix-swarm helix-swarm-triggers helix-
swarm-optional
```

Important

If you are upgrading from Swarm 2019.3 to Swarm 2021.1 or later, remove the Swarm PHP 7.1 files:

```
yum remove $(yum list installed | grep php71 | awk '{print $1}')
```

- c. **Optional, ImageMagick:** when ImageMagick is enabled, Swarm can display the following image formats: BMP, EPS, PSD, TGA, TIFF.
 - i. Install ImageMagick:


```
yum install ImageMagick
```
 - ii. Restart the web server so that ImageMagick is picked up.


```
systemctl restart httpd
```
 - d. Check that SELinux is working correctly for your system. For instructions on configuring SELinux on CentOS/RHEL, see ["SELinux configuration" on page 149](#).
2. PHP 7 or PHP 8 and Apache 2.4 are now installed, configure the Swarm Redis service on the Swarm machine, see ["Configure Redis on the Swarm machine" below](#).

Configure Redis on the Swarm machine

Follow the instructions for the Swarm version you are upgrading from:

- ["Upgrade from Swarm 2019.1 and earlier" below](#)
- ["Upgrade from Swarm 2019.2 and later" on page 288](#)

Upgrade from Swarm 2019.1 and earlier

Swarm requires Redis to manage its caches and by default Swarm uses its own Redis server on the Swarm machine. Swarm caches data from the Helix Core Server to improve the performance of common searches in Swarm and to reduce the load on the Helix Core Server. Redis is included in the Swarm Tarball installation.

This section describes how to configure the Swarm Redis service on the Swarm machine. Do not change the default values of the following configuration files if you are using the Swarm Redis server.

Tip

If required, you can use your own Redis server instead of the Swarm Redis server. For instructions on how to configure Swarm to use your own Redis server, see ["Use your own Redis server" on page 235](#).

Swarm has two Redis binaries in `SWARM_ROOT/p4-bin/bin.linux26x86_64/`:

- `redis-server-swarm`
- `redis-cli-swarm`

1. Configure the Redis cache database, enter the following details in the **redis-server.conf** file in **/opt/perforce/etc/**:

```
bind 127.0.0.1
port 7379
supervised auto
save ""
dir /opt/perforce/swarm/redis
```

Default values:

Tip

- The default settings are shown below, the **redis-server.conf** file contains more detailed information about the Redis configuration for Swarm.
- On Swarm systems with a large number of users, groups, and projects, start-up time can be improved by persisting the memory cache. You can persist the memory cache by disabling background saves and enabling append saves, see the **redis-server.conf** file comments for detailed information.

- **bind 127.0.0.1** - Redis server IP address (loopback interface)
- **port 7379** - Redis server port number
- **supervised auto** - detects the use of **upstart** or **systemd** automatically to signal that the process is ready to use the supervisors
- **save ""** - background saves disabled, recommended.
- **dir /opt/perforce/swarm/redis** - the directory the Redis cache database is stored in.

Tip

To fine-tune your Redis server settings, make your changes in the Laminas Redis adapter. For information about the Laminas Redis adapter, see the [Laminas Redis Adapter documentation](#).

2. The Redis cache database must be running before Swarm starts so we recommend setting it up as a service so that it starts automatically at boot time.

The following example script will:

- run the Redis service as the Perforce user
- use the configuration file in **/opt/perforce/etc/redis-server.conf**
- assume the database server is installed in **/opt/perforce/swarm/p4-bin/bin.linux26x86_64/**

To configure Redis as a service, add a script with the following content in **/etc/systemd/system**

redis-server-swarm.service

```
[Unit]
Description=Redis Server for Swarm
After=network.target

[Service]
ExecStart=/opt/perforce/swarm/p4-bin/bin.linux26x86_64/redis-server-swarm /opt/perforce/etc/redis-server.conf
ExecStop=/opt/perforce/swarm/p4-bin/bin.linux26x86_64/redis-cli-swarm shutdown
Restart=always
RestartSec=10
StandardOutput=syslog
StandardError=syslog
SyslogIdentifier=redis
User=perforce
    Group=perforce

[Install]
WantedBy=multi-user.target
```

3. Create the **`SWARM_ROOT/data/config.php`** file if it does not already exist. The **redis** block of the **`SWARM_ROOT/data/config.php`** file contains the **password**, **namespace**, **host** and **port** details of the Swarm Redis server:

```
<?php
// this block should be a peer of 'p4'
'redis' => array(
    'options' => array(
        'password' => null, // Defaults to null
        'namespace' => 'Swarm',
        'server' => array(
            'host' => 'localhost', // Defaults to 'localhost'
            or enter your Redis server hostname
            'port' => '7379', // Defaults to '7379 or enter
```

```

your Redis server port
    ),
    ),
    'items_batch_size' => 100000,
    'check_integrity' => '03:00', // Defaults to '03:00' Use
one of the following two formats:
                                // 1) The time of day that
the integrity check starts each day. Set in 24 hour format with
leading zeros and a : separator
                                // 2) The number of seconds
between each integrity check. Set as a positive integer. Specify
'0' to disable the integrity check.
    'population_lock_timeout' => 300, // Timeout for initial
cache population. Defaults to 300 seconds.
    ),

```

Configurables:

- **password**: Redis server password. Defaults to **null** and should be left at default if using the Swarm Redis server.
- **namespace**: the prefix used for key values in the Redis cache. Defaults to **Swarm** and should be left at default if using the Swarm Redis server.

Note

If you have multiple Swarm instances running against a single Redis server, each Swarm server must use a different Redis **namespace**. This enables the cache data for the individual Swarm instances to be identified. The **namespace** is limited to ≤ 128 characters.

If one or more of your Swarm instances is connected to multiple Helix Core Servers, the Redis **namespace** includes the **server label** and the character limit is reduced to ≤ 127 characters, see ["Multiple Helix Core Server instances" on page 646](#).

- **host**: Redis server hostname. Defaults to **localhost** and should be left at default if using the Swarm Redis server.
- **port**: Redis server port number. Defaults to **7379**. Swarm uses port **7379** as its default to avoid clashing with other Redis servers that might be on your network. It should be left at default if using the Swarm Redis server. The default port for a non-Swarm Redis server is **6379**.

- **items_batch_size**: Maximum number of key/value pairs allowed in an **mset** call to Redis. Sets exceeding this will be batched according to this maximum for efficiency. Defaults to **100000**.

Note

The default value of **100000** was chosen to strike a balance between efficiency and project data complexity. This value should not normally need to be changed, contact support before making a change to this value.

- **check_integrity**: In some circumstances, such as when changes are made in the Helix Core Server when Swarm is down or if errors occur during updates, the Redis cache can get out of sync with the Helix Core Server. Swarm can run a regular integrity check to make sure that the Redis caches and Helix Core Server are in sync. If an integrity check finds an out of sync cache file, Swarm automatically updates the data in that cache.

The **check_integrity** configurable specifies when the Redis cache integrity check is run. Set as a specific time of day (24 hour format with leading zeros) or a number of seconds (positive integer) between checks. Disable the integrity check with **'0'**. Defaults to **'03:00'**.

- **population_lock_timeout**: specifies the timeout, in seconds, for initial cache population. If you have a large Swarm system, increase this time if the initial cache population times out. Defaults to **300** seconds.

Now that the Swarm Redis service is configured for the Swarm machine, start your Swarm upgrade, see ["Upgrade Swarm" on the next page](#).

Upgrade from Swarm 2019.2 and later

Swarm requires Redis to manage its caches and by default Swarm uses its own Redis server on the Swarm machine. Swarm caches data from the Helix Core Server to improve the performance of common searches in Swarm and to reduce the load on the Helix Core Server. Redis is included in the Swarm Tarball installation.

Note

If you are using your own Redis server instead of the Swarm Redis server, upgrade your Redis server to the version supplied with your Swarm upgrade.

Swarm has two Redis binaries in ***SWARM_ROOT/p4-bin/bin.linux26x86_64/***:

- **redis-server-swarm**
- **redis-cli-swarm**

If the Redis binaries supplied with your new version of Swarm have been updated, your Redis server is automatically upgraded to the new Redis server version. Your original Redis configuration files will not be changed. Do not change the default values of the Redis configuration files if you are using the Swarm Redis server.

Restart the Redis server:

If the Redis binaries have been updated and you are [running Redis as a service](#) (recommended), you must restart the Redis server to use the new Redis server version.

1. Run the following command to restart the Redis server:

```
systemctl restart redis-server-swarm
```

2. You are now running the updated Redis server version.

Tip

On Swarm systems with a large number of users, groups, and projects, start-up time can be improved by persisting the memory cache. You can persist the memory cache by disabling background saves and enabling append saves, see the `redis-server.conf` file comments for detailed information.

3. Now that the Swarm Redis service is configured for the Swarm machine, start your Swarm upgrade, see "[Upgrade Swarm](#)" below.

Upgrade Swarm

The steps in this section describe how to upgrade Swarm using the provided archive file. `SWARM_ROOT` refers to the current Swarm installation.

1. Download the new TAR file from <https://www.perforce.com/downloads/helix-swarm>.
2. Expand the new `swarm.tgz`:

```
tar -zxvf swarm.tgz
```

The contents of `swarm.tgz` are expanded into a top-level folder named `swarm-version`, where `version` corresponds to the version downloaded. This directory is identified as `SWARM_NEW` below.

3. Move `SWARM_NEW` to be a peer of `SWARM_ROOT`:

```
mv SWARM_NEW SWARM_ROOT/..
```

4. Copy the `SWARM_ROOT/data/config.php` file from `SWARM_ROOT` to `SWARM_NEW`:

```
cp -p SWARM_ROOT/data/config.php SWARM_NEW/data/
```

5. Create the queue token directory:

```
mkdir SWARM_NEW/data/queue
```

6. Copy the existing trigger token(s):

```
sudo cp -pR SWARM_ROOT/data/queue/tokens SWARM_NEW/data/queue/
```

7. Configure the ImageMagick `policy.xml` file to allow Swarm to create thumbnails if it is not already configured:
 - a. Open the ImageMagick `policy.xml` file for editing:
`/etc/ImageMagick-6/policy.xml`
 - b. Disable lines with the pattern below for the following `<filetype>`s: `PS`, `PS2`, `EPS`, `PDF`, or `XPS`:
`<policy domain="coder" rights="none" pattern="<filetype>" />`
 - c. Save the `policy.xml` file.
8. Assign correct ownership to the new Swarm's data directory:

```
sudo chown -R www-data SWARM_NEW/data
```

Note

The `www-data` user above is an example of what the web server user name might be, and can vary based on distribution or customization. For example, the user is typically `apache` for Red Hat/CentOS and `www-data` for Debian/Ubuntu.

9. Update your triggers, see "[Update your triggers](#)" below.

Update your triggers

Important

If you are using Swarm Helix Core Server Triggers, the Swarm Helix Core Server extension must not be installed.

1. Copy the new Swarm trigger script to your Helix Core Server machine. The trigger script is `SWARM_ROOT/p4-bin/scripts/swarm-trigger.pl`, and requires installation of Perl 5.08+ (use the latest available) on the Helix Core Server machine. If Swarm is using SSL, then the triggers also require the `IO::Socket::SSL` Perl module.

Warning

Do not overwrite any existing trigger script at this time. Give the script a new name, for example: `swarm-trigger-new.pl`.

2. Configure the Swarm trigger script by creating, in the same directory on the Helix Core Server machine, `swarm-trigger.conf`. It should contain:

Note

If you already have a `swarm-trigger.conf` file, no additional configuration is required.

```
# SWARM_HOST (required)
# Hostname of your Swarm instance, with leading "http://" or
# "https://".
SWARM_HOST="http://my-swarm-host"

# SWARM_TOKEN (required)
# The token used when talking to Swarm to offer some security. To
# obtain the
# value, log in to Swarm as a super user and select 'About Swarm'
# to see the
# token value.
SWARM_TOKEN="MY-UUID-STYLE-TOKEN"

# ADMIN_USER (optional) Do not use if the Workflow feature is
# enabled (default)
# For enforcing reviewed changes, optionally specify the normal
# Perforce user
# with admin privileges (to read keys); if not set, will use
# whatever Perforce
# user is set in environment.
ADMIN_USER=

# ADMIN_TICKET_FILE (optional) Do not use if the Workflow feature
# is enabled (default)
# For enforcing reviewed changes, optionally specify the location
# of the
# p4tickets file if different from the default
# ($HOME/.p4tickets).
# Ensure this user is a member of a group with an 'unlimited' or
# very long
# timeout; then, manually login as this user from the Perforce
# server machine to
# set the ticket.
ADMIN_TICKET_FILE=
```

```
# VERIFY_SSL (optional)
# If HTTPS is being used on the Swarm web server, then this
controls whether
# the SSL certificate is validated or not. By default this is set
to 1, which
# means any SSL certificates must be valid. If the web server is
using a self
# signed certificate, then this must be set to 0.
# set the ticket.
VERIFY_SSL=1
```

Fill in the required **SWARM_HOST** and **SWARM_TOKEN** variables with the configuration from any previous Swarm trigger script, typically **swarm-trigger.pl**.

Tip

The **ADMIN_USER** and **ADMIN_TICKET** variables were used by the '**enforce triggers**' in Swarm 2019.1 and earlier. They can be removed unless you are explicitly [disabling workflow](#) and using the deprecated '**enforce triggers**'.

Note

Swarm 2015.4 and earlier: Swarm trigger script files were available as shell scripts in these earlier Swarm versions, typically **swarm-trigger.sh**.

Swarm must now use a Perl trigger script file, typically **swarm-trigger.pl**.

3. **On Linux:** ensure that the script is executable:

```
sudo chmod +x swarm-trigger-new.pl
```

4. Rename the new trigger script:

On Linux:

```
mv swarm-trigger-new.pl swarm-trigger.pl
```

On Windows:

```
ren swarm-trigger-new.pl swarm-trigger.pl
```

5. Update the triggers in your Helix Core Server.

Warning

- The **swarm.shelvedel shelve-delete** trigger line was added to Swarm in version 2018.1 and updated in version 2020.1.

- **Upgrading from Swarm 2017.4 and earlier:** add the `swarm.shelvedel shelve-delete` trigger line to the Helix Core Server trigger table if it is not already present, see "[Update the Helix Core Server triggers table to run the trigger script.](#)" on page 211
- **Upgrading from Swarm 2018.x and 2019.x:** replace the existing `swarm.shelvedel shelve-delete` trigger line in the Helix Core Server trigger table with the one supplied in the Swarm version you are upgrading to.

■ **Workflow feature:**

The [Workflow feature](#) is enabled by default in Swarm 2019.2 and later. The trigger lines required when workflow is enabled are different to those required when workflow is disabled:

- **Workflow feature [enabled](#) (default):**
 - Comment out the `swarm.enforce.1`, `swarm.enforce.2`, `swarm.strict.1`, and `swarm.strict.2` trigger lines in the Helix Core Server trigger table if they are present, see "[Update the Helix Core Server triggers table to run the trigger script.](#)" on page 211
 - Add the `swarm.enforce change-submit`, `swarm.strict change-content`, and `swarm.shelvesub shelve-submit` trigger lines to the Helix Core Server trigger table if they are not already present, see "[Update the Helix Core Server triggers table to run the trigger script.](#)" on page 211
- **Workflow feature [disabled](#):**

Comment out the `swarm.enforce change-submit`, `swarm.strict change-content`, and `swarm.shelvesub shelve-submit` trigger lines in the Helix Core Server trigger table if they are present, see "[Update the Helix Core Server triggers table to run the trigger script.](#)" on page 211

- Run the Swarm trigger script to capture (using **Ctrl+C** on Windows and Linux) the trigger lines that should be included in the Perforce trigger table:

On Linux:

```
./swarm-trigger.pl -o
```

On Windows:

```
path/to/perl swarm-trigger.pl -o
```

- As a Perforce user with *super* privileges, update the Perforce trigger table by running **p4 triggers** command and replacing any `swarm.*` lines with the previously captured trigger line output (using **Ctrl+V** on Windows and Linux).

Important

If you previously customized the Swarm trigger lines, perhaps to apply various ["Trigger options" on page 735](#), be sure to repeat those customizations within the updated trigger lines.

Replace your old Swarm instance with your new Swarm instance

Important

Downtime occurs in this step.

Replace your old Swarm instance with your new Swarm instance. Run the following command:

```
sudo apache2ctl stop; mv SWARM_ROOT SWARM.old; mv SWARM_NEW SWARM_ROOT; sudo apache2ctl start
```

Note

If you see the following error message when you start Swarm, Swarm is using the wrong version of P4PHP. The latest version of P4PHP is included in the Swarm tarball but you must configure Swarm to use that version of P4PHP. For instructions about how to configure Swarm to use the new version of P4PHP, see ["PHP configuration" on page 190](#).

Swarm has detected a configuration error

Problem detected:

- The Perforce PHP extension (P4PHP) requires upgrading. Found 2016.2, only 2018.2 or later is supported.

The php.ini file loaded is /etc/php5/apache2/php.ini.

Other scanned php.ini files (in /etc/php5/apache2/conf.d) include:

- /etc/php5/apache2/conf.d/imagick.ini
- /etc/php5/apache2/conf.d/pdo.ini
- /etc/php5/apache2/conf.d/perforce.ini

Please investigate the below PHP error:

The Perforce PHP extension (P4PHP) requires upgrading.
Found 2016.2, only 2018.2 or later is supported.

/var/www/deployments/1703508/swarm/public/index.php on line 104

For more information, please see the [Setting Up](#) documentation; in particular:

- [Initial Installation](#)
- [Runtime dependencies](#)
- [PHP configuration](#)
- [Swarm configuration](#)

Please ensure you restart your web server after making any PHP changes.

Reload the Swarm config cache

As an *admin* user, delete and reload the Swarm configuration cache from the **Cache Info** tab of the **System Information** page. Navigate to the **User id** dropdown menu, select **System Information**, select the **Cache Info** tab, and select **Reload Configuration**. For more information, see "[Cache Info](#)" on page 731.

Validate your upgrade

Tip

When Swarm starts it verifies the Redis cache, during this time you cannot log in to Swarm. The time taken to verify the Redis cache depends on the number of users, groups, and projects Swarm has. Start-up time can be improved by persisting the memory cache. You can persist the memory cache by disabling background saves and enabling append saves in the `redis-server.conf` file, see "[Redis server configuration file](#)" on page 676.

Check that your new Swarm instance is working correctly by doing the following:

1. Create a new changelist that:
 - a. Contains at least one modified file
 - b. Contains the `#review` keyword in the changelist description
2. Right click on the new changelist in P4V and click **Shelve Files...**

Important

Do not select **Request New Swarm Review...** because this method uses the API and will not fully test the Swarm Helix Core Server extension.

This is also true if you are using Swarm Helix Core Server Triggers instead of the Swarm Helix Core Server extension.

3. Check that a new Swarm review is created for the changelist.
 - If a Swarm review is created, the Swarm Helix Core Server extension is working. If you are using Swarm Helix Core Server Triggers instead of the Swarm Helix Core Server extension and the review is created, the triggers are working.
 - If a Swarm review is not created, see "[Review not created](#)" on page 240.

Swarm index upgrade

If you are upgrading from Swarm 2017.2 or earlier you should run the index upgrade, this ensures that the review activity history is displayed in the correct order on the [Dashboard](#), and [Reviews list](#) pages.

Note

If you are upgrading from Swarm version 2017.3 or later, the index upgrade step is not required.

The index upgrade process can be configured to suit your Swarm system specifications. See [Upgrade index](#) for details.

Run the upgrade as an *Admin* user by visiting the following URL:

```
http://SWARM-HOST/upgrade
```

Note

After the Swarm upgrade, on the first visit to some Swarm pages, users might see a message to perform a browser hard refresh. This happens because we are updating the UI and content of some Swarm pages, so the user's page cache is no longer valid and requires a hard refresh to load the upgraded page. For example, for Chrome on Windows/Linux **[CTRL]+[F5]**.

All done!

Moving your Swarm instance

This section describes how to move your Swarm instance to another server. This can be useful if you want to move Swarm to a different server to improve performance or if you want to run Swarm on a newer operating system. Swarm data is held on the Helix Core Server and not in Swarm so there is no need to migrate any data to your new Swarm instance.

Note

If you are installing a newer version of Swarm than your original Swarm instance, you might need to configure options if they are not already present in your original version of Swarm.

For example, Swarm workflow, updated trigger scripts, the introduction of Extensions, the introduction of a Redis cache, and changes to custom modules. For information on important changes to Swarm, see the ["Important information" on page 24](#) section of the Swarm ["What's new" on page 23](#) page.

Summary

Moving your Swarm installation to a new server results in downtime, but the following process minimizes it. The process for moving your Swarm installation to another server is summarized below:

1. Install Swarm on a new server.
2. Copy your original Swarm `config.php` file and authentication token to your new Swarm instance.
3. **Only required if you are moving to a newer version of Swarm:** Configure your new Swarm instance to receive event notifications from the Helix Core Server.
4. Shutdown your original Swarm instance and start your new Swarm instance .

Warning

Downtime occurs at this step.

5. Validate that your new Swarm instance is working correctly.

For step by step instructions for moving your Swarm instance, start by ["Installing Swarm on your new server" below](#).

Installing Swarm on your new server

Choose one of the following installation methods (we recommend the package installation method whenever possible):

- ["Install and configure Swarm from a package \(recommended\)" on page 115](#). Do not follow the **Post-installation configuration** steps. Copy your original Swarm `config.php` file to your new Swarm instance, see ["Copying your existing Swarm configuration" below](#).
- ["Install and configure Swarm manually from a Tarball" on page 181](#), configure [Redis](#), configure [Apache](#), and configure [PHP](#). Copy your original Swarm `config.php` file to your new Swarm instance, see ["Copying your existing Swarm configuration" below](#).

Copying your existing Swarm configuration

1. Copy your original `SWARM_ROOT/data/config.php` file to your new Swarm instance.
2. Create the queue token directory on your new Swarm instance:

```
mkdir SWARM_ROOT/data/queue
```
3. Copy your original trigger token(s) from the `SWARM_ROOT/data/queue/tokens` to your new Swarm instance.
4. Assign correct ownership to the new Swarm instance data directory:

```
sudo chown -R www-data SWARM_ROOT/data
```

Note

The `www-data` user above is an example of what the web server user name might be, and can vary based on distribution or customization. For example, the user is typically `apache` for Red Hat/CentOS and `www-data` for Debian/Ubuntu.

5. Do one of the following:
 - **If moving to an instance with the same Swarm version:** go to ["Replacing your original Swarm instance with your new instance" on page 303](#).
 - **If moving to an instance with a newer version of Swarm:** go to ["Configuring Helix Core Server event notification" below](#).

Configuring Helix Core Server event notification

Important

The steps in this section are only required if you are moving your Swarm instance to a newer version of Swarm. If the new and original Swarm instances are the same Swarm version, go to ["Replacing your original Swarm instance with your new instance" on page 303](#).

Swarm needs to know about a number of Helix Core Server events to operate correctly, this can be done by using Helix Core Server Extensions or Helix Core Server Triggers.

You must configure your new Swarm instance to receive event notifications from the Helix Core Server to complete the Swarm move.

Do one of the following so that Swarm is notified about events on the Helix Core Server:

- **If your old Swarm instance is using Helix Core Server Extensions:** update the Swarm Helix Core Server extension configuration, see ["Updating your Swarm Helix Core Server extension" below](#)
- **If your old Swarm instance is using Helix Core Server Triggers:** update your Swarm Triggers, see ["Updating your triggers" below](#)

Updating your Swarm Helix Core Server extension

Important

You must be a user with *super* user permissions to configure Helix Core Server Extensions.

If you are moving a Swarm instance that is already using Helix Core Server Extensions, there is no need to reinstall the extension because it is already installed on Helix Core Server.

To update your Swarm Helix Core Server extension:

1. Open the extension configuration file with:

```
p4 extension --configure Perforce::helix-swarm
```

The spec file opens in your text editor
2. Edit the configuration for your new Swarm instance as required.
3. **Swarm-URL:** do one of the following:
 - **If the hostname of your Swarm server has changed:** set the **Swarm-URL** to the URL of your new Swarm server instance.
 - **If the Swarm server hostname has not changed:** you do not need to change the **Swarm-URL** in the configuration file.
4. Save the configuration changes.
5. Continue with the next task, see ["Replacing your original Swarm instance with your new instance" on page 303](#).

Updating your triggers

Important

If you are using Swarm Helix Core Server Triggers, the Swarm Helix Core Server extension must not be installed.

1. Copy the new Swarm trigger script to your Helix Core Server machine. The trigger script is **SWARM_ROOT/p4-bin/scripts/swarm-trigger.pl**, and requires installation of Perl 5.08+ (use the latest available) on the Helix Core Server machine. If Swarm is using SSL, then the triggers also require the **IO::Socket::SSL** Perl module.

Warning

Do not overwrite any existing trigger script at this time. Give the script a new name, for example: **swarm-trigger-new.pl**.

2. Configure the Swarm trigger script by creating, in the same directory on the Helix Core Server machine, **swarm-trigger.conf**. It should contain:

Note

If you already have a **swarm-trigger.conf** file, no additional configuration is required.

```
# SWARM_HOST (required)
# Hostname of your Swarm instance, with leading "http://" or
# "https://".
SWARM_HOST="http://my-swarm-host"

# SWARM_TOKEN (required)
# The token used when talking to Swarm to offer some security. To
# obtain the
# value, log in to Swarm as a super user and select 'About Swarm'
# to see the
# token value.
SWARM_TOKEN="MY-UUID-STYLE-TOKEN"

# ADMIN_USER (optional) Do not use if the Workflow feature is
# enabled (default)
# For enforcing reviewed changes, optionally specify the normal
# Perforce user
# with admin privileges (to read keys); if not set, will use
# whatever Perforce
# user is set in environment.
ADMIN_USER=
```



```
# ADMIN_TICKET_FILE (optional) Do not use if the Workflow feature
is enabled (default)
# For enforcing reviewed changes, optionally specify the location
of the
# p4tickets file if different from the default
($HOME/.p4tickets).
# Ensure this user is a member of a group with an 'unlimited' or
very long
# timeout; then, manually login as this user from the Perforce
server machine to
# set the ticket.
ADMIN_TICKET_FILE=

# VERIFY_SSL (optional)
# If HTTPS is being used on the Swarm web server, then this
controls whether
# the SSL certificate is validated or not. By default this is set
to 1, which
# means any SSL certificates must be valid. If the web server is
using a self
# signed certificate, then this must be set to 0.
# set the ticket.
VERIFY_SSL=1
```

Fill in the required **SWARM_HOST** and **SWARM_TOKEN** variables with the configuration from any previous Swarm trigger script, typically **swarm-trigger.pl**.

Tip

The **ADMIN_USER** and **ADMIN_TICKET** variables were used by the **'enforce triggers'** in Swarm 2019.1 and earlier. They can be removed unless you are explicitly [disabling workflow](#) and using the deprecated **'enforce triggers'**.

Note

Swarm 2015.4 and earlier: Swarm trigger script files were available as shell scripts in these earlier Swarm versions, typically **swarm-trigger.sh**.

Swarm must now use a Perl trigger script file, typically **swarm-trigger.pl**.

3. **On Linux:** ensure that the script is executable:

```
sudo chmod +x swarm-trigger-new.pl
```

4. Rename the new trigger script:

On Linux:

```
mv swarm-trigger-new.pl swarm-trigger.pl
```

On Windows:

```
ren swarm-trigger-new.pl swarm-trigger.pl
```

5. Update the triggers in your Helix Core Server.

Warning

- The **swarm.shelvedel shelve-delete** trigger line was added to Swarm in version 2018.1 and updated in version 2020.1.

- **Upgrading from Swarm 2017.4 and earlier:** add the **swarm.shelvedel shelve-delete** trigger line to the Helix Core Server trigger table if it is not already present, see "[Update the Helix Core Server triggers table to run the trigger script.](#)" on page 211
- **Upgrading from Swarm 2018.x and 2019.x:** replace the existing **swarm.shelvedel shelve-delete** trigger line in the Helix Core Server trigger table with the one supplied in the Swarm version you are upgrading to.

- **Workflow feature:**

The [Workflow feature](#) is enabled by default in Swarm 2019.2 and later. The trigger lines required when workflow is enabled are different to those required when workflow is disabled:

- **Workflow feature [enabled](#) (default):**
 - Comment out the **swarm.enforce.1**, **swarm.enforce.2**, **swarm.strict.1**, and **swarm.strict.2** trigger lines in the Helix Core Server trigger table if they are present, see "[Update the Helix Core Server triggers table to run the trigger script.](#)" on page 211
 - Add the **swarm.enforce change-submit**, **swarm.strict change-content**, and **swarm.shelvesub shelve-submit** trigger lines to the Helix Core Server trigger table if they are not already present, see "[Update the Helix Core Server triggers table to run the trigger script.](#)" on page 211
- **Workflow feature [disabled](#):**

Comment out the **swarm.enforce change-submit**, **swarm.strict change-content**, and **swarm.shelvesub shelve-submit** trigger lines in the Helix Core Server trigger table if they are present, see "[Update the Helix Core Server triggers table to run the trigger script.](#)" on page 211

- a. Run the Swarm trigger script to capture (using **Ctrl+C** on Windows and Linux) the trigger lines that should be included in the Perforce trigger table:

On Linux:

```
./swarm-trigger.pl -o
```

On Windows:

```
path/to/perl swarm-trigger.pl -o
```

- b. As a Perforce user with *super* privileges, update the Perforce trigger table by running **p4 triggers** command and replacing any **swarm.*** lines with the previously captured trigger line output (using **Ctrl+V** on Windows and Linux).

Important

If you previously customized the Swarm trigger lines, perhaps to apply various ["Trigger options" on page 735](#), be sure to repeat those customizations within the updated trigger lines.

6. Continue with the next task, see ["Replacing your original Swarm instance with your new instance" below](#).

Replacing your original Swarm instance with your new instance

Warning

Downtime occurs in this step.

Replace your original Swarm instance with your new Swarm instance.

1. Stop your original Swarm instance:

```
sudo apache2ctl stop
```

2. Start your new Swarm instance:

```
sudo apache2ctl start
```

Note

If you see the following error message when you start Swarm, Swarm is using the wrong version of P4PHP. The latest version of P4PHP is included with Swarm.

If you are installing Swarm from a tarball, you must configure Swarm to use that version of P4PHP. For instructions about how to configure Swarm to use the new version of P4PHP, see "PHP configuration" on page 190.

Swarm has detected a configuration error

Problem detected:

- The Perforce PHP extension (P4PHP) requires upgrading. Found 2016.2, only 2018.2 or later is supported.

The php.ini file loaded is /etc/php5/apache2/php.ini.

Other scanned php.ini files (in /etc/php5/apache2/conf.d) include:

- /etc/php5/apache2/conf.d/imagick.ini
- /etc/php5/apache2/conf.d/pdo.ini
- /etc/php5/apache2/conf.d/perforce.ini

Please investigate the below PHP error:

The Perforce PHP extension (P4PHP) requires upgrading.
Found 2016.2, only 2018.2 or later is supported.

/var/www/deployments/1703508/swarm/public/index.php on line 104

For more information, please see the [Setting Up](#) documentation; in particular:

- [Initial Installation](#)
- [Runtime dependencies](#)
- [PHP configuration](#)
- [Swarm configuration](#)

Please ensure you restart your web server after making any PHP changes.

3. Continue with the next task, see "Validating your new Swarm instance" on the next page.

Validating your new Swarm instance

Tip

When Swarm starts it verifies the Redis cache, during this time you cannot log in to Swarm. The time taken to verify the Redis cache depends on the number of users, groups, and projects Swarm has. Start-up time can be improved by persisting the memory cache. You can persist the memory cache by disabling background saves and enabling append saves in the `redis-server.conf` file, see ["Redis server configuration file" on page 676](#).

Check that your upgraded Swarm instance is working correctly by doing the following:

1. Create a new changelist that:
 - a. Contains at least one modified file
 - b. Contains the `#review` keyword in the changelist description
2. Right click on the new changelist in P4V and click **Shelve Files...**

Important

Do not select **Request New Swarm Review...** because this method uses the API and will not fully test the Swarm Helix Core Server extension.

This is also true if you are using Swarm Helix Core Server Triggers instead of the Swarm Helix Core Server extension.

3. Check that a new Swarm review is created for the changelist.
 - If a Swarm review is created, the Swarm Helix Core Server extension is working. If you are using Swarm Helix Core Server Triggers instead of the Swarm Helix Core Server extension and the review is created, the triggers are working.
 - If a Swarm review is not created, see ["Review not created" on page 240](#).
4. Do one of the following:
 - **If your new Swarm instance is working correctly:** Retire your original Swarm instance.
 - **If your new Swarm instance is not working correctly:** Shutdown the new Swarm instance, remove any updates you made to the triggers or extensions on the Helix Core Server, and switch back to your original Swarm instance. Check the configuration and connectivity of your new Swarm instance. If you can't find the issue, you can submit a [Support Request](#).

4 | Basics

This chapter covers the basic operations provided by Swarm. These include:

My Dashboard	306
Global Dashboard	310
Activity	321
Files	322
Commits	332
Jobs	335
Changelists	338
Diffs	341
Comments	348
Users	363
Groups	373
Projects	381
Workflows	388
Work-in-progress tag	393
Tests	395
Notifications	398
Log in/Log out	405
Notable minor features	407
Markdown	414

My Dashboard

The purpose of the dashboard is to allow you to focus on reviews that need to be done, so that other users are not blocked. The dashboard lists the most recently modified reviews first, and shows your role in the review.

Your dashboard is available by clicking **My Dashboard** in the menu or by clicking the Swarm  icon above the main menu.

Note

Since it is tied to the logged in user, **My Dashboard** is only populated if you are logged in.

My Dashboard						
Project <input type="text"/> Role <input type="text"/> Author <input type="text"/>			There are new changes to your review list Clear all		<input type="text" value="Search for a review"/>	
Description	Your Role	Last Activity	State	Tests	Complexity	Votes
Updated README to have a bit more information. Jack Boone requested a pre-commit review 365 for mercury:dev	Reviewer	44 minutes ago			137	0 0
Updated notes for the new year. Jack Boone requested a post-commit review 422 for jplugin:main	Required reviewer	2 months ago			18	0 0
Add comment to method. Steve Russell requested a post-commit review 328 for jplugin:main	Required reviewer	3 months ago			6	2 0
All available reviews are displayed						

A review is displayed on your dashboard if any of the following criteria are met:

- You are a reviewer or required reviewer, the review status is **Needs Review** and you have not already voted on it.
- You are a member of a reviewer group or a required reviewer group, the review status is **Needs Review** and you have not already voted on it. The review will remain on your dashboard even if the group has met its criteria if you have not already voted on it.
- You are the review author and the review status is **Needs Revision**, or **Approved** (only if the review is approved but not committed).
- You are a moderator or a member of a moderator group, the review status is **Needs Review**, the **Minimum up votes** requirement for the branch is satisfied, and one of the following is true:
 - There are no required reviewers.
 - All of the required reviewers have up-voted the review.
 - You are the last remaining required reviewer for the review.

Tip

If moderator behavior is configured to require approval from one moderator per branch and the review spans multiple moderated branches:

- You will see the review in your dashboard if the review has not been approved by another moderator from your branch.
- You will be able to **Approve and Commit** the review if it has been approved by moderators from all of the other branches.

Filtering

Project <input type="text"/> Role <input type="text"/> Author <input type="text"/>			There are new changes to your review list Clear all		<input type="text" value="Search for a review"/>	
--	--	--	---	--	--	--

The filters are used to filter your dashboard reviews:

- **Project:** filter by the project the review is part of, limiting results to **My Projects** (projects you are an owner of or a member of) or specific projects. To select multiple projects, click the projects you want to filter by. The **Project** filter will only show projects for which there are reviews in your dashboard.
- **Role:** filter by your specific role in a review, limiting results to reviews for which you are the author, a reviewer, a required reviewer, or a moderator. The **Role** filter will only show roles for which there are reviews in your dashboard.
- **Author:** filter the reviews to only those that have been authored by a certain user. Select the author from the drop down list of users. The **Author** filter will only show authors for which there are reviews in your dashboard.
- **Clear all:** click to reset all dashboard filters back to their defaults.
- **Search:** filter the reviews by searching the review descriptions.

Refresh button

Refresh button (only displayed when new content is available for your dashboard): click **Refresh** to update your dashboard with the new content.


Note

By default, when your dashboard is open, Swarm checks for new content every five minutes.

Review summaries

The dashboard displays a summary for each review.




Description	Your Role	Last Activity	State	Tests	Complexity	Votes
 Updated README to have a bit more information. Jack Boone requested a  pre-commit review 365 for mercury:dev	Reviewer	about an hour ago			 137	 2  0

- **Avatar** displays the avatar of the review author, click to view the profile of the author, see ["Viewing another user's profile" on page 371](#).
- **Description** contains:
 - first line of the review description, click to open the review. If the review description is too long it is truncated, click on the ellipsis  to expand it in the list page.
 - review author, click to view the profile of the author.
 - review type, either a [pre-commit review](#) or a [post-commit review](#).
 - review number, click to open the review, see ["Review display" on page 433](#).
 - review project branches, click to open the project.
- **Your role:** displays your role in the review and is the reason the review is in your dashboard. This can be Author, Reviewer, Required Reviewer, or Moderator.

- **Last activity:** displays the last time that any changes were made to the review, including votes, comments, commits, and file changes.
- **State:** a review can be in one of the following states:
 - **Needs review:** The review has started and the changes need to be reviewed.
 - **Needs revision:** The changes have been reviewed and the reviewer has indicated that further revisions are required.
 - **Approved:** The review has been approved. The changes may need to be committed.

Note

The Approved state only applies to review authors, it is only shown for a review that has been approved but has not been committed.

- **Tests:** shows the test suite state for the review, either tests in progress , tests passed , or tests failed .
- **Complexity:** a traffic light icon and number shows the relative complexity of the review and the total number of lines changed in the review. [Complexity can be configured](#) by your Swarm administrator but, by default:
 - **Red:** ≥ 300 changes
 - **Amber:** < 300 and > 30 changes
 - **Green:** ≤ 30 changes

Tip

- Review complexity is only calculated for a review when the review is updated and the file content has changed.
- Review complexity is only stored for the current version of a review.

Hover over the complexity icon to display more detailed information:



- **Votes:** displays the number of up votes and down votes for the review.

Global Dashboard

Note

Global Dashboard is only available if your Swarm administrator has configured Swarm to connect to more than one Helix Core Server instance. For instructions on how to configure Swarm to connect to multiple Helix Core Server instances, see "[Multiple Helix Core Server instances](#)" on page 646.

Your Global Dashboard displays a list of reviews that you may need to act on for each of the Helix Core Server instances connected to Swarm.

Open your global dashboard:

1. Enter the basic Swarm URL without a Helix Core Server instance name, for example:
`https://swarm.company.com`
2. If the log in dialog is displayed:

Tip

Single Sign-On (SSO)

If Helix Authentication Service is configured for your Helix Core Server and Swarm, [SSO can be configured](#) by your Swarm admin as either mandatory or optional:

- If SSO is mandatory, you will be directed to the sign-in process used by your Identity Provider (IdP).
- If SSO is optional, log in to Swarm using one of the following methods:
 - Click **Log in with SSO**, you will be directed to the sign-in process used by your Identify Provider (IdP).
 - Type in your username and password, and click **Log in**.

a. Server:

- If your Helix Core Server instances do not share log in credentials, select an individual Helix Core Server instance from the dropdown list.
- If all of your Helix Core Server instances share the same log in credentials, leave set to **All available servers**.

Note

If Helix Authentication Service is configured for your Helix Core Server and Swarm, you must log into the servers individually.

- b. Enter your username and password.

Select the **Remember Me** check box if you prefer to stay logged in between browser restarts.

Note

Helix Core Servers can enforce maximum login times. You may become logged out even if **Remember Me** is selected. Swarm administrators can change the maximum login time, see ["Sessions" on page 700](#) for details.

c. Click **Connect**.

Swarm will populate the global dashboard from the Helix Core Server instances it connects to.

The global dashboard is displayed:

Tip

- Since it is tied to the logged in user, the global dashboard is only populated for the Helix Core Server instances you are logged in to in Swarm.
- If you are already logged in to a Helix Core Server instance in Swarm, the dashboard for that server will be automatically populated when you open or refresh the global dashboard.

The screenshot shows the Global Dashboard interface. On the left is a sidebar with 'Instances 2/2 connected' listing 'main' (ssl:1666) and 'assets' (ssl:2666). The main area has a toolbar with filters like 'All Instances', 'All Projects', 'All Roles', and 'Authored by'. Below this is a table of activities:

Author	ID	Description	Project(s)	Your Role	State	Last activity
main						
[Avatar]	277	Added endpoint for initiating server shutdown.	mercury: dev	Reviewer	0/0	24 minutes ago
[Avatar]	279	Keep track of all out clients we have active.	mercury: dev	Author	0/0	9 days ago
[Avatar]	271	Added first set of icons for the application.	mercury: main	Reviewer	0/0	14 days ago
[Avatar]	135	Tidying up some of the code to fix some of the ed...	jplugin: main	Reviewer Moderator	1/0	14 days ago
[Avatar]	249	Adding description to P4Credentials & P4Credent...	jplugin: main	Moderator	1/0	14 days ago
assets						
[Avatar]	4	Updated the README with some context.	mercury: main	Author	0/0	29 minutes ago

The global dashboard is made up of three main areas:

- **"Toolbar" on the facing page:** log in to Helix Core Server instances, log out of all Helix Core Server instances, and open the Swarm help from the global dashboard toolbar.
- **"Sidebar" on page 314:** Log in/logout of individual Helix Core Server instances and view your profile for any instance that you are logged in to.

- ["Helix Core Server dashboards" on page 317](#): view, filter and search the dashboards of the Helix Core Server instances you are logged in to. Jump directly to a Helix Core Server, review, or project by clicking on a link.

Toolbar

Log in to one or more Helix Core Server instances

Log In is only displayed in the global dashboard toolbar if you are not logged in to any Helix Core Server instances. If **Log In** is not available, see ["Log in to a Helix Core Server instance" on page 315](#).

1. In the global dashboard toolbar, click **Log In**.

Tip Single Sign-On (SSO)

If Helix Authentication Service is configured for your Helix Core Server and Swarm, [SSO can be configured](#) by your Swarm admin as either mandatory or optional:

- If SSO is mandatory, you will be directed to the sign-in process used by your Identity Provider (IdP).
- If SSO is optional, log in to Swarm using one of the following methods:
 - Click **Log in with SSO**, you will be directed to the sign-in process used by your Identify Provider (IdP).
 - Type in your username and password, and click **Log in**.

×

Welcome to Global Dashboard

You must be connected to at least one Instance in order to access your Dashboard across all of your Instances. [Need more help?](#)

Server

assets x


x

▼

User Name

Password

☐ Remember Me

 **Connect**

2. Server:

- If all of your Helix Core Server instances share the same log in credentials, leave set to **All available servers**.
- If your Helix Core Server instances do not share log in credentials, select an individual Helix Core Server instance from the dropdown list.

Note

If Helix Authentication Service is configured for your Helix Core Server and Swarm, you must log into the servers individually.

Tip

If you have a number of Helix Core Server instances that use the same login credentials, select them one at a time from the dropdown list.

3. Enter your username and password.

Select the **Remember Me** check box if you prefer to stay logged in between browser restarts.

Note

Helix Core Servers can enforce maximum login times. You may become logged out even if **Remember Me** is selected. Swarm administrators can change the maximum login time, see ["Sessions" on page 700](#) for details.

4. Click **Connect**.

Swarm will populate the global dashboard from the Helix Core Server instances it connects to.

Logout of all Helix Core Server instances

Logout from all instances is only displayed in the global dashboard toolbar if you are logged in to at least one Helix Core Server instance.

1. In the global dashboard toolbar, click **Logout from all instances**.
2. When prompted, click **Yes** to confirm that you want to logout of all of the Helix Core Server instances.

Tip

- If Helix Authentication Service is configured for your Helix Core Server, logging out of Swarm will not invalidate your Identity Provider (IdP) login status. If you try to log back in to Swarm while your IdP status is still valid, you will not be prompted to complete the log in steps.

If **require_login** is also enabled, Swarm will return you to the login and your IdP will automatically log you back in. In this case log out from your Identity Provider page before logging out from Swarm.


- If a custom redirect has been configured by your Swarm administrator, you are logged out of all of the Helix Core Server instances by Swarm and then redirected to the URL specified by the administrator.

The custom redirect can be set to any internal or external URL, for example:

- Company intranet, extranet, internet, FTP, or Web-mail page
- Industry news website
- Identity Provider page to invalidate your IdP log in status


Sidebar

The Helix Core Servers that Swarm can connect to are listed in the collapsible sidebar on the left of the page. From the server list you can log in/logout from individual Helix Core Servers and view your user profile.




Log in status


- Not logged in to the Helix Core Server instance:



- Logged in to the Helix Core Server instance:



Log in to a Helix Core Server instance

1. In the sidebar, hover over the Helix Core Server you want to connect to.
2. Click the **Log in** button .

Tip

- **Single Sign-On (SSO)**

If Helix Authentication Service is configured for your Helix Core Server and Swarm, [SSO can be configured](#) by your Swarm admin as either mandatory or optional:

- If SSO is mandatory, you will be directed to the sign-in process used by your Identity Provider (IdP).
- If SSO is optional, log in to Swarm using one of the following methods:
 - Click **Log in with SSO**, you will be directed to the sign-in process used by your Identity Provider (IdP).
 - Type in your username and password, and click **Log in**.
- If you are already logged in to another Helix Core Server that is configured for Helix Authentication Service, your IdP status is valid and you will not be prompted to complete the log in steps.

Log In

assets


ssl:2666

User Name


Password

☒ Try to log in to all available servers with these credentials.

☐ Remember Me

 Connect

3. Type in your username and password for the Helix Core Server.

- Select the **Try to log in to all available servers with these credentials** checkbox if you use these credentials for more than one of the Helix Core Server instances. Swarm will not try to log in to any Helix Core Server instances that are configured for Helix Authentication Service, log in to them individually using the instance **Log in** button  in the sidebar.
- Select the **Remember Me** check box if you prefer to stay logged in between browser restarts.

Note


Helix Core Servers can enforce maximum login times. You may become logged out even if **Remember Me** is selected. Swarm administrators can change the maximum login time, see ["Sessions" on page 700](#) for details.

4. Click **Connect**.

Swarm will populate the global dashboard from the Helix Core Server.

If you selected the **Try to log in to all available servers with these credentials** checkbox, Swarm will populate the global dashboard for the other servers in the list that it successfully connects to.

Logout of a single Helix Core Server instance

1. In the sidebar, hover over the Helix Core Server you want to log out from.
2. Click the **Gear** button  to open the dropdown menu.
3. Click **Logout** in the dropdown menu.

Tip

- If Helix Authentication Service is configured for your Helix Core Server, logging out of Swarm will not invalidate your Identity Provider (IdP) login status. If you try to log back in to Swarm while your IdP status is still valid, you will not be prompted to complete the log in steps.

If **require_login** is also enabled, Swarm will return you to the login and your IdP will automatically log you back in. In this case log out from your Identity Provider page before logging out from Swarm.

- If a custom redirect has been configured by your Swarm administrator, you are logged out of the Helix Core Server instance by Swarm. You are only redirected to the URL specified by the administrator if you are not logged in to any other Helix Core Server instances on the global dashboard.

The custom redirect can be set to any internal or external URL, for example:

- Company intranet, extranet, internet, FTP, or Web-mail page
- Industry news website
- Identity Provider page to invalidate your IdP log in status

Helix Core Server dashboards

The purpose of the global dashboard is to allow you to focus on reviews that need to be done, so that other users are not blocked. The global dashboard lists reviews by Helix Core Server according to the most recently modified first, and shows your role in the review.

A review is displayed on your global dashboard if any of the following criteria are met:

- You are a reviewer or required reviewer, the review status is **Needs Review** and you have not already voted on it.
- You are a member of a reviewer group or a required reviewer group, the review status is **Needs Review** and you have not already voted on it. The review will remain on your dashboard even if the group has met its criteria if you have not already voted on it.
- You are the review author and the review status is **Needs Revision**, or **Approved** (only if the review is approved but not committed).
- You are a moderator or a member of a moderator group, the review status is **Needs Review**, the **Minimum up votes** requirement for the branch is satisfied, and one of the following is true:



















- There are no required reviewers.
- All of the required reviewers have up-voted the review.
- You are the last remaining required reviewer for the review.

Tip

If moderator behavior is configured to require approval from one moderator per branch and the review spans multiple moderated branches:

- You will see the review in your dashboard if the review has not been approved by another moderator from your branch.
- You will be able to **Approve and Commit** the review if it has been approved by moderators from all of the other branches.

Example global dashboard showing a number of Helix Core Server dashboards:

<div> <div>All Instances ▾</div> <div>All Projects ▾</div> <div>All Roles ▾</div> <div>Authored by</div> <div>Search</div> </div>									
Author	ID	Description	Project(s)	Your Role	State				Last activity
main									
	277	Added endpoint for initiating server shutdown.	mercury: dev	Reviewer	 			0 / 0	24 minutes ago
	279	Keep track of all out clients we have active.	mercury: dev	Author	 			0 / 0	9 days ago
	271	Added first set of icons for the application.	mercury: main	Reviewer	 			0 / 0	14 days ago
	135	Tidying up some of the code to fix some of the ed...	jplugin: main	Reviewer Moderator	 			1 / 0	14 days ago
	249	Adding description to P4Credentials & P4Credent...	jplugin: main	Moderator	 			1 / 0	14 days ago
assets									
	4	Updated the README with some context.	mercury: main	Author	 			0 / 0	29 minutes ago

Helix Core Server header bars

The dashboard for each Helix Core Server instance is displayed under the header bar for that instance. The header bars are collapsible allowing you to temporarily hide Helix Core Server instances you are not currently interested in.

Tip

- Click on the Helix Core Server name in the header to open Swarm for that instance in a new tab.

- Click on the header to the right of the Helix Core Server to collapse the dashboard for that instance. Click again to expand the dashboard for the instance .

Filtering

The global dashboard can be filtered to display only reviews from a particular Helix Core Server instance, project, authored by a particular user, or matching a role. The filter buttons filter all of the instances on the global dashboard. The filter buttons are always in view as you scroll up and down the page so that you can quickly modify the filters and see just the reviews you want to see. Click the **Reset** button to reset these filters.

Filtering options are:

- **Instances**

You can filter by the Helix Core Server instance, limiting results to **All Instances** or to an individual instance.

- **Projects**

You can filter by the project the review is part of, limiting results to **All Projects** or to an individual project. The **Project** filter will only show projects for which there are reviews in your dashboard.

- **Roles**

You can filter by your specific role in a review, limiting results to reviews for which you are the author, a reviewer, a required reviewer, or a moderator. The **Role** filter will only show roles for which there are reviews in your dashboard.

- **Authored by**

You can filter the reviews to only those that have been authored by a certain user. Type in this field to get a drop down list of users to filter by.

- **Reset** (only displayed if one or more filters are set):









Clicking the **Reset** button resets all dashboard filters back to their defaults.

- **Search**

Typing in the search field filters the reviews by description and review ID.

Review fields

The dashboard for each Helix Core Server shows a summary of the information for each review.

Author	ID	Description	Project(s)	Your Role	State	 	 	Last activity
main 								
	277	Added endpoint for initiating server shutdown.	mercury: dev	Reviewer			0 / 0	1 hour ago

Reviews that appear here are those which are waiting for action from you. The information presented should help you prioritize what to work on next.

- **Author**

The author of this review.

- **ID**

The ID of this review. Click on this to go to the review page.

- **Description**

The review description. It may be truncated if it is too long, in which case click on the description to expand it.

- **Project(s)**

List of project branches this review covers. A review may span multiple branches and projects. Click on one of them to navigate to the project page for that branch.

- **Your role**

The reason this review is in your dashboard. This can be Author, Reviewer, Required Reviewer, or Moderator.

- **State**

The current status of the review. This can be Needs Review, Needs Revision, or Approved.

Note

The Approved state only applies to review authors, it is only shown for a review that has been approved but has not been committed.

- **Type**

The type of review. This can be Pre-commit or Post-commit.

- **Votes**

The double column of votes displays the number of up votes and down votes for the review.

- **Last activity**

The last time that any changes were made to the review, including votes, comments, commits, and file changes.

Navigating directly to a specific Helix Core Server instance in Swarm

If you want to quickly visit a specific Helix Core Server instance in Swarm without going via the global dashboard, include the server name in the URL, for example:

`https://swarm.company.com/serverA.`

Once you are viewing the Helix Core Server in Swarm, Swarm works as a standard single Helix Core Server Swarm system.

Tip

- To browse jobs on **serverA**, navigate to:
`https://swarm.company.com/serverA/jobs`
- To browse reviews on **serverB**, navigate to:
`https://swarm.company.com/serverB/reviews`
- To view the dashboard for **serverB**, navigate to
`https://swarm.company.com/serverB/#actionable-reviews`

Activity

Tip

Activity streams are also shown on the [User](#), [Group](#), and [Project](#) pages. These activity pages show the activity for the user, project, or group you are viewing.


An activity stream displays a list of events that have occurred recently in the associated Helix Core Server.


For example, an event is added to the activity stream whenever:

- Changelists are checked in
- Jobs are created
- Code reviews are updated
- Comments posted
- Projects are created or edited
- Workflows are created, edited, or deleted
- Tests are created, edited, deleted, pass, or fail

To view the main activity stream:


1. Log in to Swarm.
2. Click **Activity** in the Swarm menu.


All Activity ▾ Reviews Commits Comments Jobs 




paula.boyle updated description of [review 135 \(revision 2\)](#) for [main](#)
Tidying up some of the code to fix some of the edge conditions. 2 minutes ago


Also putting some support for future features in place. #review [@@trriage](#)

 [Add a comment](#)




claire.brevia voted up [review 185 \(revision 1\)](#) for [main](#)
These browsers are no longer supported by the product and the code
should be removed. 4 minutes ago


 [Add a comment](#)



alex.randolph cleared their vote on [review 135 \(revision 2\)](#) for [main](#)
Tidying up some of the code to fix some edge conditions. 6 minutes ago


Also putting some support for future features in place. #review [@@trriage](#)

 [Add a comment](#)



alex.randolph updated description of [review 135](#) for [main](#)
Tidying up some of the code to fix some edge conditions. 13 days ago

Also putting some support for future features in place. #review [@@trriage](#)

 [Add a comment](#)

Activity dropdown menu (main activity stream only):

- **All Activity:** Select to view all activity on the Helix Core Server
- **Followed Activity:** Select to view just the activity of the projects and users that they are following

Activity streams can be filtered to only display events related to **Reviews**, **Commits** (changes), **Comments**, or **Jobs**. Click a filter label to enable that filter. Click the label again to disable the filter, or click a different filter to change filters. When active, each filter label displays a distinct background color that matches the color stripe on the right side of each event in the activity stream, to help quickly identify each event's type.




(RSS) button: Click to subscribe to an RSS (Really Simply Syndication) feed for a particular activity stream so that events can be monitored in your favorite feed reader.

Events within an activity stream contain links to their resources. Click the links to visit users, changelists, projects, comments, and more.

Each activity stream starts with as many as 50 events. As you scroll down the page, Swarm fetches additional events in batches of 50 until the stream is exhausted. For a long-running server, or a server with significant activity, it could take quite a while to receive all of the events.




Files

Helix Core Server's primary task is to version files, so Swarm makes it easy to browse the depot. Start browsing by clicking **Files** in the menu.



- Swarm displays a list of breadcrumb links to help you quickly navigate to higher level directories quickly.

[//](#) [depot](#) / [Jam](#) / [MAIN](#) / [src](#) / [Jam.html #2](#)

- Links with folder icons represent directories of files within the depot. Click a directory link  to display the contents of that directory.
- Click the .. link with the up-arrow icon  .., when it appears, to navigate to the current directory's parent.
- Links with dog-eared page icons  represent individual files within the depot. See ["File display" on the facing page](#) for more information.
- Click the **Commits** tab to display the list of changes made to files in the current directory, or any directories it contains. See ["Commits" on page 332](#) for more information:



Swarm creates links for files and directories wherever they appear in the Swarm UI.

Note


Directories that start with a period, for example `.foo`, are sorted to appear at the end of the list of directories. The `..` link, when it appears, always appears first.

Similarly, files that start with a period, for example `.htaccess`, are sorted to appear at the end of the list of files.

Browsing deleted files and folders

When the **Show Deleted Files** button is clicked, Swarm toggles the inclusion of deleted folders and files in the file display.

Deleted folders and files are presented slightly muted compared to non-deleted entries:



File display

When Swarm is asked to display a file, if the file is a type that Swarm can display, Swarm presents the file's contents.

Buttons: the actual buttons displayed depends on the file type you are viewing.

- Click the **Shorten URL** button to display a short-link to the file.
- Click the **Blame** button to add a column to the display that identifies the userid responsible for each line of the file.
- Click the **Open** button to display the file content with no surrounding page markup.
- Click the **Edit** button to edit the file, see ["File edit" on page 329](#).
- Click the **Download .zip** button to download a compressed version of the file, see ["Download files as a ZIP archive" on page 330](#).
- Click the **Download** button to download the file.

When a file is opened it will by default be truncated to 1 MB in size. This limit can be increased (or removed) via the Swarm configurables. See [Files configuration](#) for details. This limit does not apply to downloaded files.

Along with the file name, Swarm displays the version number for the currently displayed file in the breadcrumb. For example, the following breadcrumb indicates that version 2 of **Jam.html** is being displayed:

[// depot](#) / [Jam](#) / [MAIN](#) / [src](#) / [Jam.html #2](#)

If the version of a file being previewed has been deleted, the version number is displayed in red.

Every version of a file is available on the [Commits](#) tab.

Text files

Swarm displays the contents of text files (including the Helix Core Server filetypes unicode and UTF16) with line numbers. When possible, syntax highlighting is applied to make identification of various elements within the file easier.

For more information on the supported extensions for syntax highlighting in the Review page, see ["Supported syntax highlighting in the Review page" on page 331](#).

For more information on Helix Core Server filetypes, see [File Types](#) in *Helix Core Command-Line (P4) Reference*.

Tip

Swarm can be configured to display some non-UTF8 characters, see ["Non-UTF8 character display" on page 630](#) for details.

```
// helix-swarm / development / module / Application / src / Application / Router / Regex.php#2

1. <?php
2. /**
3.  * Perforce Swarm
4.  *
5.  * @copyright  2012 Perforce Software. All rights reserved.
6.  * @license    Please see LICENSE.txt in top-level folder of this distribution.
7.  * @version    <release>/<patch>
8.  */
9.
10. namespace Application\Router;
11.
```


Click **Blame** to add a column to the display that identifies the userid responsible for each line of the file:

```
// helix-swarm / development / module / Application / src / Application / Router / Regex.php#2

1. allison.c... <?php
2. allison.c... /**
3. allison.c...  * Perforce Swarm
4. allison.c...  *
5. allison.c...  * @copyright  2012 Perforce Software. All rights reserved.
6. allison.c...  * @license    Please see LICENSE.txt in top-level folder of this distribution.
7. allison.c...  * @version    <release>/<patch>
8. allison.c...  */
9. allison.c...
10. allison.c... namespace Application\Router;
11. allison.c...
12. allison.c... class Regex extends \Zend\Mvc\Router\Http\Regex
```

Each userid presented is a link that, when clicked, displays the changelist that provided the associated text. Muted userids indicate that the associated text is from the same changelist as the line above. For example, the userid *allison.c* is responsible for lines 1, 10, and 12 in the screenshot above.

When you hover your pointer over a userid in the blame column, a tooltip appears displaying the associated changelist description:



When there are no lines displayed, for example when you are viewing empty or shelved files, the **Blame** button is disabled.

Markdown files

Swarm displays Markdown files in two tabs:


- **Markdown:** by default, Markdown content is displayed but it is limited to prevent the execution of raw HTML and JavaScript content. This can be configured by the Swarm administrator, see ["Markdown" on page 639](#). This tab is only displayed for Markdown files.
- **View:** displays the contents of Markdown files as plain text in the same way as ["Text files" on the previous page](#) are displayed.

Tip

Valid Markdown file extensions are: **md**, **markdown**, **mdown**, **mkdn**, **mkd**, **mdwn**, **mdtxt**, **mdtext**.

Markdown tab

Click the **Markdown** tab to view the file content in Markdown:



View tab

Click the **View** tab to view the Markdown file in plain text:



Images

Swarm displays web-safe images:



The checkerboard background in this example is not part of the logo; it helps identify where transparency exists.

Many browsers can display SVG images with no additional plugins, so Swarm attempts to display SVG images rather than displaying the image's definition. When you use a browser that cannot natively display SVG images, you see the broken image icon.

When imagick (an optional module that integrates ImageMagick into PHP) is installed, Swarm can also display the following image formats: BMP, EPS, PSD, TGA, TIFF.

3D models

Important

You can view the 3D model file types only in the classic view of the Swarm review page. Use [Helix DAM](#), a digital asset management software solution by Perforce to view 3D model files. For more information about how to use Helix DAM, see [Helix DAM User Guide](#).

Swarm includes support for displaying select 3D model file types in the browser:



Supported file types include:

- **DAE** - including any referenced web-safe texture images.
- **STL** - both binary and ASCII versions of the format.
- **OBJ** - including any referenced MTL files, and web-safe texture images.

When Swarm can display a 3D model, it renders a generic grid *stage* and places the model in the center, scaled to make viewing straightforward. A toggle control appears in the top right: when enabled, you can control the view with the mouse, and when not enabled, permits auto-rotation to occur (when possible).

1. Click and hold the **left mouse button** to begin rotating the view. Drag while holding the left mouse button to rotate the view.
2. Click and hold the **right mouse button** to begin panning the view. Drag while holding the right mouse button to pan the view.
3. Roll the **mouse wheel** up or down to adjust the magnification of the view.

When possible, a second control appears allowing you to toggle between showing the model with surfaces, or just showing the model's wireframe.

Note

For systems with hardware acceleration, if your browser supports WebGL and hardware acceleration is enabled, Swarm renders the model and enables auto-rotation.

For systems without hardware acceleration or WebGL, but your browser supports HTML5 canvas elements and JavaScript TypedArrays, Swarm renders the model but auto-rotation is disabled. Rendering is likely to be slow and rendering quality is likely to be low.

For browsers without HTML5 canvas elements and JavaScript TypedArrays, no rendering is attempted; instead, users see a message indicating that the browser is not supported.

Other file types

It is possible to view other file types in Swarm, through the addition of additional modules, or by installing "[LibreOffice](#)" on [page 560](#) on the Swarm host.

When the file is a type that Swarm cannot display, Swarm presents the file's [history](#), along with the **Download** button:

// helix-swarm / development / collateral / build-utils / ant-contrib.jar #1				
Commits			Download .zip	
			Download (219 KB)	
#	Change	User	Description	Committed
#1	541195	steve.russell	Initial set of build infrastructure: * Ant build.xml to drive it all * Ant Contrib for s...	6 years ago

File edit

Important

File edit is a Technology Preview feature.

Features offered in Technology Preview are experimental and not guaranteed to always work as expected. If you have feedback and functionality suggestions, email techpreview@perforce.com.

Note

Requirements:

- You must have permissions to edit the file.
- If the file version changes while you are editing it or if you are not editing the Head revision of the file, committing your file changes will overwrite the current head revision of the file.
If you are unsure, shelving your change is a safer option because it will not overwrite the head revision of the file and you can do a manual resolve when you commit the shelf.
- File edit from the **Files** browser page is enabled by default but can be disabled by your Swarm administrator, see "[allow_edits](#)" on [page 623](#).

You can edit a file from the Swarm **Files** browser page and shelve or commit your file changes.

To edit a file from the Swarm Files page:

1. Navigate to the file on the **Files** page.
2. Click the **Edit** button above the file to open the file editor.

Edit file

projects / acme / main / README.md#1

1

ACME

2

3

This is the ACME project. It is a nice project.

4

5

Other Things

6

7

Nothing much here.

8

9

A nice new line has been added here.

Changelist description

Added new line of text at line 9.
#review

Discard changes

Commit

Shelve

3. Edit the file and add a **Changelist description**.

Tip

Optional: you can create a review for your change or add it to an existing review by including a review keyword in the **Change description**. For instruction on creating a review and adding a changelist to a review using review keywords, see ["Create a review" on page 682](#) and ["Add a changelist to a review" on page 682](#).

4. Click **Commit** to commit the file or **Shelve** to shelve the file.

To discard your changes, click **Discard changes** and confirm that you want to discard your changes when requested.

Download files as a ZIP archive

The **Download .zip** button is used to download the selected files or folders in the Helix Core Server as a ZIP archive. This makes it easy to get a copy of files without having to setup a client.

Note

The **Download zip** option is not displayed if the zip command-line tool is not installed on the Swarm server. For information about installing, and configuring the zip command-line tool, see [Zip archive](#).

When you select the **Download zip** option, Swarm performs the following steps:

1. Scans the files/folders:
 - Checks that you have permission to access their contents, according to the Helix Core Server protections.
 - Checks that the total file size is small enough to be processed by Swarm.
2. Syncs the file contents to the Swarm server from the Helix Core Server.
3. Creates the ZIP archive by compressing the file content.
4. Starts a download of the generated ZIP archive.

Note

- You might not see all of the above steps; Swarm caches the resulting ZIP archives so that repeated requests to the same files/folders can skip the sync and compress steps whenever possible.
- If an error occurs while scanning, syncing, or compressing, Swarm indicates the error.

Supported syntax highlighting in the Review page

This section outlines a list of all the supported extensions for syntax highlighting in the Review page.

File type	File extension
c	c, h
java	java
python	py
cpp	cc, cpp, hpp
csharp	cs, dotnet
actionscript for flash	fla, as
batch	bat
powershell	ps1, psm1, ps, psm
yaml	yaml, yml

File type	File extension
markup	xml, htm, html, phtml, md, markdown
visualBasic	vb
javascript	js, javascript
jsx	jsx
php	php
matlab	mat
lua	lua
vhd1	vhd1
rust	rust, rs, rlib
docker	dockerfile
shell	sh
ruby	rb
perl	pl
typescript	ts
tcl	tcl
unreadsript	uc

Commits

Whenever a new version of a file is checked into the Helix Core Server, a commit record is created. Begin browsing the history of commits by clicking **Commits** in the menu.

When you are viewing a particular file or directory, clicking the **Commits** tab displays the commits for that location in the depot.

HelixSwarm

Commits

// projects / mercury / dev

Browse Commits

Range User

Change	User	Description	Committed
276	Paula Boyle	Added endpoint for initiating server shutdown. #review-277	2 years ago
263	Steve Russell	Add functionality to report on the current status of the server. This is so the client ca ...	2 years ago
239	Steve Russell	Adding in basic client api for getting client information. #review-240	3 years ago

Range filter

The **Range** field lets you filter the list of changes for the depot path being viewed. When you click the **Range** field, a dropdown syntax guide appears providing sample commit filtering expressions.

Range User

Examples

@0,@now	All Time (Default)
@80,@90	Between changes 80 and 90
@80	Up to change 80
@=80	Just change 80
@>80	Newer than change 80
@80,@now	Change 80 to Now
@label-name	Part of label 'label-name'
@2014/11/30	Up to Nov 30, 2014
@2014/11/30,@now	Nov 30, 2014 to Now

The expressions that can be used within the **Range** field include:

- @0,@now (the default): displays all commits for the current depot path.
- @80,@90: displays all changes between 80 and 90.


- `@80`: displays all changes up to change 80.
- `@=80`: displays only change 80. Change 80 might not involve the current depot path, so there may be no commits to display.
- `@>80`: displays all changes newer than change 80.
- `@80,@now`: displays all changes from change 80 to now.
- `@label-name`: any changes represented by label *label-name*.
- `@2014/11/30`: displays all changes up to November 30, 2014.
- `@2014/11/30,@now`: displays all changes from November 30, 2014 to now.

File Commits

A file's commit history presents each version of a file that the Helix Core Server knows about, including the change number, userid, change description, time ago, along with **Open** and **Download** buttons.

// projects / mercury / candidate / README.md #1					
View		Commits		Open Download (57 B)	
#	Change	User	Description	Committed	
#1	238	steve.russell	Branched the main line to both a new candidate and dev branch.	5 months ago	Open
▼ //projects/mercury/main/README.md					
#1	5	jack.boone	Moved project to projects depot.	7 months ago	Open
▼ //depot/projects/mercury/main/README.md					
#2	3	jack.boone	Updated format of the project README.	7 months ago	Open
#1	2	jack.boone	Added README to the Mercury project main branch.	7 months ago	Open

Swarm also displays *contributing commits* when available, such as when a file has been renamed, or integrated from another location in the Helix Core Server.

If a commit represents a deleted revision, the **Open** and **Download** links are replaced with a trashcan icon  to indicate that this version is no longer available.

Remote depot commits


When your Helix Core Server has a remote depot configured, you can browse the contents of the remote depot, but remote depots do not share their commit history. If you attempt to view the commits of a file provided by a remote depot, Swarm displays:

// builds	
Browse	Commits
Range User	
Remote depot (change details are not available).	

Jobs

Jobs are a component of Helix Core Server's defect tracking system and a record of bugs found or improvement requests. Jobs can be associated with [changelists](#) to create *fix* records, indicating the work that solved the problem or provided the requested feature. Begin browsing jobs by clicking **Jobs** in the menu.

You can search available jobs by entering a job filter in the search box. Words, phrases, and **field=value** pairs can be entered. For example, entering **reportedby=steve.russell** **swarm** displays jobs that user *Steve Russell* has reported that also contain the word *swarm* in the description.



The screenshot shows the HelixSwarm interface. On the left is a sidebar menu with options: Reviews, Projects, Files, Commits, **Jobs** (highlighted), Groups, Workflows, and Tests. The main area is titled 'Jobs' and contains a search bar with the filter 'reportedby=steve.russell swarm'. Below the search bar is a table with the following data:

Job	Status	Reported By	Description	Modified Date
job093521	Inprogress	steve.russell	Remove the bash and vbs versions of the swarm-trigger. With this change we need to be careful not t...	about 3 hours ago
job093520	Open	steve.russell	The post format dropdown on project settings can appear in the wrong spot (to the right of post bod...	about 4 hours ago

The fields you can search for depend on the jobspec defined in your Helix Core Server.

Edit the Job columns


Configure the columns that are displayed:



1. Click the **Select Columns** button beside the search field to display the dropdown menu. The dropdown menu shows all of the available jobspec fields.
2. Select the columns to display using the checkboxes in the dropdown menu.
3. Drag the column label up or down the list to change the order the columns are displayed in.
4. Click the **Select Columns** button again, or a blank portion of the page, to hide the menu.

Change the column order from the **Jobs** page:

1. Click a column heading and drag it to the left or right to move the column to a new position.
2. The column display updates as columns are dragged.



The screenshot shows the HelixSwarm interface with the 'Jobs' page. The search bar still has the filter 'reportedby=steve.russell swarm'. The table now displays an additional column, 'Owned By', which has been moved from its original position to the left of 'Description'. The data in the table is as follows:

Job	Status	Reported By	Owned By	Description	Modified Date
job093521	Inprogress	steve.russell	jack.boone	Remove the bash and vbs versions of the swarm-trigger. With this change we need to be careful not t...	about 4 hours ago
job093520	Open	steve.russell	jack.boone	The post format dropdown on project settings can appear in the wrong spot (to the right of post bod...	about 4 hours ago

For more information on customizing job specifications, see [Customizing Helix Core Server: Job Specifications](#) in *Helix Core Server Administrator Guide*.


Job display

Jobs are typically identified with the word job followed by six digits. For example, `job000123`.




View a specific job by clicking on a linked job identifier, or by visiting the URL:

`https://myswarm.url/jobs/jobid`.

When Swarm displays a job, the presentation is similar to:

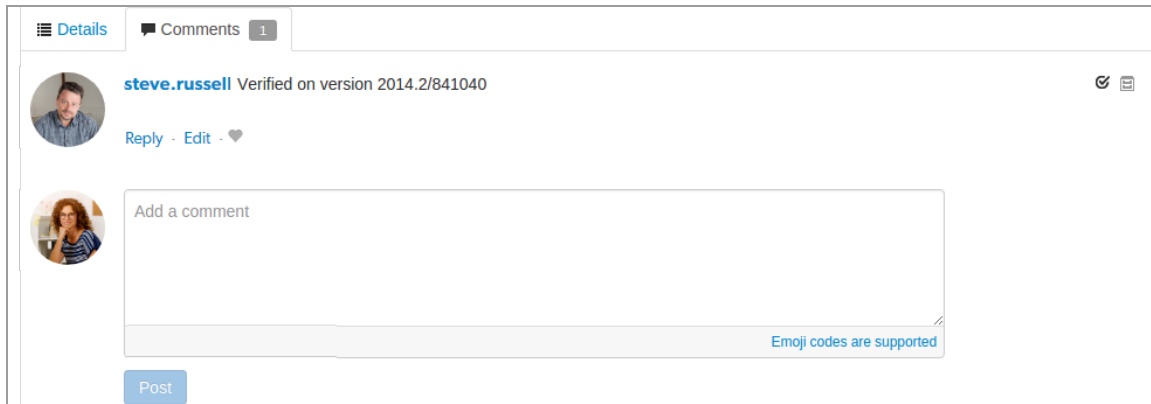


The upper portion of the job presentation includes:


- The avatar and userid of the user that created the job
- The job's [creation time](#)
- If changes have been made to the job, the modifying userid and [time](#)
- A status indicator
- The job's description
- If any changelists have been submitted that *fix* the job, a list of those changelists and their descriptions. Each associated changelist includes an icon to represent their type:  (review),  (pending),  (committed)

The lower portion of the job presentation lists all of the keys configured in your Helix Core Server's jobspec. Swarm inspects the jobspec and enhances the presentation of fields it recognizes. For example, date fields display as *time ago*, and links are created for *userids*.

Click the ["Comments" on page 348](#) tab to view any comments added to the job, or to add a comment.



Adding a comment sends a [notification](#). The **Comments** tab displays the number of open comments associated with the job. If you hover your mouse over the comment count, a tooltip is displayed showing how many comments are archived:



For more information on customizing job specifications, see [Customizing Helix Core Server: Job Specifications](#) in *Helix Core Server Administrator Guide*..


Note

The default Helix Core Server job specification contains very few fields. Adding fields to record additional information, such as the modification time and userid, reporting time and userid, can assist Swarm use appropriate terminology when describing the current disposition of a job.

Add jobs

Swarm does not provide the ability to create new jobs in the Helix Core Server, but jobs can be added to changelists or reviews.

Add a job to a changelist

1. Navigate to a changelist.
2. Click the **Add Job** link  [Add Job](#).
3. Scroll through the available jobs, or enter job search criteria to search available jobs.
For more information on job search criteria, see [Jobs](#) in *Helix Core Command-Line (P4) Guide*.
4. If you find the job you want to add, click its row to highlight it and then click **Select**. Or, double-

click the desired job to add it.

5. If you do not find the appropriate job, click **Cancel**.

Add a job to a review

1. Click **Jobs (n)** (where **n** is the number of jobs that are already part of the review).
2. Click **Add a job**.
3. Scroll through the available jobs, or enter job search criteria to search available jobs.

For more information on job search criteria, see [Jobs](#) in *Helix Core Command-Line (P4) Guide*.

4. If you find the job you want to add, click it to highlight it and then click **Select** to add it.

Note

If you attempt to add a job to a review that affects a single project, Swarm applies the project's job view filter to display only jobs that affect the project. It is not currently possible to expand the filter to include jobs outside of the project.

Unlinking jobs

Swarm does not provide the ability to delete jobs from the Helix Core Server, but jobs can be unlinked from changelists or reviews.

Unlink a job from a changelist

1. Navigate to a changelist that has an associated job.
2. Click the **X** button beside the job.
3. When prompted for confirmation, click **Unlink** to unlink the job.

Unlink a job from a review

1. Click **Jobs (n)** (where **n** is the number of jobs that are already part of the review).
2. Click **X** next to the job you want to unlink.

The job is immediately unlinked from the review.

Changelists

Changelists are the basic unit of versioning work in Helix Core Server. A changelist is a list of files, their revision numbers, and the changes made to those files. Commits is a shorter synonym used throughout Swarm.

More information is available here:


- [Browsing changelists](#)
- [Browsing a user's shelved changelists](#)
- [Swarm's internal use of changelists to facilitate code reviews](#)

Changelist display

View a specific *changelist* by clicking on a linked changelist number, or by visiting the URL:
`https://myswarm.url/changes/changelist number`.

When Swarm displays a change, the presentation is similar to:

Change 238



steve.russell committed this change 21 days ago into [//projects/mercury](#) [Request Review](#)

Branched the main line to both a new candidate and dev branch.

[Add a Comment](#)

[Add Job](#)

Files 8 Comments 0

0 edited 8 added 0 deleted

- > + candidate/README.md#1
- > + candidate/src/api/jamfile#1
- > + candidate/src/api/jamfile.api#1
- > + candidate/src/api/p4api.cc#1
- > + dev/README.md#1
- > + dev/src/api/jamfile#1
- > + dev/src/api/jamfile.api#1
- > + dev/src/api/p4api.cc#1

Tip: Use **n** and **p** to cycle through the changes.

Swarm supports stream specs in your workspace using the [Private editing of streams](#) feature. If a changelist or review contains a stream spec, it will be displayed first in **Files** with the prefix **stream: //**, for example: **stream: //MyStreamDepotName/MyStreamSpecLocationName**. A changelist/review can only contain one stream spec.

The changelist display includes:

- The avatar and userid of the user who made the change
- The time the change was made
- The common depot location containing all the files included in the change

- The **Download .zip** button, used to download a ZIP archive that contains all of the files in the changelist:

Note

The **Download zip** option is not displayed if the zip command-line tool is not installed on the Swarm server. For information about installing, and configuring the zip command-line tool, see [Zip archive](#).

When you select the **Download zip** option, Swarm performs the following steps:

1. Scans the files/folders:
 - Checks that you have permission to access their contents, according to the Helix Core Server protections.
 - Checks that the total file size is small enough to be processed by Swarm.
2. Syncs the file contents to the Swarm server from the Helix Core Server.
3. Creates the ZIP archive by compressing the file content.
4. Starts a download of the generated ZIP archive.

Note

- You might not see all of the above steps; Swarm caches the resulting ZIP archives so that repeated requests to the same files/folders can skip the sync and compress steps whenever possible.
- If an error occurs while scanning, syncing, or compressing, Swarm indicates the error.

- If the change was involved in a code review, a link to the review along with a **View Review** button.
- The description of the change
- A list of jobs that this change *fixes*, if any. You can [add jobs](#) and [unlink jobs](#) here.
- The list of files included in the change, including any folders between the common depot location and the file, and the file's version number.
- A tab to review any comments made regarding the change, or any of its files.

Each file is presented in a diff display, showing you whether the file was added, modified, or deleted. For text-based and image files, Swarm can display any changes made within the file. For changes with only a single file, the diff display is the default; otherwise each file is listed. Click the filename to see the diff display. See ["Diffs" on the next page](#) for more information.


The **Request Review** button indicates the current state of this change; no Review record has been created. Clicking **Request Review** starts a code review for this change. For more information, see ["Start a review" on page 465](#).

Important

If your Helix Core Server is configured as a commit-edge deployment, and your normal connection is to an edge server, Swarm refuses to start reviews for shelved changes that have not been promoted to the commit server.

Within Swarm, this means that the **Request Review** button does not appear for unpromoted shelved changes. Outside of Swarm, attempts to start reviews for unpromoted shelved changelists appear to do nothing. Ask your Helix Core Server administrator for assistance if you cannot start a review.

An administrator of the Helix Core Server can automatically promote shelved changes to the commit server by setting the configurable `dm.shelve.promote` to 1.


When a file in a changelist has one or more associated comments, an icon  appears near the far right of the file's entry.

Diffs

Tip
Code review page

The information on this page only applies to the classic view of the Swarm Code review page. For information about diffs on the default review page, see ["Files tab" on page 453](#). To switch between the default review page and the classic review page, use the **Classic view** toggle switch at the top of the review page.

When you view a changelist or code review, the associated files are presented as *diffs*, short for differences, showing you how they have changed.



Swarm supports stream specs in your workspace using the [Private editing of streams](#) feature. If a changelist or review contains a stream spec, it will be displayed first in **Files** with the prefix **stream: //**, for example: **stream: //MyStreamDepotName/MyStreamSpecLocationName**. A changelist/review can only contain one stream spec.

Note

If you append a changelist with a stream spec to a review that already contains a stream spec, the spec in the changelist replaces the original one in the review. If it is a different spec from the original spec in the review, Swarm cannot display the diff between them and displays **File content unchanged**.

The review version selectors are used to specify which versions of a review you want to *diff*, they are located in the **Files** tab above the list of files.

- **Left dropdown selector:** the base version for the review is selected by default, base is the revision of the file that was checked out of the depot before it was changed for this review. The current review version in the depot, sometimes called Head can be selected, if there are multiple versions of the review a specific version can be selected.
- **Right dropdown selector:** the latest version of the review is selected by default. If there are multiple versions of the review, a specific version of the review can be selected.

Tip

If a review consists of one or more Swarm-managed changelists. When comparing versions of a review, Swarm is showing any differences between the selected versions, not the review author's personal changelist. See "[Internal representation](#)" on page 424 for details.

The first row of buttons above the files allow you to (left to right):






- **Show Comments** button: toggles the display of comments to inline in files or only in the **Comments** tab.
- **Show Diffs In-Line** button: displays all diffs as inline.
- **Show Diffs Side-by-Side** button: displays all diffs as side-by-side.
- **Toggle Show Whitespace** button: toggles the display of whitespace characters (such as space, tab, and newline) for all files.
- **Toggle Ignore Whitespace** button: toggles the highlighting of whitespace changes in all of the file diffs.
 - **Highlight whitespace changes:** makes it easier to identify changes in file types where whitespace is important. This is the default value.
 - **Ignore whitespace changes:** whitespace changes are not highlighted, this makes it easier to see the important changes in file types where whitespace changes are not important.
- **Collapse All** button: collapses all files
- **Expand All** button: expands all files. By default this button is disabled if there are more than 10 files in the review. For details, see [Expand All Limit](#).

Tip

The default states for the **Show Comments**, **Show Diffs Side-by-Side**, **Toggle Show Whitespace**, and **Toggle Ignore Whitespace** buttons are set in your user settings, see [User Settings](#).

Each file is presented with an icon indicating whether the file was:

-  Added/Branched/Imported
-  Edited/Integrated
-  Deleted

The file's presentation can be controlled with (left to right):

Tip


Only the **Show Full Context** button is available for stream specs.



- **Comments in File** icon (only displayed if the file contains comments): indicates that the file contains comments.
- **Show In-line** button: highlights line additions, modifications, and removals in a single pane.
- **Show Side-by-Side** button: highlights additions, modifications, and removals in two panes, with the older revisions of the file on the left, and the newer revision on the right.
- **Show Whitespace** button: makes whitespace characters more visible; spaces show up as dots, tabs show up as arrows that point to a bar, and line endings show up as down-pointing arrows.
- **Ignore Whitespace** button: toggles the highlighting of whitespace changes in a file, making it easier to identify non-formatting changes. Normally, whitespace is not ignored.
- **Toggle Ignore Whitespace** button: toggles the highlighting of whitespace changes in the file diff.
 - **Highlight whitespace changes:** makes it easier to identify changes in file types where whitespace is important. This is the default value.
 - **Ignore whitespace changes:** whitespace changes are not highlighted, this makes it easier to see the important changes in file types where whitespace changes are not important.
- **Show all diffs** button (edited and integrated files only): toggles between displaying all the diffs for the file and only the first few. The limit of what are shown by default is configurable by the administrator (see [Max Diffs](#)) and defaults to 1500. If there are fewer than that then this button is hidden.
- **Show Full Context** button (edited and integrated files only): toggles between displaying only the portions of the file that have changed and the full file.
- **Open File** button (edited or integrated files only): opens a new browser tab/window display the full file (where possible), provide access to its history, and a button to download the file.

- **Mark file as read** button (displayed for code reviews only): helps you (and others) keep track of which files have been reviewed. This is particularly useful when a code review consists of many files.

When clicked, the button color inverts and the associated file is visually muted, to make it easy to distinguish read files from unread files:



If a file has been marked as read, click the button a second time to reset the status to unread.

Viewing a diff

When you view a diff, the changes are highlighted:

- Red indicates lines that have been removed.
- Blue indicates lines that have been modified.
- Green indicates lines that have been added.

151		151	
152	# Initialize variables	152	# Initialize variables
153	#	153	#
154	# "default =" - set only if unset	154	# "default =" - set only if unset
155		155	
156	OSFULL = \$(OS)\$(OSPLAT)\$(OSVER) \$(OS)\$(OSPLAT) \$(OS)	156	OSFULL = \$(OS)\$(OSPLAT)\$(OSVER) \$(OS)\$(OSPLAT) \$(OS)
157		157	
158	#	158	#
159	# OS specific variable settings	159	# OS specific variable settings
160	#	160	#
161		161	
162	switch \$(OS)	162	switch \$(OS)
163	{	163	{
164	case AIX : LINKLIBS default = -lbsd ;	164	case AIX : LINKLIBS default = -lbsd ;
165	case DGUX : RANLIB default = "" ; RELOCATE =	165	case DGUX : RANLIB default = "" ; RELOCATE =
166	case HPUX : RANLIB default = "" ;	166	case HPUX : RANLIB default = "" ;
167	INSTALL default = "" ;	167	INSTALL default = "" ;
168	case IRIX : RANLIB default = "" ;	168	case IRIX : RANLIB default = "" ;
169	INSTALL default = "" ;	169	INSTALL default = "" ;
170	case MVS : RANLIB default = "" ; RELOCATE =	170	case MVS : RANLIB default = "" ; RELOCATE =
171	case NEXT : AR default = libtool -o ;	171	case NEXT : AR default = libtool -o ;
172	RANLIB default = "" ;	172	RANLIB default = "" ;
173	case NCR : RANLIB default = "" ;	173	case NCR : RANLIB default = "" ;
174	INSTALL default = "" ;	174	INSTALL default = "" ;
175	case PTX : RANLIB default = "" ;	175	case PTX : RANLIB default = "" ;
176	case QNX : INSTALL default = "" ;	176	case QNX : INSTALL default = "" ;
177	case SCO : RANLIB default = "" ;	177	case SCO : RANLIB default = "" ;
210	CCFLAGS default = -nosyspath	209	CCFLAGS default = -nosyspath
211	C++ default = \$(CC) ;	210	C++ default = \$(CC) ;
212	C++FLAGS default = -nosyspath ;	211	C++FLAGS default = -nosyspath ;
213	FORTTRAN default = "" ;	212	FORTTRAN default = "" ;
214	LIBDIR default = /boot/develop	213	LIBDIR default = /boot/develop
215	LINK default = mwld ;	214	LINK default = \$(CC) ;
216	LINKFLAGS default = "" ;	215	LEX default = "" ;
217	LEX default = "" ;	216	MANDIR default = /boot/develop
218	MANDIR default = /boot/develop	217	NOARSCAN default = true ;
219	NOARSCAN default = true ;	218	RANLIB default = "" ;
220	RANLIB default = "" ;	219	STDHDRS default = /boot/develop
221	STDHDRS default = /boot/develop		
1884	actions quietly updated piecemeal together RmTe		
1885	{		
1886	{		
1887	\$(RM) \$(>)		
1888	}		
1889			
1890	actions Shell		
1891	{		
1892	copy \$(>) \$(<)		
1893	}		
1894	}		
1895		1769	
1896		1770	
1897	#	1771	#
1898	# Backwards compatibility with jam 1, where rules w	1772	# Backwards compatibility with jam 1, where rules w
1899	#	1773	#

Show more context buttons








The diff presentation displays a concise view of where changes were made within a file, showing the changed lines and only a few lines before and after the change. Sometimes, this concise view needs to be expanded to fully understand the context of the change, use the **Show More** and **Show All** buttons to display extra lines around the change:

Note

By default, the number of extra lines displayed when you click the **Show More Lines Above** and **Show More Lines Below** buttons is 10. This setting can be changed by your Swarm administrator, see ["More context lines" on page 692](#).


Tip

The following buttons are not available for stream specs.


- **Show the file list panel**  button (only displayed on the vertical bar next to the file list panel): click to view the file list panel.
- **Show the information panel**  button (only displayed at the bottom of the information panel): click to view the detailed information panel.
- **Show All Lines to Start of File**  button (only displayed for the first change in the file): click to show all of the lines up to the start of the file.
- **Show More Lines for the Code Below**  button: click to show 10 more lines above the change, the extra lines are displayed in the pane below the button.
- **Show Entire Section**  button: click to show all of the lines between the changes that are above and below the button, the two changes and the lines between them are displayed in a single pane.
- **Show More Lines for the Code Above**  button: click to show 10 more lines below the change, the extra lines are displayed in the pane above the button.
- **Show All Lines to End of File**  button (only displayed for the last change in the file): click to show all of the lines down to the end of the file.

Comments

Comments can be displayed within the body of a file and appear immediately below the line the commenter targeted for comment. See ["Comments" on page 348](#) for more details.




Note

The *comment here* icon  appears in the line number column whenever there is a comment. Click the icon to display the comment and click the icon again to collapse the comment. This is useful when the comment display is toggled off.

Inline diff view

When viewing a diff in-line, the line numbers for the old version are first, the line numbers for the new version are second, and they are followed by the file content. Some users find it easier to use this view to locate an area that has changed, but they then switch to the side-by-side view to understand the change better. Swarm maintains the scroll position, you do not lose your place in the file after toggling the diff view.




Note

Press **n** on your keyboard to scroll to the next changed area within a file. Press **p** to scroll to the previous change.

Per-file toolbar

When a diff contains multiple files, the changes can be taller than your browser window. Swarm keeps the *per-file toolbar* in view for each file as you scroll, this means that you can identify the file and control its presentation with the buttons on the right of the toolbar.



Comments

Comments are the primary feedback mechanism provided by Swarm. Review comments can be flagged as *tasks* for a lightweight workflow within a review that helps authors and reviewers prioritize review feedback, see ["Tasks" on page 351](#) for details. By default, comment notifications are delayed to allow you to add or edit comments as you progress through a review without sending a notification for each individual comment on the review. Comment notifications are rolled up into a single notification that you can either leave to be sent automatically, or you can send manually, see ["Comment notification delay" on page 355](#).

You can also like comments, add links, add attachments, and add Emojis, see ["Comment features" on page 355](#) for details.

Tip

- Markdown content is displayed in review comments, but Markdown support is limited to prevent execution of raw HTML and JavaScript content. For information about Markdown, see ["Markdown in comments and review descriptions" on page 420](#)
- If you use Markdown styles in your review comment, Swarm renders them when you post the comment.

You can add comments to:

- [A changelist, review, or job](#)
- [A review description](#)
- [A changelist description](#)
- [A line of a text file in a changelist, or review](#)
- [A file in a changelist, or review](#)

Access comments for a changelist, review, or job, by clicking the **Comments** tab. The number of open (non-archived comments) is displayed in the tab. Hover your mouse pointer over the comment count, the tooltip shows how many comments are archived. See "[Archiving comments](#)" on page 362 for details.

Adding comments

This section describes how to add comments.

Related comment features:

- [Comment notification delay](#)
- [Reply to comment](#)
- [Emoji](#)
- [Links](#)
- [Attachments](#)
- [Markdown](#)

Tip

You can use "[Links in descriptions and comments](#)" on page 409 in comments. An @mention includes the specified user in the review, and they will receive a notification whenever there is an update to the review.

Commenting on a changelist, review or job

1. From the changelist, review, or job page: click **Comments** to view the **Comments** tab.
2. Add your comment in the text area.
3. Click **Post**.

Commenting on a review description

1. From the changelist or review page description area: click **Comments (n)** (where **n** is the number of comments that already exist).
2. Click **Add a comment**.
3. Add your comment in the text area.
4. Click **Post**.

Tip

To hide the description comments, click **Comments n** (where **n** is the number of description comments that exist). To display the comments again, click **Comments n**.

Commenting on a changelist description

1. From the changelist or review page description area: click **Add a Comment** or **n Comments** (where **n** is the number of comments that already exist).
2. Add your comment in the text area.
3. Click **Post**.

Tip

To hide the description comments, click **n Comments** (where **n** is the number of description comments that exist). To display the comments again, click **n Comments**.

Commenting on a specific line in a file in a changelist or review

1. From the changelist or review page: click **Files** to view the **Files** tab.
2. Click on the line you want to comment on.
3. Add your comment in the text area.
4. **Optional** (Review page only): Select the **Flag as Task** checkbox to mark the comment as a task that needs to be addressed.
5. Click **Post**.

Commenting on a file in a changelist or review

1. From the changelist or review page: click **Files** to view the **Files** tab.
2. If there are multiple files, click the file you want to comment on to expand its view.
3. Click the **Add a comment** link in the footer of the file display.
4. Add your comment in the text area.
5. **Optional** (Review page only): Select the **Flag as Task** checkbox to mark the comment as a task that needs to be addressed.
6. Click **Post**.

Editing comments

This section describes how to edit comments.


Related comment features:

- [Comment notification delay](#)
- [Reply to comment](#)
- [Emoji](#)
- [Links](#)

- [Attachments](#)
- [Markdown](#)

Note

- You can only edit comments that you have created.
- If a comment is edited, all of the likes for that comment are removed.

1. Click the comment **Edit comment**  link.
2. Edit the comment text.
3. Add attachments or remove attachments. For instructions on adding and removing attachments, see "[Comment attachments](#)" on page 358.

Tip

If you remove the wrong file attachment but have not saved the comment yet, click **Cancel**, and the file will remain attached to the comment.

4. Click **Save** to save the edited comment.

Swarm sends a notification to everyone involved in the review, including the review author and the reviewers, but not the editor of the comment.

The comment timestamp is marked as (*edited*) to show that the comment has been changed.

Note

Swarm does not provide a mechanism to see older versions of edited comments.

Tasks


Flagging review comments as tasks is a lightweight workflow within a review that helps authors and reviewers prioritize review feedback. Any comment on a code review can be flagged as a task, indicating to the code review's author that the described issue needs to be addressed, and that the review is unlikely to be approved without a fix.


Note

Changelist and Job comments cannot be flagged as tasks.

Flag a comment as a task

To flag a comment as a task, select the **Flag as Task** checkbox when posting a comment, or click the

Comment actions  button in the upper right of an existing comment and select **Flag as Open task** in the drop-down menu.




Note

If you do not have permission to archive comments, you do not have permission to flag comments as tasks. Anonymous users never have permission to archive comments, and can only view current task states.

Set a task to Task addressed or Not a task


Once a comment is flagged as a task, it is considered to be an *open task*. Click the **Red flag** button to display a drop-down menu with the following options:



- **Task addressed:** usually used by the author of the review to indicate that the issue has been fixed.
- **Remove task:** used to correct comments that have been flagged as tasks by mistake.

Verify a task, verify and archive a task, or reopen a task


A comment with a green check indicates that the task has been addressed. Click the **Green check mark** button to display a drop-down menu with the following options:



- **Verify task:** usually used by the author of the comment, or another reviewer, after confirming that the issue is fixed.
- **Verify and Archive** (only supported in the original Swarm review page): used to both indicate that the issue has been fixed, and to archive the comment so that it is hidden from view. Archived tasks, whether they are open, addressed, or verified, are not included in the task counts for the code review.
- **Reopen task:** used if the issue needs further work after it has been marked as addressed.

Reopen a task

A comment with a blue double-check indicates that the task has been verified. Click the **Blue double-check mark** button to display a drop-down menu with the following option:



- **Reopen task:** used if the issue needs further work post-verification, or if verification was made by mistake.

Task details


A summary of the number and status of comments flagged as tasks is displayed in the **Information panel** to the right of the review.

Note

Archived comments that are flagged as tasks are not included in the summary or the **Tasks** dialog.



- **Red Flag:** Displays the number of open tasks on the review.
- **Green check mark:** Displays the number of addressed tasks on the review.
- **Blue double-check mark:** Displays the number of addressed and verified tasks on the review.

Show Task details  : Click to display a dialog listing all of the tasks associated with the review:

Tasks

Reporter

Status		Reporter	Description
Addressed	✓	Allison Clayborne	I have passed this to QA.
Addressed	✓	Allison Clayborne	We will need to update this for the next revision.
Open	🚩	Claire Brevia	This description needs more detail. I feel the lack of detail makes it very hard to review this work. When will you get some time to change this?
Verified	✓✓	Claire Brevia	This makes no sense, please update

Within the **Tasks** dialog, you can filter the tasks by the **Reporter** (the userid of the user who created the task), and/or by task state using the buttons at the top of the dialog:

- Click the **Red flag** button to display only open tasks (comments that need to be addressed).
- Click the **Green check mark** button to display only addressed tasks (comments that have been addressed).
- Click the **Blue double-check mark** button to display only verified tasks (comments that have been addressed and verified).

To change the state of a task from the **Tasks** dialog, click the task state dropdown for the task and select the new task state.

To view the full comment text for a task, click the ellipses to the right of the task.

Approve a review with open tasks


Flagging a comment as a task provides a visual indication that there is an identified issue that needs to be addressed before the review can be approved, see ["Set a task to Task addressed or Not a task" on page 352](#).

Note

If Swarm is configured to prevent approval of reviews with open tasks and a review has open tasks, the **Approve**, and **Approve and commit** options will not be available for the review. This option is configured by an administrator. See ["Disable approve for reviews with open tasks" on page 688](#).

To approve, or approve and commit a review with open tasks, you must address the tasks first and then set them to **Task Addressed**, or **Not a Task**. See ["Set a task to Task addressed or Not a task" on page 352](#) for details.

If you select **Approve** or **Approve and Commit** for a review that has open tasks, a warning message is displayed in the **Update Review** dialog. The warning is only advisory, either click **Cancel** and address the open tasks or click **Approve** to approve the review. Archived open tasks will not trigger the warning message.



The image shows a dialog box titled "Update Review" with a close button (X) in the top right corner. Inside the dialog, there is a yellow warning banner that reads "Warning! There are 2 open tasks on this review." Below the banner is a large text area with the placeholder text "Optionally, provide a comment". At the bottom right of the dialog, there are two buttons: "Approve" (in blue) and "Cancel" (in white with a grey border).

Comment features

Comment notification delay

By default, comment notifications are delayed to allow reviewers to add or edit comments as they progress through a review without sending a notification for each individual comment on the review. Comment notifications are rolled up into a single notification and sent either manually by the reviewer, or automatically after the notification delay time has been exceeded.

The delay countdown is reset each time the reviewer adds or edits a comment on the review, by default the notification delay time is set to 30 minutes.

- If you are commenting on more than one review, each of the reviews that you are commenting on has its own notification delay countdown that only applies to the comments that *you* make on *that* review.
- If another reviewer is making comments on the same review as you, *that* reviewer has their own notification delay timer for *that* review.
- If you manually send a delayed comment notification, the notification will only contain the comments that *you* made on *that* review.

Tip

The comment notification delay does not delay the posting of the comments, only the comment notification is delayed.

Note

Comment notifications are only delayed for comments on reviews. Comments on commits or jobs produce notifications immediately.

Note

The notification delay time is a global configuration setting configured by the Swarm administrator, see [Comment notification delay](#).

Manually send the comment notification immediately:

1. Add or edit your comment as normal.
2. Click the **Post and notify (n)** button to the right of the comment box.
Where **(n)** is the number of delayed comment notifications in the queue waiting to be sent, this number does not include the current comment you are working on.

Tip

- If you forget to click the **Post and notify (n)** button, click the **Send All Notifications (n)** button below the review description to send all of your notification for the review immediately.
- Only the comments that you have made on this review are rolled up into the notification that is sent when you click **Post and notify (n)** or **Send All Notifications (n)**.

Reply to comments

By default, you can reply to comments, replies are displayed in a thread below the parent comment. Comment thread depth is set to 4 by default, this means you can have up to 4 levels of replies for a parent comment. The **Reply** link is not displayed for replies at or above the maximum thread depth set for Swarm. If the parent comment is archived, replies are archived with the parent, see ["Archiving comments" on page 362](#).


Note

Comment replies can be disabled, and the thread depth can be increased or reduced by a Swarm administrator, see ["Comment threading" on page 586](#).

Tip


If the thread depth is reduced by a Swarm administrator, earlier replies at a deeper level will continue to be displayed but you cannot reply to them.


To reply to a comment:

1. Click **Reply**  below the comment you are replying to.
2. Add your comment in the text area.
3. Click **Post**.

Files **6**
Comments **4**
Activity
Send All Notifications (0)


5 archived comments


Steve Russell commented on review 644 (version 1) //projects/mercury/dev/src/api/db.c, line 51 - 11 months ago
✓ read ✓



Allison Clayborne commented on review 644 (version 1) //projects/mercury/dev/src/api/db.c, line 105 - 10 months ago
👁️ ...

```
+ */
+int TalisAliquam(long quandoScelerisqueSem, float count, double litora) {
+ // Peristo velit parco.
+ parco = volutpat / 1;
+ printf("Aquam arcu vehicula faucibus umbra magna.");
```

@steve.russell what do you think of this?
👍


Steve Russell commented 4 months ago
👁️ ...

This looks good to me. I am happy for it to be checked in.
👍


Allison Clayborne commented 4 months ago
👁️ ...

Thanks for the confirmation.
👍

Add a comment

Drag and drop files to attach them or [Choose your files](#)

Post

☐ Flag as Task
☒ Post and notify (0)

Emoji

Swarm comments support Emoji shorthand. So when you save a comment, emoticon text like

`:smile:` is displayed as: 😊

Emoji emoticons are listed in the [Emoji Cheat Sheet](#).

Links in comments

Whenever you include a URL in a comment, it is automatically made into a link.

If the link points to an image, or a YouTube video, that resource is displayed at the end of the comment. For information about linking images and videos, see ["Common text styles" on page 415](#).

Comment attachments


Note

- Swarm must be [configured to enable comment attachments](#). Once the configuration is complete, the comment area will include the following text **Drop files here to attach them**.
- By default, the maximum file size of a single comment attachment is limited to:
 - **Ubuntu, RHEL 7, and CentOS:** 8Mb
 - **RHEL 8 and later:** 2Mb

Your Swarm administrator can increase the maximum attachment size if required. See ["Increasing the maximum attachment file size" on page 581](#).

Files can be attached to comments. This is useful for sharing files in code reviews, such as screenshots of error conditions, reference code, and documents.

When a file is attached to a comment, you can:

- Hover over the attachment to see its filename and file size.
- Click the attachment to open the file in a new browser tab.
- Hover over the attachment and click the **Download**  button to download the file.

Adding attachments to a comment


Tip

Multiple files can be attached to a comment, either one at a time or multiple files at the same time.

To attach a file to an open comment:

1. Do one of the following:
 - Drag the files from your file browser and drop them on the comment.
If you drag a folder onto a comment, all of the files in the folder and the files in any subfolders are added to the comment.
 - Click **Choose your files** and select the files from the browse dialog.
You cannot add folders to a comment using the file picker.
2. If you attach a file to the comment by mistake, you can remove it by clicking the **X** button on the attachment.
3. Click **Post**.
The file attachments are displayed below the comment text and image thumbnails are displayed for supported image formats.


Removing an attachment from a comment

1. Click the comment **Edit comment**  link.
2. Hover over the attachment you want to remove and click **X** to remove the file.
3. Confirm you want to remove the attachment when prompted.
4. Click **Save** to save the comment.

Tip

If you remove the wrong file attachment but have not saved the comment yet, click **Cancel**, and the file will remain attached to the comment.

Reacting to comments

As an authenticated user, you can *like* a comment by clicking the thumbs up  icon beneath the comment.

When you like a comment, a [Notification](#) is sent to the author of the comment, and the thumbs up icon changes to yellow to indicate that you have liked the comment. The number of likes the comment has is displayed next to the thumbs up icon. If a comment is edited, all of the likes for that comment are removed.

Hover your mouse pointer over the number of likes to display the usernames of everyone that has liked the comment.

Click the thumbs up icon again to *unlike* a comment. You can like/unlike a comment that has been archived.

Comment context

When comments are added to files in a review, on lines that have been changed, Swarm records several lines of context before the line receiving the comment. This helps makes sense of the comments should later changes remove those lines.


The screenshot shows a code review interface. At the top, a code diff is displayed with line numbers 90 to 105. Lines 97 to 105 are highlighted in green, indicating changes. The code is in C and includes a function `TalisAliquam`. Below the code, a comment thread is visible. The first comment is by Allison Clayborne, dated 10 months ago, asking "@steve.russell what do you think of this?". Below it, a comment by Steve Russell, dated 4 months ago, says "This looks good to me. I am happy for it to be checked in.". Below that, another comment by Allison Clayborne, dated 4 months ago, says "Thanks for the confirmation.". At the bottom of the comment thread, there is a button labeled "Add a comment".

Each comment associated with that line has a record of the context, but only the first comment displays that context.

Mark comments as read

When you mark a comment as read, the comment is rolled up into a single line to save space and make it easier for you to find comments you have not read. The read flag is remembered independently for each user.

Mark a single comment as read:

1. Click the **Mark comment as read**  button to the right of the comment.
2. The comment is rolled up into a single line to save space.

The screenshot shows a comment by Vernon Renno, dated "commented on review 1805 (version 3) - about 3 hours ago". To the right of the comment, there is a green "read" status with a checkmark icon.


Note








Marking a parent comment as read will not mark the child comments as read.

Mark all of the comments on a review as read:

Tip

Mark all comments read will mark all the comments as read including the description comments.

1. Click the **Review actions**  button.
2. Select **Mark all comments read** from the dropdown menu.
3. All of the comments in the review are rolled up into single lines to save space.


	Steve Russell commented on review 644 (version 1) //projects/mercury/dev/src/api/db.c, line 51 - 7 months ago	read 
	Allison Clayborne commented on review 644 (version 1) //projects/mercury/dev/src/api/db.c, line 105 - 7 months ago	read 
	Allison Clayborne commented on review 644 (version) - 10 days ago	read 
	Steve Russell commented on review 644 (version 2) //projects/mercury/dev/src/api/db.c, line 67 - 3 days ago (edited 3 days ago)	read 
	Jack Boone commented on review 644 (version 2) - about 15 hours ago	 read 

Mark comments as unread

Marking a comment as unread expands the comment so that you can view the comment content. If a comment is marked as read and the comment changes, Swarm will automatically clear the read flag so that you can see that the comment has changed. Changes that automatically mark a comment as unread are:

- Comment text is edited
- Comment attachments are added or removed
- Comment is marked as a task
- Task is reopened
- Comment or task is unarchived

Mark a single comment as unread:

1. Click the **read**  button to the right of the comment.
2. The comment content is expanded.


Note

Marking a parent comment as unread will not mark the child comments as unread.

Mark all of the comments on a review as unread:

Tip

Mark all comments unread will expand the content of all the comments including the description comments.

1. Click the **Review actions**  button.
2. Select **Mark all comments unread** from the dropdown menu.
3. The content of all of the comments in the review are expanded.

Archiving comments


Tip

When you archive a comment, it is archived for all of the Swarm users.

As a code review progresses, comments made on earlier versions of a file might become less useful. Archiving these comments tidies up the comment view and makes it easier to find the more important comments.

The **Archive** option is only available for top level comments, if a comment has replies they are archived with the parent comment.

To archive a comment:

1. Click the **Comment actions**  button at the top right of the comment you want to archive.
2. Select **Archive** from the dropdown menu.

Archived comments are hidden from view, the number of archived comments that exist for the review is displayed in the **archived comments** button at the top of the **Comments** tab.

Files **1** Comments **7** Activity







2 archived comments

To display and hide archived comments:

Click on the **archived comments** button to toggle the archive comment display on and off.

Files **1** **Comments 8** Activity

2 archived comments

-  **Vernon Renno** commented on review 1710 (version 1) - about a minute ago  
- Please run this past QA before release.
-  **Vernon Renno** commented on review 1710 (version 1) - less than a minute ago  
- Thanks

Restoring comments

Tip

When you restore a comment, it is restored for all of the Swarm users.







Archived comments can be restored. If archived comments have replies, the replies are also restored.


To restore an archived comment:

1. Click the **archived comments** button at the top of the **Comments** tab to display all of the archived comments.

Files **1** **Comments 8** Activity

2 archived comments

-  **Vernon Renno** commented on review 1710 (version 1) - about a minute ago  
- Please run this past QA before release.
-  **Vernon Renno** commented on review 1710 (version 1) - less than a minute ago  
- Thanks

2. Find the comment you want to restore.
3. Click the **Restore** button  in the top right of the comment to restore it.


Users

Swarm's users are based on the users configured in the Helix Core Server.

- "Viewing your user profile" on the facing page
- "Viewing another user's profile" on page 371
- "Viewing another user's profile when you have admin or super user privileges" on page 372

Viewing your user profile


Display your own user profile when you are logged in by clicking on your userid in the header and selecting **My Profile**.



The user profile page displays your [avatar](#), your full name, your email address, an **Unfollow all projects and users for your-username** button, a list of users following your activity, the users or [projects](#) that you are following, and the list of projects of which you are an owner or member, including an eye icon (👁️) beside any private projects you belong to.



Users/Allison.Clayborne


allison.clayborne




2 FOLLOWERS
1 FOLLOWING
6 PROJECTS

FULL NAME
Allison Clayborne
EMAIL ADDRESS
allison.clayborne@nowhere.xxjzzj.com
Unfollow all Projects and Users for allison.clayborne

FOLLOWERS



FOLLOWING


PROJECTS
 Blue Book
Owner, Member
JPlugin
Owner, Member
Maker
Owner, Member
Test Data
Member
Test project for workflow
Member
Mercury
Following

Activity Shelves Settings Notifications Reviews Commits Comments Jobs

Allison Clayborne updated project (JPlugin)
A Java plugin for continuous integration. 4 days ago
Allison Clayborne updated project (JPlugin)
A Java plugin for continuous integration. 4 days ago
Allison Clayborne updated project (JPlugin)
A Java plugin for continuous integration. 4 days ago
Allison Clayborne commented on review 264 (revision 1) (client.cc, line 37) for mercury:dev
The strings should be declared as constants. 25 days ago
3 comments
Allison Clayborne rejected review 232 (revision 1) for test-data:data
Updating tests for test-18.txt, about a month ago
Add a comment
Allison Clayborne approved review 230 (revision 1) for test-data:data
Updating tests for test-17.txt, about a month ago
Add a comment
Allison Clayborne approved review 215 (revision 1) for test-data:data
Updating tests for test-09.txt, about a month ago
Add a comment
Allison Clayborne approved review 199 (revision 1) for test-data:data
Updating tests for test-01.txt, about a month ago
Add a comment
Allison Clayborne approved review 197 (revision 1) for test-data:data
Updating tests for test-00.txt, about a month ago
Add a comment
Allison Clayborne approved review 194 (revision 1) for test-data:data
Updating test data to the latest format. about a month ago
Add a comment
Allison Clayborne (on behalf of steve.russell) committed change 309 into test-data:data
Updating test data to the latest format. about a month ago
Add a comment
Allison Clayborne requested review 307 (revision 1) for jplugin:main
Record and return the number of ticks per thread. about a month ago
Add a comment
Allison Clayborne requested review 303 (revision 1) for jplugin:main
The user property should be based on the email address of the user about a month ago

Unfollow all projects and users for yourself

When you are viewing your own user profile, you can unfollow all projects and users that you are following.

Note


This action cannot be undone.

1. Click the **Unfollow all projects and users for *your-username*** button.
2. Click **OK** when the confirmation dialog is displayed to complete the unfollow action.
3. You will no longer be following any projects or users.

Activity tab


Click the **Activity** tab to display the [activity stream](#) for events you have created.

Activity
Shelves
Settings
Notifications
Reviews
Commits
Comments
Jobs




Allison Clayborne updated project (JPlugin)
 A Java plugin for continuous integration.

4 days ago




Allison Clayborne updated project (JPlugin)
 A Java plugin for continuous integration.

4 days ago




Allison Clayborne updated project (JPlugin)
 A Java plugin for continuous integration.

4 days ago




Allison Clayborne commented on [review 264 \(revision 1\)](#) ([client.cc](#), [line 37](#)) for [mercury:dev](#)
 The strings should be declared as constants.

25 days ago
 3 comments



Allison Clayborne rejected [review 232 \(revision 1\)](#) for [test-data:data](#)
 Updating tests for test-18.txt,

about a month ago
 Add a comment



Allison Clayborne approved [review 230 \(revision 1\)](#) for [test-data:data](#)
 Updating tests for test-17.txt,

about a month ago
 Add a comment

Shelves tab

Click the **Shelves** tab to display a list of your shelved changelists.

<div> <div>Activity</div> <div>Shelves</div> <div>Settings</div> <div>Notifications</div> </div>			
Change	Description	Created	
1380317	#review-1380318 Update the coverage of My Reviews to cover the new look/behaviour. @ph...	4 days ago	<div> <div></div> <div>View Review</div> </div>
1379450	Update the phone numbers listed in Swarm's contact page to match those (a...	5 days ago	<div> <div></div> <div>Request Review</div> </div>
1378653	#review-1378654 Remove language indicating that Swarm is at an early stage. @pmclary s...	7 days ago	<div> <div></div> <div>View Review</div> </div>
1378432	#review-1378433 Add coverage of the new p4 -> auto_create_url configuration item, whic...	7 days ago	<div> <div></div> <div>View Review</div> </div>
1371898	#review-1371899 Update the admin configuration to capture the change in project settings...	26 days ago	<div> <div></div> <div>View Review</div> </div>
1371895	#review-1371896 Remove the note fro the JIRA integration page about manual configuration...	26 days ago	<div> <div></div> <div>View Review</div> </div>
1371884	#review-1371886 Update delayed notifications description to make it clear that the checkb...	26 days ago	<div> <div></div> <div>View Review</div> </div>
1357393	#review-1357327 My tweaks to Jenny's review.	2 months ago	<div> <div></div> <div>View Review</div> </div>

A *shelved* changelist is a pending changelist that has a copy of one or more files from within the changelist stored on the server. Shelved files are not versioned. If you update the shelved files, the update replaces any existing files on the changelist's shelf.

Note

Swarm can use multiple shelved changes to record the history of reviews. See "[Internal representation](#)" on [page 424](#) for details.

Swarm uses shelved changelists as the basis of its code review feature. However, not all shelved changelists are reviews. Users may shelf files for other reasons, including ensuring that the Helix Core Server has a copy of work in progress, or as a way to move temporary work from one workspace to another.

Click **Request Review** to start a Swarm review for any shelved changelist that is not already involved in a review.

Click **View Review** to view the Swarm review associated with shelved changelists when a review has already started.

Important

By default, when you delete files from a shelved changelist, the files are not removed from the associated review.

However, Swarm can be configured to remove files from a review when they are deleted from an associated shelf, see ["Process shelf file delete when" on page 693](#). To delete a shelved changelist without removing all of the shelved files from the associated review, see ["Deleting shelves" on page 469](#).

Do not use the Swarm user that is configured in the [Swarm configuration file](#) when deleting shelves, or deleting files from shelves. The Swarm logic processes the *shelve-delete* trigger event, if the event is invoked by the Swarm user it is rejected. The delete operations will fail.

Settings tab

The **Settings** tab allows you to configure how [diffs](#) are initially displayed when you view them, and how time is displayed by Swarm.

Note

Defaults for these options are configured by the Swarm administrator, see [Users](#). The defaults set by the Swarm administrator are used until you change your settings.

Activity
Shelves
Settings
Notifications

Profile Settings

Reset to default

Reviews

Sets default values for how diffs are displayed on all reviews.

Show comments in files	<input checked="" type="checkbox"/>
View diffs side-by-side	<input checked="" type="checkbox"/>
Show space and newline characters	<input type="checkbox"/>
Ignore whitespace when calculating differences	<input type="checkbox"/>

Time Display

Sets how you would like to display time across Swarm.

Display the time in Timestamp

Save
Cancel

Toggle the **Reviews** settings to control how diffs are initially displayed on changelists and reviews by default.

Choose the **Time Display** settings to configure how time is displayed by Swarm. You can configure Swarm to display timestamps based on your local machine browser time, or by how long ago events happened.

Tip

If timestamp is selected, the date format used by Swarm matches the date format configuration of the local machine browser.

Click the **Save** button to save the settings. Click **Reset to default** to reset the settings back to system defaults.

Notifications tab

The **Notifications** tab allows you to configure which notifications you receive when events occur within Swarm. This allows you to limit the number of emails you receive to just those you are interested in. The settings apply across all projects.

Note

Defaults for these options are configured by the Swarm administrator. Your Swarm administrator may force some of these options to on or off, see ["Global settings" on page 664](#).

[Activity](#)
[Shelves](#)
[Settings](#)
[Notifications](#)

Adjust when notifications are sent to you about reviews that you're associated with (as an author, reviewer, project member or moderator).

Email me when:
Reset to default

Check all ☐

I change the state of a review	<input type="checkbox"/>
I am the author, and	
a review is requested	<input checked="" type="checkbox"/>
files in the review are updated	<input checked="" type="checkbox"/>
tests on the review have finished	<input checked="" type="checkbox"/>
a vote is cast on a review	<input checked="" type="checkbox"/>
the state of the review changes	<input checked="" type="checkbox"/>
a review or change is committed	<input checked="" type="checkbox"/>
a comment is made on the review or change	<input checked="" type="checkbox"/>
a comment on the review or change is updated	<input checked="" type="checkbox"/>
Someone likes one of my comments	<input checked="" type="checkbox"/>
I am a member, and	
a review is requested	<input checked="" type="checkbox"/>
a review or change is committed	<input checked="" type="checkbox"/>
I am a reviewer, and	
files in the review are updated	<input checked="" type="checkbox"/>
tests on the review have finished	<input checked="" type="checkbox"/>
a vote is cast on a review	<input checked="" type="checkbox"/>
the state of the review changes	<input checked="" type="checkbox"/>
someone joins or leaves the review	<input checked="" type="checkbox"/>
a review or change is committed	<input checked="" type="checkbox"/>
a comment is made on the review	<input checked="" type="checkbox"/>
a comment on the review is updated	<input checked="" type="checkbox"/>
I am a moderator, and	
files in the review are updated	<input checked="" type="checkbox"/>
tests on the review have finished	<input checked="" type="checkbox"/>
a review or change is committed	<input checked="" type="checkbox"/>

Save
Cancel

Toggle notifications for each event on or off to control whether you receive an email when that event occurs.

Click the **Save** button to save the settings. Click **Reset to default** to reset the settings back to system defaults.

Viewing another user's profile

View the profile pages of other users displayed anywhere in Swarm by clicking on their [avatar](#), *userid*, or by visiting the URL: `https://myswarm.url/users/userid`.

Note

When viewing another user's profile page you can see the user's **Settings** tab but you cannot make changes to it.

Note

Currently, Swarm does not provide an overall list of users.

Follow or unfollow the user

- **If you are not following the user:** click the **Follow** button below the user's avatar to start following the user.
- **If you are following the user:** click the **Unfollow** button below the user's avatar to stop following the user.

Users/Alex.Randolph

alex_randolph

Activity Shelves Settings Notifications Reviews Commits Comments Jobs

Alex Randolph committed change 295 into mercury:main
A few minor changes.
Add a comment 4 months ago

Alex Randolph approved review 177 (revision 1)
Adding basics.notifications to documentation.
#review @@Idocs
Add a comment about a year ago

Alex Randolph approved review 165 (revision 1)
Adding basics.files to documentation.
#review @@Idocs
1 comment about a year ago

Alex Randolph approved review 131 (revision 1)
Adding admin.search to documentation.
#review @@Idocs
2 comments about a year ago

Alex Randolph approved review 129 (revision 1)
Adding admin.reviews to documentation.
#review @@Idocs
3 comments about a year ago

Alex Randolph approved review 179 (revision 1)
Adding basics.projects to documentation.
#review @@Idocs
2 comments about a year ago

Alex Randolph approved review 175 (revision 1)
Adding basics.notable to documentation.
#review @@Idocs
Add a comment about a year ago

Alex Randolph voted down review 119 (revision 1)
Adding admin.projects to documentation.
#review @@Idocs
9 comments about a year ago

Alex Randolph cleared an issue on review 135 (revision 1) (P4Browser.java, line 17) for jplugin:main
Shouldn't the graph be enabled by default?
3 comments about a year ago

FULL NAME
Alex Randolph
EMAIL ADDRESS
alex.randolph@nowhere.xxjzzj.com

FOLLOWING

PROJECTS
 Blue Book
Owner, Member
Mercury
Member
Test Data
Member
Test project for workflow
Member

Viewing another user's profile when you have admin or super user privileges

View the profile pages of other users displayed anywhere in Swarm by clicking on their [avatar](#), [userid](#), or by visiting the URL: `https://myswarm.url/users/userid`.

Unfollow all projects and users for another user

When a user has been removed from the Helix Core Server but they are still following projects and users, it is useful to be able to remove all of their follows. This helps to keep the project and user follower lists up to date.

Warning

This action cannot be undone.

1. Click the **Unfollow all projects and users for *username*** button located below the user's email address.
2. Click **OK** when the confirmation dialog is displayed to complete the unfollow action.
3. The user will no longer be following any projects or users.

Users/Alex.Randolph

alex.randolph

Activity Shelves Settings Notifications Reviews Commits Comments Jobs

Alex Randolph committed [change 295](#) into [mercury:main](#) 4 months ago
A few minor changes.
Add a comment

Alex Randolph approved review 177 (revision 1) about a year ago
Adding basics.notifications to documentation.
#review @@Idocs
Add a comment

Alex Randolph approved review 165 (revision 1) about a year ago
Adding basics.files to documentation.
#review @@Idocs
1 comment

Alex Randolph approved review 131 (revision 1) about a year ago
Adding admin.search to documentation.
#review @@Idocs
2 comments

Alex Randolph approved review 129 (revision 1) about a year ago
Adding admin.reviews to documentation.
#review @@Idocs
3 comments

Alex Randolph approved review 179 (revision 1) about a year ago
Adding basics.projects to documentation.
#review @@Idocs
2 comments

FULL NAME
Alex Randolph
EMAIL ADDRESS
alex.randolph@nowhere.xojzzj.com

Unfollow all Projects and Users for alex.randolph

0 FOLLOWERS 1 FOLLOWING 4 PROJECTS

Groups

Groups are a feature of Helix Core Server that makes it easier to manage permissions for users.

It is important to be aware of certain aspects of Helix Core Server groups:


- Groups can have both users and sub-groups. A user that is a member of a sub-group is also a member of the parent group.
- Groups have owners, but owners are not members of their groups. A group owner can be added as a member of a group.
- Groups can be created by users with *super* privileges in Helix Core Server (**p4d**). If p4d is at version 2012.1 or newer, users with *admin* privileges can also add groups.
- Groups can be edited by their owners, or by users with *super* privileges in Helix Core Server.
- Groups have a separate namespace from users; you can have a user named **fish** that is a member of a group named **fish**.
- Groups can be project moderators
- Groups can be reviewers, see [Review display](#) and [@@mention](#) for details.






- You can **@mention** a group in a code review comment to ensure the mentioned group receives **notifications** of code review events. This also creates a link that displays the group's details when clicked. See **@mention** for details.
- Groups cannot be project owners.

More information on adding, editing, and removing groups is included in the **"Groups"** on page 491 chapter.

Listing groups

Begin browsing groups by clicking **Groups** in the menu.




Name	Description	Owners	Members
Marketing marketing	Marketing team		3
triage	Triage team.	 	5
docs	Docs team		3
swarm-group	Swarm group		1
swarm-support	Swarm support team		3


For each group, the group's name, description, list of owner avatars, and a membership count (which has a darker background when you are a member). Click on a group name to display details for that group.

The groups listing is sorted by multiple criteria:

- Groups that you own, or belong to, are listed first. A thicker row border indicates where your group ownership/membership ends; you are not an owner/member of any group below the thicker border.
- Groups are then sorted by name.

You can search available groups by entering some text in the **Search** field and clicking **Search**. Any group names or descriptions that match the entered text are displayed.



Name	Description	Owners	Members
swarm-group	Swarm group		1
swarm-support	Swarm support team		3

If you have *super* privileges in Helix Core Server (P4D), or have admin privileges in P4D version 2012.1 or newer, Swarm displays the **+ Add Group** button. Click **+ Add Group** to add a new group.

Note

If your Helix Core Server has a large number of groups, the time required to display the list of groups might be notable.

Viewing a group

View a specific group by clicking on a linked group name, or by visiting the URL:

`https://myswarm.url/groups/group-id`.

Activity page

The group **Activity** page shows the **activity stream** for events generated group members to the right of the group sidebar.

The screenshot shows the 'docs' group page in Helix Core Swarm. On the left is the group sidebar, and on the right is the activity stream.

Group Sidebar (Left):

- Group name: docs
- Group avatar: A purple circular logo with three stylized bees.
- Group description: Docs team
- Owner: 1 (indicated by a large number 1)
- Members: 3 (indicated by a large number 3)
- Email address: docs.team@nowhere.xxjzj.com
- Owners: One owner profile picture is shown.
- Members: Three member profile pictures are shown.

Activity Stream (Right):

- Tab: Activity (selected)
- Activity items (all from 7 days ago):
 - 1. **paula.boyle** updated files in [review 249 \(revision 8\)](#) for `jplugin:main`. Adding description to P4Tickets. [Add a comment](#)
 - 2. **paula.boyle** updated files in [review 249 \(revision 7\)](#) for `jplugin:main`. Adding description to P4Basecredentials and P4Credentials. [Add a comment](#)
 - 3. **paula.boyle** updated files in [review 249 \(revision 6\)](#) for `jplugin:main`. Adding description to P4Basecredentials and P4Credentials. [Add a comment](#)
 - 4. **paula.boyle** updated files in [review 249 \(revision 5\)](#) for `jplugin:main`. Minor update in description of P4Password. [Add a comment](#)
 - 5. **paula.boyle** updated files in [review 249 \(revision 4\)](#) for `jplugin:main`. Adding description to P4Credentials & P4CredentialsImpl. [Add a comment](#)
 - 6. **paula.boyle** updated files in [review 249 \(revision 3\)](#) for `jplugin:main`. Adding description to P4Basecredentials and P4Credentials. [Add a comment](#)
 - 7. **paula.boyle** requested [review 249 \(revision 1\)](#) for `jplugin:main`. Adding description to P4Credentials & P4CredentialsImpl. [Add a comment](#)

The following information is displayed in the group sidebar:

- Group name
- Group avatar, if it has one
- Group description, if it has one

- Count of the number of owners and members of the group
- Group email address, if it has one
- Avatars of the group owners
- Avatars of the group members

Reviews page

The group **Reviews** page shows a list of code reviews authored by any of the group members. It lists code reviews that are in progress and code reviews that are complete.

ID	Description	Project	Created	Status
368	Updated definitions (again)	mercury:dev	5 months ago	6 / 0
330	Added a missing comment.	jplugin:main	7 months ago	0 / 0
328	Add comment to method.	jplugin:main	7 months ago	0 / 0
264	Add functionality to report on the current status of the server. This is so the client ca...	mercury:dev	2 years ago	3 / 0
119	Adding admin.projects to documentation. #review @@!docs		3 years ago	9 / 1
105	Optimised some of the error checking to improve performance. #review @@!triage	mercury:main	3 years ago	2 / 0
77	Adding admin.files to documentation. #review @@!docs		3 years ago	2 / 0

The following information is displayed in the group sidebar:








- Group name
- Group avatar, if it has one
- Group description, if it has one
- Count of the number of owners and members of the group
- Group email address, if it has one
- Avatars of the group owners
- Avatars of the group members

Reviews list


The **Opened** and **Closed** tabs display:

- **Opened tab:** displays a list of all code reviews that have started, are being reviewed, are awaiting revisions, or need to be committed.
- **Closed tab:** displays a list of all code reviews that have been approved and committed, rejected, or archived.

Each reviews list displays a summary for each review:

ID	Description	Project	Created						
368	 Updated definitions (again)	mercury:dev	5 months ago						

- The review id
- The avatar of the review author
- The review description
- The associated project name and branch
- **Created or Last activity** dropdown: click to change the order the reviews are displayed in, options are:
 - **Created:** Reviews sorted by when they were created
 - **Last activity:** Reviews sorted by when they were last updated



Note

- If the **Result order** button is not displayed reviews are sorted by when they were created.
- The **Result order** button display is a global setting controlled by the Swarm administrator. See ["Reviews filter" on page 686](#) for details.
- When review results are older than 24 hours they are displayed in numerical order within each day.

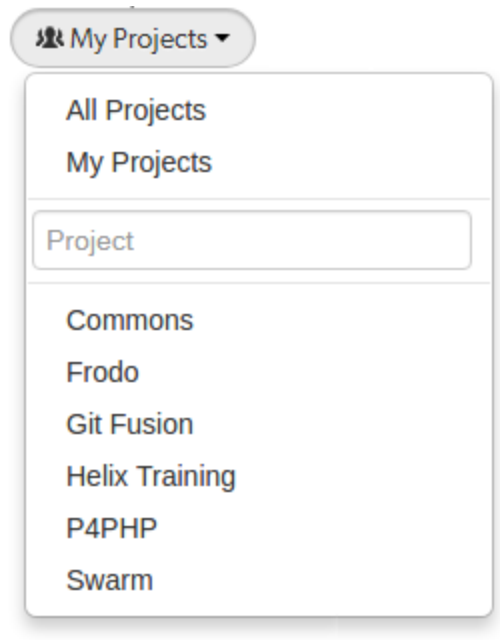
- An icon indicating the current review state
- An icon indicating the review type. This can be Pre-commit or Post-commit.
- An icon indicating the test suite state for the review, either tests in progress, tests passed, or tests failed.

- A counter for the number of open (non-archived) comments that are associated with the review. Hover your mouse over the comment count to display a tooltip showing the number of archived comments associated with the review.
- An indicator showing the number of up votes and down votes


Filtering reviews

The following filtering options are available for code reviews:

- **Project:** a dropdown menu that lets you filter which reviews to display based on project:



- **All Projects:** all reviews are displayed for all projects.
- **My Projects:** all reviews for all of the projects you are participating in, as a member, owner, moderator, or follower.
- **Project Search:** an auto-complete search field that allows you to choose one of the projects defined in Helix Core Server. Once specified, only reviews for the selected project are displayed. Click the **X** button to remove a *projectid* after it has been specified.
- **Select Project:** selecting a project name displays only reviews for that project.
When you select one of the available options, the list of options updates to match the currently selected filter, and the **Projects** dropdown indicates the current filter: **All Projects**, **My Projects**, or *projectid*.
- **Users:** a dropdown menu that lets you filter which reviews to display based on user involvement:



- **All Reviews:** displays all reviews.
- **Reviews I've Authored:** displays reviews that you have authored.
- **Reviews I'm Participating In:** displays reviews that you are a reviewer of, but not an author of.
- **Reviews I've Authored Or Am Participating In:** displays reviews that you have authored, or are a reviewer of.
- **Specific User:** An auto-complete search field that allows you to choose one of the user accounts defined in the Helix Core Server. Once specified, only reviews authored by the user are displayed. Click the **X** button to remove a *userid* after it has been specified.

When you select one of the available options, the list of options updates to match the currently selected filter, and the **Users** dropdown indicates the current filter: **All**, **Author**, **Participant**, or *userid*.

■ **Reviewer presence (Opened tab only):**



- **Has Reviewers:** displays reviews that have one or more reviewers.
- **No Reviewers:** displays reviews that have no reviewers.

■ **Review state (Opened tab only):**




- **Needs review:** the review's changes need to be reviewed.
- **Needs revision:** the review's changes have been reviewed, but further revisions are required before the review can be accepted.
- **Approved:** the review's changes have been approved, and should be committed.

■ **Review state (Closed tab only):**



- **Approved:** the review's changes have been approved and committed.
- **Rejected:** the review's changes have been rejected.
- **Archived:** the review's changes have been put aside.

■ **Test status:**



- **Tests pass:** when automated tests are enabled for the associated project, and the test suite execution succeeds, Swarm updates the review accordingly.
- **Tests fail:** similar to the Tests pass state, except that the test suite execution has failed. Check with your test suite to determine why the tests failed.

■ **Vote status:**



- **Voted up:** I have voted the review up.
- **Voted down:** I have voted the review down.
- **Not voted:** I am a participant but have not voted on the review.

Note

Filters for voting only apply to reviews which you are a participant of. Commenting on or voting on a review will automatically add you as a participant. If you leave the review after commenting on it, then this review will not be included in the list.

■ **Comment status:**



- **Has comments:** I have commented on the review.
- **Does not have comments:** I have not commented on the review.

Note

Filters for commenting only apply to reviews which you are a participant of. Commenting on or voting on a review will automatically add you as a participant. If you leave the review after voting on it, then this review will not be included in the list.

- **Bookmark:** review filters can be remembered by bookmarking the page, and this icon acts as a reminder of that.



Swarm updates the URL in your browser to reflect filtering options. This makes it easy to bookmark or share review list URLs. Swarm maintains the current filtering if you click on a review link and then use your browser **Back** button to return to the review list.


- **Search term:** where review descriptions match your search string.

Projects

A Swarm project is made up of a group of Helix Core Server [users](#) who are working together on one or more codelines within the Helix Core Server. A project's definition includes one or more branches of code, and optionally a job filter, [automated test integration](#), and [automated deployment](#). This section provides an introduction to the interactions users have with projects. For details about managing projects, see ["Projects" on page 498](#).

Listing projects



To view a list of projects, click **Projects** in the menu.



Logged-in users can choose which projects to display by clicking on the **My Projects** tab or the **All Projects** tab. Anonymous users only see the public projects, the **My Projects** tab is not available to anonymous users:

- The **My Projects**: tab lists all of the projects that you are an owner of or a member of.
- The **All Projects**: tab lists all of the available projects.

Tip


- The **Project Owner** icon  is displayed for projects you own.
- The **Private Project** icon  is displayed for private projects you are authorized to view. For details, see ["Private projects" on page 387](#).
- The number in the icon to the right of the project name displays the total number of members for the project.

Search project names and project descriptions

You can search for Swarm project names and for the content of the project description from the projects page:

1. Enter text in the search box.
2. Click **Search**.

Any project names, and descriptions that match the search text are displayed in the project tab:



The screenshot shows the Swarm Projects interface. At the top, there is a search bar with the text 'co' entered. To the right of the search bar is a 'Search' button, which is highlighted with a blue border. Further right is a '+ Add Project' button. Below the search bar are two tabs: 'My Projects' and 'All Projects'. The 'All Projects' tab is selected. Below the tabs, there are two project cards. The first card is for 'JPlugin', described as 'A Java plugin for continuous integration.', and has a '3' in a grey box to the right of the project name. The second card is for 'Mercury', described as 'Mercury project for code management.', and also has a '3' in a grey box to the right of the project name.

Tip

Switch to the other project tab to display search results for that tab.

Viewing a project

View a project by doing one of the following:

- Click on the project name in the projects list on the Swarm **Projects** page.
- Click on the project name or branch identifier in an [activity stream](#).
- Visit the URL: `https://myswarm.url/projects/project-name`.

Overview page

The project **Overview** page shows a description of the project.

Important

The project's **Overview** page is only displayed if there is a **README** markdown file in the project's mainline.

The **README** file can contain Markdown text, allowing a formatted description of the project to be provided. For a description of the type of formatting that is supported, see ["Markdown in projects" on page 420](#). For details on how to configure the mainline of a project, see ["Mainline branch identification" on page 636](#).

HelixSwarm Projects/Jplugin

About
A Java plugin for continuous integration.

3 MEMBERS **1** FOLLOWER **2** BRANCHES

OWNERS

MODERATORS

MEMBERS

FOLLOWERS

BRANCHES
MAIN
CANDIDATE

P4 Plugin

Jenkins plugin for a Perforce Helix Versioning Engine (P4D).

Contents

- [Release notes](#)
- [Setup guide](#)
- [Notes page](#)
- [Jenkins page](#)

Requirements

- Jenkins 1.642.3 or greater.
- Helix Versioning Engine 2012.1 or greater.
- Minimum Perforce Protection of **open** for the Jenkins user.
- Review Build feature requires Swarm 2014.2 or greater.

Install

1. Open Jenkins in a browser; e.g. http://jenkins_host:8080
2. Browse to 'Manage Jenkins' --> 'Manage Plugins' and Select the 'Available' tab.
3. Find the 'P4 Plugin' or use the Filter if needed
4. Check the box and press the 'Install without restart' button

If you are unable to find the plugin, you may need to refresh the Update site.

1. Select the 'Advanced' tab (under 'Manage Plugins')
2. Press the 'Check now' button at the bottom of the page.
3. When 'Done' go back to the update centre and try again.

Building

To build the plugin and run the tests use the following:

The following information is displayed in the project sidebar:

- A short description of the project
- A **Follow** button, click to follow the project. This button is not available if you are a member of the project.
- A list of project owners
- A list of project moderators
- A list of project members
- A list of project followers
- A list of the branches defined for the project

Tip

Hover over a branch name to view the moderators on that branch.

Activity page

The project **Activity** page shows the [activity stream](#) for the project

The screenshot shows the HelixSwarm interface for the 'JPlugin' project. On the left is a sidebar with a 'MAIN MENU' containing 'Overview', 'Activity' (selected), 'Reviews', 'Files', 'Commits', and 'Settings'. The 'Activity' section shows counts: 3 MEMBERS, 1 FOLLOWER, and 2 BRANCHES. Below these are lists for OWNERS, MODERATORS, MEMBERS, FOLLOWERS, and BRANCHES (MAIN and CANDIDATE). The main area displays the activity stream with tabs for Reviews, Commits, Comments, and Jobs. The activity stream shows several updates and reviews by Allison Clayborne and Steve Russell, with timestamps ranging from 4 days ago to about a month ago.

The following information is displayed in the project sidebar:

- A short description of the project
- A **Follow** button, click to follow the project. This button is not available if you are a member of the project.
- A list of project owners
- A list of project moderators
- A list of project members
- A list of project followers
- A list of the branches defined for the project

Tip

Hover over a branch name to view the moderators on that branch.

Reviews page


The project **Reviews** page shows a list of code reviews specific to the project.

The screenshot shows the HelixSwarm interface for the JPlugin project. The left sidebar contains a 'MAIN MENU' with options: Overview, Activity, Reviews (selected), Files, Commits, and Settings. The main content area is titled 'Reviews' and has tabs for 'Opened' (5) and 'Closed' (9). A search bar is present. Below the tabs, there are filters for Branch, Role, Reviewers, States, Tests, Comments, and Votes. The reviews table has the following data:



Description	Created	State	Tests	Complexity	Comments	Votes
Added a missing comment. Steve Russell requested a pre-commit review 330, 6 months ago for jplugin:main	6 months ago	Open	0	6	0	0/0
Add comment to method. Steve Russell requested a post-commit review 328, 6 months ago for jplugin:main	6 months ago	Open	1	18	0	1/0
Updated some of the documentation. Allison Clayborne requested a pre-commit review 311, 6 months ago for jplugin:main	6 months ago	Open	0	137	1	3/0

For more details on browsing, filtering, and searching reviews, see ["Reviews list" on page 428](#).

Files page

The project **Files** page shows a list of files for the project, starting with a folder view representing each branch. Branches are designated with the *branch* icon .

The screenshot shows the HelixSwarm interface for the JPlugin project. The left sidebar contains a 'MAIN MENU' with options: Overview, Activity, Reviews, Files (selected), Commits, and Settings. The main content area is titled 'Projects/Jplugin' and has a search bar. Below the search bar, there are tabs for 'Browse' and 'Commits'. The 'Browse' tab is active, showing a list of branches. The list has the following data:

Name	Modified	Size
 main		
 candidate		

The project's *main* branch, identified by using a name such as *main*, *mainline*, *master*, *trunk*, is sorted to the top of the list of branches and appears in bold. The list of names can be configured, see ["Mainline branch identification" on page 636](#) for details.

For more information on browsing files, see ["Files" on page 322](#).

Commits page

The project **Commits** page shows a list of changes made to the project.

Projects/Jplugin

MAIN MENU

JPlugin

- Overview
- Activity
- Reviews
- Files
- Commits**
- Settings

Browse Commits

Range User

Change	User	Description	Committed
300	Claire Brevia	Update candidate from main. #review-301	about a month ago
286	Claire Brevia	Updated message text for the application. #review-287	about a year ago
241	Steve Russell	Created candidate branch for the plugin, so that we know what is ready to be released to ...	2 years ago
134	Allison Clayborne	Tidying up some of the code to fix some edge conditions. Also putting some support for f ...	2 years ago

For more details on history browsing, see ["Commits" on page 332](#).

Jobs page

The project **Jobs** page shows a list of jobs associated with the project.

Important

The **Jobs** page is only displayed when the project configuration includes a *job filter*. See ["Specify a Job filter." on page 513](#) for details.

Projects/Jplugin

MAIN MENU

JPlugin

- Overview
- Activity
- Reviews
- Files
- Commits
- Jobs**
- Settings

Query Search

Job	Status	Reported By	Description	Modified Date
job086341	Fixed	p4dtguser	Enable "require_login" by default. This is due to the way the new Groups feature interacts with anon...	37 minutes ago
job086335	Open	p4dtguser	[doc] Update coverage/mentions of require_login to note that it is now enabled by default. This f...	about 4 hours ago

For more details on browsing and searching jobs, see ["Jobs" on page 335](#).

Settings page

The project **Settings** page shows the settings associated with the project.

Tip

Open the **Branches** tab and select a branch to view the moderators on that branch.

Name	Type	
Amy Crover	User	
Delbert Henkel	User	
My Group	Group	

Rows per page: 25 1-3 of 3 < >

For more information about project settings, see ["Project settings" on page 499](#).

Private projects

Private projects, introduced in Swarm 2016.2, provide a way to make specific projects and their activity less visible to Swarm users. When a project is made private, only the projects owners, moderators, and members, plus users with *super* privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews.

If you are logged in as an owner, moderator, or member of a private Swarm project, that project appears on the Swarm project page with an eye icon to indicate that it is private and has limited visibility:

Similarly, the eye icon appears beside the project's title when viewing the project. When you hover your mouse over the eye icon, a tooltip appears indicating that this project is indeed private.

Caveats

The following are important caveats regarding private projects:

- While Swarm can mask the existence of projects, their activity, and reviews, Swarm honors each user's access to files within the Helix Core Server. This means that users that have access to the files in a private project's branches can browse to those branch paths within the depot and see the files and any committed changes.
- If you need to prevent access to important files, your Helix Core Server administrator is required to manage the protections table accordingly. For details, see [Authorizing Access](#) in the *Helix Core Server Administrator Guide*.
- It is possible for a user to start a review that touches files that belong to a private project. If the user is not a member, owner, or moderator of the private project, or does not have super privileges in the Helix Core Server, they cannot participate in the review.
- If a user is not a member, owner, or moderator of a private project, or does not have super privileges in the Helix Core Server, and they are added as a review participant (via an ["Links in descriptions and comments" on page 409](#) or by editing a review's participants) to a review containing files within the private project's branches, that user cannot participate in the review: the user cannot see the review, its files, or comments. Due to limitations in how Swarm sends email notifications, such users could still receive notifications for reviews they cannot participate in.
- Notifications usually include links and mentions of the associated projects. Notifications involving private projects are filtered to remove any references to those private projects.

Workflows

Important

Workflow is enabled by default and can be disabled by your Swarm administrator, see ["workflow " on page 752](#).

A workflow can be applied to a project or project branch to ensure that changelists and code reviews in the project/branch follow the rules specified in that workflow.

- Swarm workflows can be created by any Swarm user.
- Shared Swarm workflows can be viewed by any Swarm user.
- Shared Swarm workflows can be applied to a project or project branch by any Swarm user that is authorized to edit the project.
- The global workflow can be viewed by any Swarm user but can only be edited by a Swarm user with *super* or *admin* privileges, or users that have been made an owner.

Important

The Swarm administrator can enforce a minimum workflow rule setting (by setting the global rule to **Enforce**) for the entire Helix Core Server. If a project or project branch has an associated workflow, the global workflow rule is merged with the workflow rule and the most restrictive setting is used.


- For more information about how workflow rules are merged, see ["Merging multiple workflows" on page 532](#).
- For more information about setting Global Workflow rules, see ["Workflow global rules" on page 752](#).

This section provides an introduction to listing, searching, and viewing workflows.

- For an overview of workflows, see ["Workflow overview" on page 526](#).
- For instructions on how to add a workflow, see ["Add a workflow" on page 537](#).
- For instructions on how to delete a workflow, see ["Delete a workflow" on page 542](#).
- For instructions on how to add a workflow to a project, and a project branch, see ["Project settings" on page 499](#).

Listing workflows


To view a list of workflows, click **Workflows** in the menu.



Logged-in users can choose which workflows to display by clicking on the **My Workflows** tab or the **All Workflows** tab. Anonymous users only see the global and shared workflows, the **My Workflows** tab is not available to anonymous users:

- The **My Workflows**: tab lists all of the workflows that you are an owner of.
- The **All Workflows**: tab lists the global workflow, all of the shared workflows, and workflows that you are an owner of.

Tip


The **Workflow Owner** icon  is displayed for workflows you own.

Search workflow names and workflow descriptions

You can search for Swarm workflow names and for the content of the workflow description from the workflows page:

1. Enter text in the search box.
2. Click **Search**.

Any workflow names, and descriptions that match the search text are displayed in the workflows tab:


Tip

Switch to the other workflow tab to display search results for that tab.

Viewing a workflow that you own

To view workflow details:

1. Click **Workflows** in the menu.
2. Click on the name of the workflow you want to view.

For information about editing the workflow, see ["Add a workflow" on page 537](#).

Note

You can edit a workflow if you are the owner of the workflow, or if you have *super* user rights.

Candidate line workflow

Name

Candidate line workflow

Description

Use this workflow for candidate lines

Owners

Add an Owner

Individuals:

bruno

Shared with others

Rules

On commit without a review

Reject

On commit with a review

Reject unless approved

On update of a review in an end state

Reject

Count votes up from

Members

Automatically approve reviews

Based on vote count

Tests

When

Blocks

Code sniff	On Update	Approve
Doc tests	On Submit	Nothing
Security check	On Demand	Approve
+Add Test		

2 projects use this workflow

jPlugin

Branches

1. Candidate

Test Data

Branches

1. Candidate branch

Save

Cancel

Tip

- The project count does not include project branches.
- If the workflow is associated with a [Private project](#), the private project name is only displayed if you are authorized to view it. Private projects are included in the project count.

Viewing a shared workflow that you do not own

Shared workflows can be viewed and used by all Swarm users.

To view workflow details:

1. Click **Workflows** in the menu.
2. Click on the name of the workflow you want to view.

Main line workflow


Name

Main line workflow

Description

Use this workflow for main branches

Owners

 Add an Owner

Individuals:

swarm

☒ Shared with others

Rules

On commit without a review

Allow

On commit with a review

Allow

On update of a review in an end state

Allow

Count votes up from

Anyone

Automatically approve reviews

Never

Tests

	When	Blocks
Code sniff	On Update	Nothing
Build tests	On Submit	Approve
Doc tests	On Demand	Nothing

2 projects use this workflow

jPlugin

Branches

1. Main

Close

Tip

- The project count does not include project branches.
- If the workflow is associated with a [Private project](#), the private project name is only displayed if you are authorized to view it. Private projects are included in the project count.

- The project count does not include project branches.
- If the workflow is associated with a [Private project](#), the private project name is only displayed if you are authorized to view it. Private projects are included in the project count.

Viewing the global workflow

The global workflow can be viewed by any Swarm user but can only be edited by a Swarm user with *super* or *admin* privileges, or users that have been made an owner. For information about editing the global workflow, see ["Workflow global rules" on page 752](#).

To view the global workflow details:

1. Click **Workflows** in the menu.
2. Select the **All Workflows** tab.
3. Click on **Global Workflow** to view it.

Tip

By default, the global workflow is named **Global Workflow**, but it can be changed by your Swarm administrator. To help you identify the global workflow it is always shown as the first workflow in the **All Workflows** tab and it has a different background color to the other workflows.

Global Workflow x

Name

Global Workflow

Description

Description

Owners

Add an Owner

Individuals:

swarm

Global Rules ?

Enforce ?

On commit without a review

Allow



On commit with a review

Reject unless approved



On update of a review in an end state

Reject



Count votes up from

Members



Automatically approve reviews

Never



Members or groups who can ignore ALL workflow rules



Member or group

Groups:

swarm-test

Tests ?

When

Blocks

Code sniff

On Update

Nothing

Build tests

On Submit

Nothing

Security check

On Update

Approve

Doc tests

On Submit

Nothing

Close

Work-in-progress tag

Add the work-in-progress tag (**#wip** by default) to a pending changelist description to tell Swarm that your work is not ready to be reviewed yet. For example, when changes are requested for a review, and you want to address them in stages but get them all reviewed once your work is complete. When you are ready for your changes to go back into review, delete **#wip** from the changelist description and update your shelf to update the review.

Note

- By default, the work-in-progress tag is **#wip**, but this can be changed by your Swarm administrator if required. See ["Tag processor" on page 732](#).
- The following instructions assume your organization uses the default **#wip** work-in-progress tag. If they are using a different work-in-progress tag, substitute that tag in this page.
- The work-in-progress tag only applies to pending changelists ([pre-commit](#) reviews).

Summary

The work-in-progress tag:

- Prevents a new review from being created for the changelist, even if you have the **#review keyword** in the changelist description.
- Prevents an existing review from being updated when you add, edit, rename, or delete files from the changelist, even if you have the **#review-1234 keyword** in the changelist description.
- Prevents an existing review from being attached to the changelist, even if you have **keywords** in the changelist description. For example, **#append-1234**, **#replace-1234**, or **#review-1234**.
- Prevents the review description from being synchronized if ["Synchronize review description" on page 694](#) is enabled. When the **#wip** tag is deleted from the changelist description, Swarm will synchronize the review description.

When the **#wip** tag is deleted from the changelist, Swarm will update the review when you next update the changelist (add, edit, rename, or delete a file). Simply deleting the **#wip** tag will not update the review.

The work-in-progress tag will not let you break the workflow rules, so it is not a method for getting around the rules set by your organization.

The work-in-progress tag does not prevent a commit from updating its associated review.

Using the work-in-progress tag

Adding a work-in-progress tag to a pending changelist

1. Add **#wip** to your pending changelist description.
2. Update the files in your changelist as often as you need to.
Swarm will not create a new review or update an existing review.

Deleting a work-in-progress tag and updating the review

1. When you are ready for your changes to be reviewed, delete **#wip** from your changelist description.
2. Update your changelist shelf by making a file change (add, edit, rename, or delete a file).
Swarm will create a review or update the existing review with your changes so the reviewers can review them.

Tests

Important

The **Tests** page is only available if Workflow is enabled, workflow is enabled by default.

Associate a test with a [workflow](#) to ensure that the test is run when a review associated with that workflow is started or updated. Associate a test with the [global workflow](#) to ensure that the test is run whenever a Swarm review is started or updated. This ensures that the global tests are enforced for all changes even if they are not part of a project.

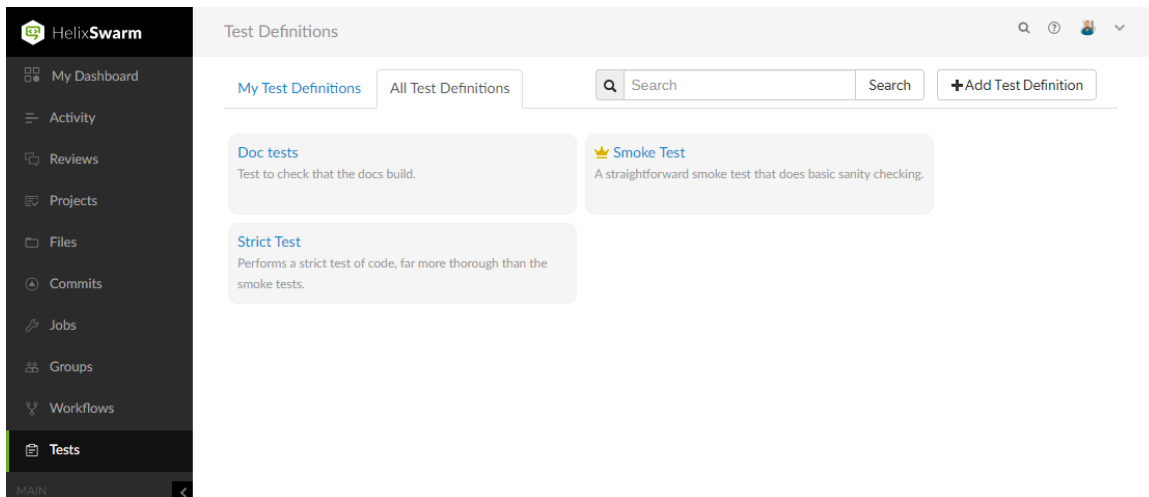
- Swarm tests can be created by any Swarm user.
- Shared Swarm tests can be viewed by any Swarm user.
- Shared Swarm tests can be added to a workflow by any Swarm user that is authorized to edit the workflow.
- Shared Swarm tests can be added to the global workflow by any Swarm user that is authorized to edit the global workflow.

This section provides an introduction to listing, searching, and viewing tests.

- For instructions on how to add a test, see ["Add a test" on page 544](#).
- For instructions on how to edit a test, see ["Edit a test" on page 551](#).
- For instructions on how to delete a test, see ["Delete a test" on page 551](#).
- For instructions on how to add a test to a workflow, see ["Add a workflow" on page 537](#).
- For instructions on how to add a test to the global workflow, see ["Global workflow" on page 753](#).

Listing tests


To view a list of tests, click **Tests** in the menu.



Logged-in users can choose which tests to display by clicking on the **My Tests** tab or the **All Tests** tab. Anonymous users only see the shared tests, the **My Tests** tab is not available to anonymous users:

- The **My Tests**: tab lists all of the tests that you are an owner of.
- The **All Tests**: tab lists the all of the shared tests, and tests that you are an owner of.

Tip

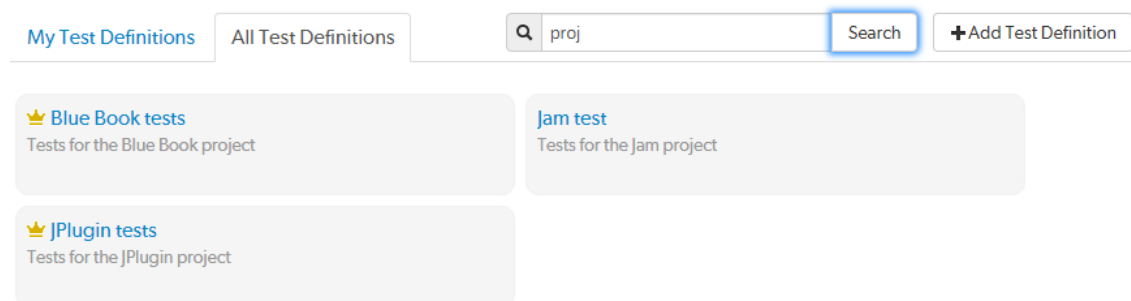
The **Test Owner** icon  is displayed for tests you own.

Search test names and test descriptions

You can search for Swarm test names and for the content of the test description from the tests page:

1. Enter text in the search box.
2. Click **Search**.

Any tests names and descriptions that match the search text are displayed in the tab:



Tip

Switch to the other test tab to display search results for that tab.

Viewing a test that you own

To view test details:

1. Click **Tests** in the menu.
2. Click on the name of the test you want to view.

For information about editing the test, see ["Add a test" on page 544](#).

Note

You can edit a test if you are the owner of the test, or if you have *super* user rights.

Doc tests x

Name

Description

Test to check that the docs build.

Owners

Individuals:

allison.clayb

☒ Shared with others

URL

Timeout(seconds)

Body

status={status}&review={review}&change={change}&update(update)

☒ URL Encoded
☐ JSON Encoded
☐ XML Encoded

☒ Iterate tests for affected projects and branches ⓘ

Header	Value
BasicAuth	YWxpY2U6QUJDRDEyMzQ1Njc4Cg==

[+Add header](#)

Standard arguments (URL and Body fields only) ⓘ

[change] the change number
[status] the status of the shelved change, shelved or committed
[review] the review identifier
[version] the version of the review
[reviewStatus] the Swarm status of the review; needsReview, needsRevision, archived, rejected, approved, approved:commit
[description] the change description of the change (request body only)
[test] the name of the test
[testRunId] the test run id
[projects] the project identifiers of projects that are part of the review
[projectNames] the project names of projects that are part of the review
[branches] the branch identifiers of branches that are part of the review
[update] the update callback URL
[pass] the tests pass callback URL
[fail] the tests fail callback URL

▼ 1 workflow uses this test definition

Strict Workflow

Viewing a shared test that you do not own

Shared tests can be viewed and used by all Swarm users.

To view test details:

1. Click **Test** in the menu.
2. Click on the name of the test you want to view.

To avoid leaking sensitive information to users that do not own the test, configuration details are hidden.

Doc tests x


Name

Doc tests


Description

Test to check that the docs build.

Owners

 Add an Owner

Individuals:



☒ Shared with others

▼ 1 workflow uses this test definition

Strict Workflow

Close

Some details of this test definition are private and cannot be displayed. Only the test definition owners have access to these details.

Notifications

Provided you have entered a working email address for your userid in Helix Core Server, Swarm sends email notifications to you when various events take place. The table below shows the notifications sent if all of the project, group, and user notifications are enabled. This is the default behavior.

Note

The system-wide default notifications can be changed by the system owner. See ["Global settings"](#) on page 664 for details.

Tip

Many of the notifications can be disabled, this allows project, group, and user notifications to be customized to limit the number and type of notifications sent.

- Project notifications are configured on the [project settings tab](#).
- Group notifications are configured on the [group Notifications tab](#).
- User notifications are configured on the [user Notifications tab](#).

Legend:

A = Author	PM = Project member or project member group	Yes = role user/group receives notification
R = Reviewer or reviewer group	PF = Project follower	No = role user/group does not receive notification
M = Moderator or moderator group	AF = Author follower	NA = role not included for this event
GM = Group member	CA = Comment author	

Event	Notification sent to one of these roles...								
	Event by you?	A	R	M	GM	PM	PF	AF	CA
A review is committed in Swarm	No	Yes	Yes	Yes See (3)	Yes See (4)	Yes See (1)	Yes See (1)	No	NA
A review is committed outside of Swarm	No	No	Yes	Yes See (3)	Yes See (4)	Yes See (1)	Yes See (1)	No	NA

Notification sent to one of these roles...									
Event	Event by you?	A	R	M	GM	PM	PF	AF	CA
A change is committed outside of Swarm See "(1) Committed change notifications" on page 403	No	No	NA	Yes See (3)	Yes See (4)	Yes See (1)	Yes See (1)	Yes See (1)	NA
A review is started See "(2) Review start notifications" on page 404	Yes	Yes	Yes	Yes See (3)	Yes See (4)	Yes See (2)	No	No	NA
A reviewer casts a vote	Yes	Yes	Yes See (5)	No	Yes See (4)	No	No	No	NA
A review's state changes	No	Yes	Yes See (5)	No	Yes See (4)	No	No	No	NA
A review's reviewers are changed	No	Yes	Yes See (5)	No	Yes See (4)	No	No	No	NA
A review's files are updated	Yes	Yes	Yes See (5)	No	Yes See (4)	No	No	No	NA

Notification sent to one of these roles...									
Event	Event by you?	A	R	M	GM	PM	PF	AF	CA
A review's automated tests have finished See "(7) Tests have finished" on page 405	No	Yes	Yes	Yes	Yes See (4)	No	No	No	NA
A review's description is changed	No	No	No	No	No	No	No	No	NA
A review comment is created See "Comment notification delay" on page 355	No	Yes	Yes See (5)	No	Yes See (4)	No	No	No	No See (6)
A review comment is edited See "Comment notification delay" on page 355	No	Yes	Yes See (5)	No	Yes See (4)	No	No	No	No See (6)
A committed change comment is created	No	Yes	NA	No	Yes See (4)	No	No	No	No See (6)
A committed change comment is edited	No	Yes	NA	No	Yes See (4)	No	No	No	No See (6)

Event	Notification sent to one of these roles...								
	Event by you?	A	R	M	GM	PM	PF	AF	CA
Someone joins or leaves a review	No	Yes	Yes See (5)	No	Yes See (4)	No	No	No	NA
Someone likes a comment you wrote	No	No	No	No	No	No	No	No	Yes See (5)
You like a comment you wrote	Yes	No	No	No	No	No	No	No	No

Any custom modules added to Swarm may also send notifications.

Important

Email delivery for events related to restricted changes is disabled by default. See ["Restricted Changes"](#) on page 703 for details on how to enable restricted change notifications.

@mention notifications

Users: @mention

Using an [@mention](#) in a review, changelist, or comment causes the referenced userid to receive a notification and be included in any future notifications regarding the associated file or review.

When a comment is added to a job, Swarm sends a notification to users listed in user fields in the job, users [@mentioned](#) in the job description, and the authors of any associated changes.

Note

Typically, you would not receive a notification when you [@mention](#) yourself. [@mentions](#) themselves do not trigger notifications; they inform who receives notifications. See ["notify_self"](#) on page 605 for details on how to enable self notification.

Groups: @@mention

Using an [@@mention](#) in a review, changelist, or comment causes the referenced groupid members to receive a notification and be included in any future notifications regarding the associated file or review.

When a comment is added to a job, Swarm sends a notification to groups listed in group fields in the job, groups @@mentioned in the job description, and the authors of any associated changes.

Note

- **Group mailing list enabled:** notifications are sent to the group email address.

When a group mailing list is enabled it overrides the group member's preference set for the `notify_self` configurable in the `SWARM_ROOT/data/config.php` file. So even when the `notify_self` configurable is set to false the group member will receive an email notification.

There is a caveat that if a user is a group member and the same user is added individually to the review and has set the `notify_self` configurable to true, the user receives two email notifications.

- **Group mailing list disabled:** notifications are sent to the group members individual email addresses.

(1) Committed change notifications

Notifications for committed changes are sent by default, but can be disabled or require users to opt-in. Please see the [notification configuration](#) and [project settings overview](#) for more details.

When committed change notifications are configured for opt-in, you need to copy the configuration's special depot path to the **Reviews:** field in your user spec within Helix Core Server. Once you have done so, Swarm can send you committed change notifications for any change that matches a depot path specified in the **Reviews:** field.

Update Reviews with p4

1. Begin editing your user spec:

```
$ p4 user
```

2. Edit the **Reviews:** field to include the special depot path, plus any other paths you want to receive notifications for.
3. Save the spec.

Update Reviews with P4V

1. Select **Connection > Edit Current User**.
2. Edit the **Reviews:** field to include the special depot path, plus any other paths you want to receive notifications for.
3. Click **OK**.

(2) Review start notifications

By default, notifications are sent when a review is started, but can be disabled for a project, group, or user.

Note

If a user or group is added as a reviewer when a new review is created, the reviewer will be notified that the review has been started. This initial notification cannot be muted.

(3) Moderator notifications

Moderators and moderator groups are associated with specific project branches. Moderators receive notifications for reviews and commits against the branches they are associated with.

(4) Group member notifications

Groups can be members of a project, moderators of a project branch, and reviewers of a review. The notifications that group members receive will depend on, the role the group has and the group notification settings that are configured for that group. Notifications received by group members also depend on whether the group has a group mailing list enabled, in this case additional notification settings are available. See *Group mailing list enabled* below for details.

Group mailing list disabled

Group members can be notified when a member of the group starts a review. Group members can also be notified when a change is committed by, or on behalf of, a changelist owner who is also a member of the group. These two notifications can be individually selected as required. See ["Add a group" on page 491](#) for details. Notifications are sent to the group email address.

Notifications are sent to the group members individual email addresses.

Group mailing list enabled

Group members can be notified when a member of the group starts a review. Group members can also be notified when a change is committed by, or on behalf of, a changelist owner who is also a member of the group. These two notifications can be individually selected as required.

When a group mailing list is enabled it overrides the group member's preference set for the `notify_self` configurable in the `SWARM_ROOT/data/config.php` file. So even when the `notify_self` configurable is set to false the group member will receive an email notification.

There is a caveat that if a user is a group member and the same user is added individually to the review and has set the `notify_self` configurable to true, the user receives two email notifications.

For more information about `notify_self` configurable, see ["notify_self" on page 605](#).

Additional notifications are available if the group mailing list address is enabled. See ["Add a group" on page 491](#) for details.

(5) Disable notifications

By default, notifications for events in a review are sent to all reviewers. Reviewers can disable notifications during a review to avoid further email notifications. Once notifications are disabled for a reviewer, they can be re-enabled when they are specifically [@mentioned](#); reviewers can disable notifications again after they have been re-enabled. See ["Disable notifications" on page 443](#) for details.

(6) Comment author notifications

By default, notifications are not sent to comment authors, but Swarm can be configured to send notifications of comments to comment authors. See ["notify_self" on page 605](#) for details.

(7) Tests have finished

By default, notifications are sent when:

- Automated tests have failed for a review.
- The first time automated tests pass for a review after a test failure for that review.

Tip

The test notifications email contains the complete URL to the tests.

Log in/Log out

When you are not logged into Swarm, certain features are unavailable to you, such as providing comments to changes or reviews, adding projects, and more.

Log in with a password

1. Click **Log in** on the right of the Swarm header.
2. Type in your username and password, appropriate for the Helix Core Server that Swarm is configured to use.
3. Select the **Remember me** check box if you prefer to stay logged in between browser restarts.

Note

The Helix Core Server can enforce maximum log in times. You may become logged out even if **Remember me** is selected. Swarm administrators can change the maximum log in time, see ["Sessions" on page 700](#) for details.

4. Click **Log in with credentials**.

Log in with SSO

1. If your organization has single sign-on (SSO) configured for Helix Core Server, you can use it to log in through your organization's identity provider.
2. Open Swarm.
3. Enter your **Email** or **Username**.
4. Click **Log in with SSO**.

You are prompted to log in to the identity provider.

Note

To log in manually to Swarm using your username and password instead, click **Log in with credentials**. If you switch to log in with credentials and then decide to use single sign-on, click **Log in with SSO**.

5. Enter your identity provider credentials.

If your login is successful, the Swarm home page opens. See ["Quickstart" on page 29](#) to get started with Swarm.

Log out of Swarm

1. Click your userid, on the right of the Swarm header.
2. Select **Log Out** from the drop-down menu.

Tip

- If Helix Authentication Service is configured for your Helix Core Server, logging out of Swarm will not invalidate your Identity Provider (IdP) login status. If you try to log back in to Swarm while your IdP status is still valid, you will not be prompted to complete the log in steps.

If **require_login** is also enabled, Swarm will return you to the login and your IdP will automatically log you back in. In this case log out from your Identity Provider page before logging out from Swarm.

- If a custom redirect has been configured by your Swarm administrator, you are logged out of Swarm and then redirected to the URL specified by the administrator.

The custom redirect can be set to any internal or external URL, for example:

- Company intranet, extranet, internet, FTP, or Web-mail page
- Industry news website
- Identity Provider page to invalidate your IdP log in status

require_login

By default, Swarm requires users to log in, which prevents anonymous users from accessing a Helix Core Server via Swarm. Users who have not logged in see a log in page immediately when visiting Swarm:

The steps to log in are identical to using the Log in dialog.

Swarm administrators can disable [require_login](#) to allow anonymous users to see commits, reviews, etc.

Note

service and *operator* users are not permitted to login. See [User types](#) in the *Helix Core Server Administrator Guide*.

Notable minor features

Quick URLs

Swarm handles URLs intelligently to reduce the amount of information you need to enter to quickly locate what you are looking for:

1. Enter a URL in the following format:

`https://myswarm.url/<identifier>`

2. Swarm checks the `<identifier>` and redirects you to the first match it finds. Swarm checks for the identifier in following order:
 - [Reviews](#)
 - [Changelists](#)
 - [Depot paths](#)
 - [Projects](#)
 - [Jobs](#)
 - [Users](#)
 - [Groups](#)
 - Depot names

Note

- Changelist and review identifiers must be numeric.
- If you enter an identifier that does not exist for any type of resource, Swarm displays a **Page Not Found** error.

Example usage

Display the latest version of a review

For example, to display review 124: enter the URL `https://myswarm.url/124`, Swarm redirects to `https://myswarm.url/reviews/124`.

Tip

- To display a specific version of a review, enter the full URL including **reviews** and append `/v<n>/` to the URL. For example, to display version 2 of review 124: enter `https://myswarm.url/reviews/124/v2/`
- To display a diff between two versions of a review, enter the full URL including **reviews** and append `/v<n,n>/` the URL. For example, to display the diff between version 2 and 4 of review 124: enter `https://myswarm.url/reviews/124/v2,4/`
- To open a review with a specific tab open, append `#<tab name>` to the end of the URL (`#<tab name>` must be the last item in the URL):
 - **Files tab:** `#files` (default if `#<tab name>` is not specified)
 - **Comments tab:** `#comments`
 - **Activity tab:** `#activity`

Display a changelist

For example, to display changelist 123: enter the URL `https://myswarm.url/123`, Swarm redirects to `https://myswarm.url/changes/123`.

Tip

To open a changelist with a specific tab open, append `#<tab name>` to the end of the URL:

- **Files tab:** `#files` (default if `#<tab name>` is not specified)
- **Comments tab:** `#comments`

Display the latest version of a file

Enter the URL `https://myswarm.url/depot/alpha/readme.txt`, Swarm redirects to `https://myswarm.url/files/depot/alpha/readme.txt`.

Tip

- To display a specific version of a file, enter the full URL including `files` and append `?v=<n>` to the URL. For example, to display version 2 of the `depot/alpha/readme.txt` file: enter `https://myswarm.url/files/depot/alpha/readme.txt?v=2`
- To open a file with a specific tab open, append `#<tab name>` to the end of the URL (`#<tab name>` must be the last item in the URL):
 - **View tab:** `#view` (default if `#<tab name>` is not specified)
 - **Commits tab:** `#commits`

Display the user profile page for jsmith

Enter the URL `https://myswarm.url/jsmith`, Swarm redirects to `https://myswarm.url/users/jsmith`.

Links in descriptions and comments

When you write a job description, changelist description, review description, or comment, you can link to users, groups, changelists, reviews, and jobs. Swarm will attempt to automatically create links to your references to make it easy to navigate to the specified resource.

In this section:

- "[@mentioning users and groups](#)" on the facing page
- "[Linking to a review](#)" on the facing page
- "[Linking to a review, changelist, or numeric userid](#)" on page 411
- "[Linking to a job](#)" on page 411

@mentioning users and groups

To link to users or groups, use `@mention` to refer to users and `@@mention` to refer to groups when you write a job description, changelist description, review description, or comment. Swarm automatically creates a link for each `@mention` and `@@mention`.

When you start a code review, including a user `@mention` or a group `@@mention` in the changelist description automatically makes them reviewers for that code review. During a code review, including a user `@mention` or group `@@mention` in a comment causes the mentioned user or group to receive [notifications](#) of code review events, even if they are not a member of your project or following you or your project.

Additional option for a user `@mention`:

- `@*mention`: include an asterisk (*) before the `userid` to make the user a required reviewer. See ["Required reviewers" on page 478](#) for details.

Additional options for a group `@@mention`:

- `@@*mention`: include an asterisk * before the groupid to make the group a required reviewer with **All votes required**. See ["Required reviewers" on page 478](#) for details.
- `@@!mention`: include an exclamation mark ! before the groupid to make the group a required reviewer with **One vote required**. See ["Required reviewers" on page 478](#) for details.

Note

- **Group mailing list enabled:** notifications are sent to the group email address.
- **Group mailing list disabled:** notifications are sent to the group members individual email addresses.

Linking to a review

Important

Adding `#review-<reviewid>` to a changelist description is a special case and does not just add a simple link to the review. It replaces all of the files in the review with the files in the changelist. For information about adding a changelist to a review, see ["Add a changelist to a review" on page 682](#).

To make a simple link to the review, use `review-<reviewid>`.

To create a link to a review, add `#review-<reviewid>` to a job description, review description, or comment. Swarm automatically creates a link to the page for the review if a match is found.

For example, when you include `#review-12345` in a review description, Swarm automatically turns that text into a link that, when clicked, displays the Review page for review 12345. The `#` is optional, `review-12345` would also work.

Linking to a review, changelist, or numeric userid

If you add `#<numericid>` to a job description, changelist description, review description, or comment. Swarm automatically creates a link to the page for the numeric id specified if a match exists.

Swarm checks the numeric id specified and links to the first match it finds. Swarm checks for the numeric id in the following order:

- Reviews
- Changelists
- Users

For example, when you include `#12345` in a comment, Swarm turns it into a link to the page for either review 12345, change 12345, or user 12345, depending on which of them exists.

Linking to a job

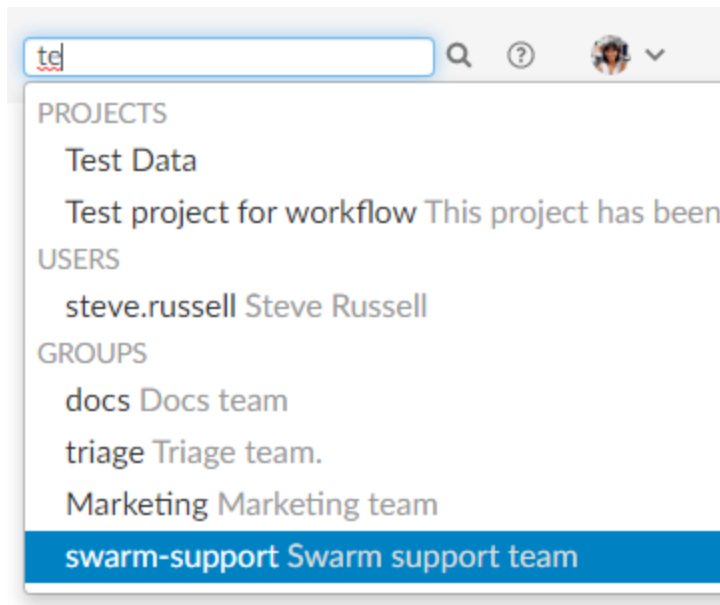
To create a link for a job, add `jobnnnnnn` or `@jobnnnnnn` to a job description, changelist description, review description, or comment. Swarm automatically creates a link to the job page if a match is found.

For example, when you include `@job12345` in a comment, Swarm turns that text into a link that, when clicked, displays the Job page for job12345.

Tip


If your job numbers are not formatted as `jobnnnnnn`, your Swarm admin can configure Swarm to recognize other keywords in addition to `job`. For example, if your job numbers are formatted as `defectnnnnnn`. For instructions on adding a new job keyword to Swarm, see ["Job keywords" on page 625](#).

Search



Swarm can search for users, groups, projects, and file paths. Enter keywords or path elements into the search box in the Swarm header. Navigate the results with the ↑ (up arrow) and ↓ (down arrow) keys. To display the details for a result, press **Enter** or click the search result.

Full-content searching is only available if your Swarm administrator installs the Helix Core Server Search Tool. See ["Search" on page 696](#) for details.



Swarm updates the search results as you type. Some results should appear within a second or two. You may have to wait a few seconds for final results to be incorporated into the results list. When Swarm does not yet have any results, it indicates such.

Jira integration

Swarm ships with a Jira module that can:

- create links to Jira when Swarm displays Jira issue identifiers in changelists, jobs, reviews, etc. Jira links in comments are only supported in the original Swarm review page.
- add links within Jira back to Swarm, for Jira issues associated with reviews or committed changes. These links reflect the current status of associated code reviews.

By default, the Jira module is disabled. For instructions about enabling Jira integration, see ["Jira" on page 552](#).

Avatars

Each event in an activity stream includes an *avatar*, an image that represents the user responsible for the event. Avatars help to visually tie together various events and personalize the history presented in the stream.

Based on the email address entered in a user's or group's Helix Core Server account, Swarm attempts to fetch an avatar from [gravatar.com](#). Swarm selects an avatar from its collection of default avatars if an avatar is not returned from [gravatar.com](#).

Hover your mouse over an avatar to display the fullname of the user or group.

Following

Whenever you see a **Follow** button, for example when you are viewing a project page or user profile, clicking the button causes Swarm to send you notifications whenever there is activity generated by the current resource.

This is useful if, in the case of a project, you are not a project member but want to know what's happening in the project. Or, in the case of a user, you want to see what activity that user generates.

To stop receiving notifications, visit the project page or user profile and click the **Unfollow** button.

Time

Swarm typically displays the time of an event, such as when a file was created, as about *X units ago*. Hover your mouse pointer over a time display to see a tooltip displaying the exact date and time of the event.

Tip

Configure Swarm to display the exact time and date of events by setting **Time Display** to **Timestamp** on your user [Settings tab](#).

Keyboard shortcuts

Swarm provides the following keyboard shortcuts:

- **n** - While viewing ["Diffs" on page 341](#), pressing **n** scrolls to the next difference.
- **p** - While viewing ["Diffs" on page 341](#), pressing **p** scrolls to the previous difference.
- **ESC** - While viewing a dialog, pressing **ESC** closes the dialog.
While entering text into a text area, pressing **ESC** stops text entry.

- **alt + n** or **alt + p** - While viewing a review, pressing **alt + n** or **alt + p** scrolls to the next file or the previous file in the list.

About Swarm

You can discover the version of Swarm you are using:

1. [Log in](#) to Swarm.
2. Click your userid, found at the right of the main toolbar, and select **About Swarm**.

A dialog appears displaying the Swarm version.

Custom error pages

If Swarm encounters an error during processing, such as when a ["Quick URLs"](#) on [page 407](#) is used that points to a non-existent resource, Swarm displays a custom error page.

Short links

Swarm provides a *short link* feature that creates shorter URLs than normal to make sharing specific views within Swarm easier. It is also possible to register or configure an alternate, shorter hostname to have even shorter URLs. See ["Short links"](#) on [page 711](#) for details.

Conceptually, this is identical to <http://tinyurl.com>, or the Twitter feature <http://t.co>, but is restricted to Swarm URLs on a hostname you control.



Swarm displays a *bookmark* button  when viewing files or folders in the depot.

Click the button to display a popup containing the short link. Press **CTRL+C** (on Windows and Linux), or **Command+C** (on macOS), to copy the short link. You can then paste the short link anywhere you'd like to share the current file or folder view in Swarm.

Markdown

You can use Markdown to add text styles to comments, review descriptions, and project overview pages. Swarm can also display Markdown files stored in the Helix Core Server.

Tip

- The [Review Activity](#) tab is rendered in GFM (Git Flavored Markdown).
- The [User](#), [Group](#), and [Project](#) page **Activity** tabs and the [Activity page](#) are rendered in Parsedown.

In this section:

- "Common text styles" below
- "Text styling example" on page 417
- "Markdown in comments and review descriptions" on page 420
- "Markdown in projects" on page 420

Common text styles

Markdown renderers supported

Swarm uses GFM (Git Flavored Markdown) or Parsedown to render text styles. The renderer used depends on where you are entering your text in Swarm:

- Review comments and review descriptions support GFM (Git Flavored Markdown), see [Basic writing and formatting syntax](#)
- Project overview pages, changelist comments, and jobs page comments support Parsedown, see the [Parsedown demo](#)

Common GFM and Parsedown styles

A number of the more common text styles are shown below:

- **bold** and *italics* can be specified with **`**bold**`** or *`*italics*`*.
- Unordered lists can be specified with asterisk (*) markers. Plus (+) and minus (-) signs also work:

```
* This
* That
* Another
```

- Ordered lists can be specified with numbered markers:

```
1. First
2. Second
3. Third
```

- Hypertext links can be specified with [Link text](http://address/page). If you don't need to add anchor text, then a URL in the text without any markup is linkified.
- You can mark inline text as code using ``backticks``.
- You can also mark blocks of code using three backticks on a line, like:

```
```
var text = "Some code text";
```
```

```
alert(text);
...
```

- Headers can be defined using hash (#) marks.

```
# Main heading
## Sub heading
### Lesser heading
```

- Images can be added as well. You can either provide the full Swarm *view* URL, or use a relative Swarm *view* URL to reference something in Swarm.

```
![alt text]
(https://swarm.company.com/view/depot/www/dev/images/jamgraph-
jam.gif "Title Text")
![alt text] (/view/depot/www/dev/images/jamgraph-jam.gif "Title
Text")
```

Tip

To find the Swarm **view** URL for an image file:

1. From Swarm, click on the **Files** tab and navigate to the image file.
2. Click on the image file so that it opens in the **View** tab.
3. Click the **Open** button for the file to open the image in your browser.
4. The URL in your browser address bar is the Swarm *view* URL for the image file.

- Link to a YouTube video using the **Share --> COPY** feature for the YouTube video. For example:

```
[![alt text]
(https://img.youtube.com/vi/wimka8j4XBk/mqdefault.jpg) ]
(https://youtu.be/wimka8j4XBk)
```

- Tables are supported by using pipe (|) separators between columns and colons (:) for justification.

```
| Tables | Look | Like this |
| ----- | ----: | :-----: |
| Left | right | center |
```

- It is also possible to blockquote paragraphs using the greater than symbol (>).

```
> Blockquotes can be displayed like this, using the
> the greater than sign at the start of the line.
```

Normal text resumes here.

Text styling example

The following example shows how Markdown text can be used in a comment but it would display the same for a project overview page.

The source and the final result are shown:

- ["Source text" on the facing page](#)
- ["Rendered view of Markdown text" on page 419](#)

Source text

`*Comments*` can include `**Markdown**` text, which allows basic styles to be applied to the text.

You can have unordered lists, like this:

```
* A line
* Another line
  * A sub list
  * Again
* Back again
```

Or ordered lists:

```
1. First
2. Second
3. Third
```

It is possible to mark text as ``code { like: this }`` or to define a block of text as code:

```
```
var i = 1;
var j = 10;
for (i = 1; j != i; i++) {
 print i * j;
}
```
```

You can also [create hyper links](<http://www.perforce.com>) which point to other places.

Use the backslash character `\` to escape Markdown syntax characters.

You can escape the following characters:

Asterisk `*`

Underscore `_`

Curly braces `\{ \}`

Square brackets `\[\]`

```
Brackets \( \)
Hash \#
Plus \+
Minus \-
Period \.
Exclamation point \!
```

Rendered view of Markdown text



claire.brevia commented 6 minutes ago (edited)

Comments can include **Markdown** text, which allows basic styles to be applied to the text.



You can have unordered lists, like this:

- A line
- Another line
 - A sub list
 - Again
- Back again

Or ordered lists:

1. First
2. Second
3. Third

It is possible to mark text as `code { like: this }` or to define a block of text as code:

```
var i = 1;
var j = 10;
for (i = 1; j != i; i++) {
  print i * j;
}
```

You can also [create hyper links](#) which point to other places.

Use the backslash character `\` to escape Markdown syntax characters.

You can escape the following characters:

Asterisk `*`
Underscore `_`
Curly braces `{ }`
Square brackets `[]`
Brackets `()`
Hash `#`
Plus `+`
Minus `-`
Period `.`
Exclamation point `!`

[Reply](#) · [Edit](#) · ❤


Markdown in comments and review descriptions

Tip

- Try to avoid using styles like headers, images, and tables in comments, they can make comments harder to read this is especially true inline comments.
- **Renderer supported:**
 - Review comments and review descriptions support GFM (Git Flavored Markdown), see [Basic writing and formatting syntax](#).
 - Changelist comments and jobs page comments support Parsedown, see the [Parsedown demo](#).

Markdown content is displayed in comments and review descriptions, but Markdown support is limited to prevent execution of raw HTML and JavaScript content.

If you use Markdown styles in your comments and review descriptions, Swarm renders them when you save the changes. The same markdown is displayed in the Swarm HTML emails.



Markdown in projects

If a project has a **README** markdown file in the root of its **MAIN** branch, that **README** markdown file is displayed as the overview page for the project.

By default, Markdown content is displayed on the project overview page but is limited to prevent the execution of raw HTML and JavaScript content. This can be configured by the Swarm administrator.

- Markdown support for displaying the project overview page can be disabled by your Swarm administrator, see ["Project readme" on page 672](#).
- The level of Markdown support can be configured by your Swarm administrator, see ["Markdown" on page 639](#).

If you need to change which branch is considered MAIN, and therefore from where the **README** is read from, see ["Mainline branch identification" on page 636](#).

Tip

- Valid Markdown file extensions are: **md**, **markdown**, **mdown**, **mkdn**, **mkd**, **mdwn**, **mdtxt**, **mdtext**.

If more than one **README** file is found, Swarm displays the first one it finds based on the order above.

- The Project overview pages support Parsedown, see the [Parsedown demo](#).

Projects/Jplugin

About

A Java plugin for continuous integration.

3 MEMBERS 1 FOLLOWER 2 BRANCHES

OWNERS

MODERATORS

MEMBERS

FOLLOWERS

BRANCHES

MAIN

CANDIDATE

P4 Plugin

Jenkins plugin for a Perforce Helix Versioning Engine (P4D).

Contents

- [Release notes](#)
- [Setup guide](#)
- [Notes page](#)
- [Jenkins page](#)

Requirements

- Jenkins 1.642.3 or greater.
- Helix Versioning Engine 2012.1 or greater.
- Minimum Perforce Protection of **open** for the Jenkins user.
- Review Build feature requires Swarm 2014.2 or greater.

Install

1. Open Jenkins in a browser; e.g. http://jenkins_host:8080
2. Browse to 'Manage Jenkins' --> 'Manage Plugins' and Select the 'Available' tab.
3. Find the 'P4 Plugin' or use the Filter if needed
4. Check the box and press the 'Install without restart' button

If you are unable to find the plugin, you may need to refresh the Update site.

1. Select the 'Advanced' tab (under 'Manage Plugins')
2. Press the 'Check now' button at the bottom of the page.
3. When 'Done' go back to the update centre and try again.

Building

To build the plugin and run the tests use the following:

5 | Code reviews

A code review is a process in which other developers can see your code and provide feedback that can suggest ways to improve the code's structure, performance, maintainability, and interaction with other code.

Benefits

Some of the benefits of code review are:

- Enforcing coding standards: code reviews can catch code that does not meet your team's coding standards. This improves the readability and consistency of your codebase.
- Knowledge and experience sharing: your team can help you learn to code better. This is particularly useful for developers new to the team.
- Early defect detection: small errors can be caught before they become problems later on.
- Code sharing: code reviews spread knowledge of the current codebase, which helps both with maintaining a mental model of the overall project, as well as defending against developer absences.
- Better personal review: knowing that someone might catch a simple coding error often increases the review developers perform of their own code.

Swarm attempts to provide these benefits without adding onerous overhead for developers.

Facilities

Swarm provides the following code review facilities. In the list, the term *author* refers to the person who creates a change to be reviewed, *reviewer* refers to any authenticated Swarm user performing code review tasks, and *required reviewer* refers to a reviewer whose up-vote is required before a review can be approved.

- Authors can request reviews, and can designate reviewers and required reviewers.
- Reviewers can start a code review on existing changes.
- Reviewers can add themselves to a review to indicate that they are participating in the review and sharing responsibility for the review.
- Reviewers can add comments to a changelist, to a specific file in a changelist, or to a specific line in a file, using [Markdown](#) text.
- Reviewers can vote on a review, to indicate their approval or disapproval.
- Required reviewers can prevent a review from becoming approved until they up-vote the review.
- Project branches with assigned moderators limit review approval to one of the moderators.

- Reviews spanning multiple project branches with assigned moderators, limit review approval to one moderator from any one of the branches by default.

Note

Swarm can be configured to require that one moderator from each branch must approve the review. For more information about moderator behavior, see ["Moderators" on page 477](#).

- Reviewers can mark changes as needing revision, approved, rejected, or to be archived for future consideration.
- Reviewers can commit approved changes if necessary.

Workflow

By default, Swarm's code review workflow is basically advisory in nature with very few restrictions imposed on the code review workflow. However, Swarm does provide a number of mechanisms to structure or restrict code review workflows, such as:

- A required reviewer, which is any user designated by a review author, project member or moderator, or is any authenticated user that joins a review and makes their vote required, can prevent reviews from being approved until they up-vote the review. See ["Required reviewers" on page 478](#) for details.
- Branch moderators, when configured for one or more branches in a project, prevent reviews from being approved (or rejected) without their involvement. For more information about moderators, see ["Moderators" on page 517](#).
- Administrators can optionally configure Swarm to prevent reviews with open tasks being approved or, approved and committed, see ["Disable approve for reviews with open tasks" on page 688](#).
- **Workflow:** on by default:
 - Global workflow rules can be configured to enforce a minimum code review workflow on the entire Helix Core Server. For more information on global workflows, see ["Workflow basics" on page 529](#).
 - Global workflow rules can be configured to enforce a minimum code review workflow on projects, and branches that don't have an associated workflow. For more information on global workflows, see ["Workflow basics" on page 529](#).
 - Individual workflows can be configured by users and can be associated with projects, and branches to enforce that code review workflow on them. For more information on workflow, see ["Workflow basics" on page 529](#).

Agile development teams should find sufficient capability within Swarm to make code reviews a regular part of their workflow. Swarm's development team has been using it regularly during development of Swarm. If you have ideas and suggestions for improvement, please [contact us](#).

Models

There are two code review models: pre-commit and post-commit. Which model you use for code reviews with Swarm is up to you.

Pre-commit model

The pre-commit model is possible due to Helix Core Server's *shelving* feature. Shelving enables you to temporarily make copies of your files available to other users without committing the changes into the depot. Shelving can be a very handy way for developers to create a backup, or to handle local workspace changes that might otherwise lose work in progress, without having to commit code that might destabilize a codebase.

Swarm uses the shelving feature in Helix Core Server to manage code reviews. Shelving allows reviewers to easily acquire a copy of the code to be reviewed, and allows updates to the reviewed code prior to submission.

For more information on shelving, see [Shelve Changelists](#) in *Helix Core Command-Line (P4) Guide*.

Post-commit model

The post-commit model can be used if your team's development processes preclude the use of shelving. Code must be committed to the Helix Core Server before code review can begin, which reduces the opportunity to fix problems before, for example, a continuous integration system notices problems. However, code reviews can be started for any existing code regardless of how long it has been committed.

Internal representation

Swarm-managed changelists

Note

Helix Core Server 2020.2 and later: any new file being shelved that has the same content as an existing shelved file refers to the existing archive file instead of creating a duplicate archive file. No Helix Core Server or Swarm configuration is required for this feature.

This Helix Core Server feature automatically reduces the space required for the Swarm-managed shelved review changelists. Swarm creates these changelists for its own internal use. Helix Core Server only updates new shelves, it does not retrospectively update your existing shelves.

A code review consists of one or more shelved changelists that Swarm manages. A shelved changelist is a pending changelist that has a snapshot of its files on a shelf associated with the changelist.

When a review is started, Swarm creates a new changelist that becomes the review changelist. What happens afterwards varies:

- If the review contains uncommitted work (the pre-commit model), Swarm copies the shelved files from the user's changelist that initiated the review into the review's changelist.
- Any time that a user's changelist associated with the review has its shelved files updated, Swarm copies the shelved files into its review changelist and creates an archive changelist. An archive changelist is no different from any other pending changelist with shelved files, but it allows Swarm to provide versioning and diffs within a review.
- If the head version of a review is committed (the post-commit model), the review's changelist is emptied of files.

The review's changelist is never actually committed; this allows the review to be opened later with additional shelved changes.

Important

Swarm's managed review changelists should only be deleted if you are uninstalling Swarm.

Swarm's review changelists maintain the history of a review and all of its feedback. The deletion of a Swarm shelved changelist causes instability and potentially data loss, and represents a scenario that can be very challenging to recover from, even with the engagement of Perforce consultants.

You can display a list of all of the Swarm-managed changelists using the **p4 changelists** command:

```
$ p4 changelists -u swarm
Change 1212285 on 2015/07/31 by swarm@swarm-96017af4-5615-9819-7af1-
6fc1fa537214 *pending* 'Add requirements and instructions'
Change 1212284 on 2015/07/31 by swarm@swarm-96017af4-5615-9819-7af1-
6fc1fa537214 *pending* 'Add requirements and instructions'
...
```

swarm is the userid with *admin*-level privileges within the Helix Core Server that Swarm is [configured to use](#). Use the appropriate userid when you run the **p4 changelists** command.

Swarm-managed workspaces

Whenever Swarm creates a changelist for a review, it uses a client workspace (or just workspace) associated with the configured Helix Core Server userid that has *admin* privileges. Whenever a user commits a change via Swarm's user interface, Swarm uses a workspace associated with that user.

To learn more about workspaces, see [Helix Core Server as a version control implementation](#) in the *Solutions Overview: Helix Core Version Control System* Guide.

The workspaces that Swarm creates and uses live in the `SWARM_ROOT/data/clients` folder.

Inside the clients folder, Swarm maintains a user-specific folder that contains any workspace folders that may be required. Each user-specific folder is named by converting their Helix Core Server userid into hexadecimal to avoid any characters that would be problematic in the filesystem, such as slashes, accents, UTF-8 characters, etc. For example, the folder for the user `steve.russell` would be named `73746576652e72757373656c6c`.

Within the user-specific folder are the folders that become the root of each workspace. Each of these folders is named with a globally-unique identifier (GUID) prefixed with `swarm-`, for example `swarm-438d482b-f107-9a35-c06c-86ac68136b00`. Accompanying each folder is a lock file with the same name plus the `.lock` extension. Finally, the user-specific clients folder contains a management lock file called `manage.lock`.

Here is an example of the folder structure:

```
SWARM_ROOT/
  data/
    clients/
      73746576652e72757373656c6c/
        manage.lock
        swarm-438d482b-f107-9a35-c06c-86ac68136b00/
        swarm-438d482b-f107-9a35-c06c-86ac68136b00.lock
        swarm-8388362a-233d-0cb9-3e90-895eaaa99f6c/
        swarm-8388362a-233d-0cb9-3e90-895eaaa99f6c.lock
      7061756c612e626f796c65/
        manage.lock
        swarm-da7de4b4-0ecb-12c8-1b35-f3e32bb18033/
        swarm-da7de4b4-0ecb-12c8-1b35-f3e32bb18033.lock
```

Here are the steps Swarm takes when it needs to use a client:

1. Convert the current connection's userid to hexadecimal.
2. Check to see whether a user-specific folder exists within `SWARM_ROOT/data/clients`; if not, create the folder.
3. Within the user-specific folder, loop over any existing workspace folders and attempt to lock each in turn:

If a lock is acquired skip to the next step. Otherwise, perform the following procedure.

Create workspace procedure:

- a. Check if the max number of clients for the current user has been reached:
 - If so, wait a short amount of time (50 milliseconds), and start [step 3](#) again.
 - If not, proceed to the next step.
- b. Take a lock on `manage.lock`.

- c. Check if the max number of clients for the current user has been reached:
 - If so, release the `manage.lock`, wait a short amount of time (50 milliseconds), and start [step 3](#) again.
 - If not, proceed to the next step.
- d. Create a new workspace folder using a GUID-based filename, and take a lock on the folder.
- e. Release the `manage.lock` lock.
4. Perform the necessary file operations using the locked workspace folder.
5. Revert the file content within the workspace folder to avoid having constantly growing disk space use.

Note

There may occasionally be stray files left; Swarm is not aggressive about cleaning up.

6. Swarm releases the lock on the workspace folder.

Most users should only require 1-2 workspaces, and those are only required if they commit from Swarm. The *admin* user that Swarm is configured to use should only use one workspace per configured worker.

By default, the number of workspaces that could be active at any given instant is two times the number of configured workers. Since the default worker count is three, Swarm would use at most six workspaces simultaneously.

If the workspace limit is reached, further file processing is blocked until a workspace becomes available. Potentially, this means that users could encounter timeouts. Configuring Swarm to use more workers could solve that issue.

Removal considerations

Administrators might wish to remove Swarm-managed workspaces. There are a few considerations that should be assessed prior to removal:

- Ideally, you should stop the web server (taking Swarm out of service) before removing a Swarm-managed workspace from the Swarm server; this eliminates the risk of removing a workspace that is in use.

If you do not stop the web server first, Swarm may encounter an error during a submit.

- Removal of a Swarm-managed workspace folder does not remove the client spec from the Helix Core Server. Unless the client spec is removed, that workspace effectively becomes orphaned. Orphaned clients are, of themselves, not a big concern as the storage and performance impact is negligible.
- Removal of a Swarm-managed workspace's corresponding client spec in the Helix Core Server can be done. However, **you should never remove a client spec that has associated shelved files.**

Usually, the only client specs that should have associated shelved files belong to the *admin* account that Swarm is configured to use. All other workspaces that may exist for other users are primarily used for submitting changes, and so should not have shelved files associated.

Reviews list

The code review list helps you keep track of code reviews that:

- Have been requested and are awaiting review
- Are underway
- Have been accepted, rejected, or archived

To see all available reviews, click **Reviews** in the menu.

The **Reviews** page lists open and closed reviews for all projects in the Helix Core Server.

Description	Created	State	Tests	Complexity	Comments	Votes
Added Windows to list of build platforms. Allison Clayborne requested a post-commit review 396, 31 minutes ago for jplugin:main	31 minutes ago	Open	0	137	1	0
Tidying up some of the code base. Allison Clayborne requested a pre-commit review 391, 3 days ago for jplugin:main	3 days ago	Open	0	18	1	0
Added some state exception checking. Jack Boone requested a pre-commit review 388, 4 days ago for jplugin:main	4 days ago	Open	0	6	2	2
Update to the 12th entry point, which is defined to measure the latency Jack Boone requested a pre-commit review 385, 4 days ago for maker:master	4 days ago	Open	0	1462	0	1
Update the README to use markdown rather than the old format. Jack Boone requested a pre-commit review 382, 4 days ago for blue-book:main	4 days ago	Open	0	2	0	0

The **Opened** and **Closed** tabs display:

- **Opened** tab: displays a list of all code reviews that have started, are being reviewed, are awaiting revisions, or need to be committed.
- **Closed** tab: displays a list of all code reviews that have been approved and committed, rejected, or archived.


Tip

- The **Opened** and **Closed** tabs display the number of reviews in each tab. Initially this can be an over estimate which Swarm dynamically corrects as more information becomes available to it.
- The **Closed** tab initially displays a – character. The – is replaced by a number when you navigate to the **Closed** tab.

Each review displays the following information:

Description	Created	State	Tests	Complexity	Comments	Votes
Added Windows to list of build platforms. Allison Clayborne requested a post-commit review 396, 31 minutes ago for jplugin:main	31 minutes ago	Open	0	137	0	0

- **Avatar** displays the avatar of the review author, click to view the profile of the author
- **Description** contains:
 - first line of the review description, click to open the review. If the review description is too long it is truncated, click on the ellipsis [...] to expand it in the list page.
 - review author, click to view the profile of the author
 - review type, pre-commit or post-commit
 - review number and version, click to open the review
 - review creation date and time
 - review project branches, click to open the project
- **Created or Last activity**: when the review was created or when the review was last updated depending on the **Result order** button setting.






To change the sort order, click the **Result order** button and select **Created** or **Last activity** from the dropdown menu:

- **Created**: Reviews sorted by when they were created
- **Last activity**: Reviews sorted by when they were last updated

Note

- If the **Result order** button is not displayed, reviews are sorted by either **Created** or **Activity**, as set by your Swarm administrator, see ["Reviews filter" on page 686](#) for details.
- **Result order** button display is a global setting controlled by the Swarm administrator. See ["Reviews filter" on page 686](#) for details.
- When review results are older than 24 hours they are displayed in numerical order within each day.

- **State**: shows the current review state
- **Tests**: shows the test suite state for the review, either tests in progress , tests passed , or tests failed .
- **Complexity**: a traffic light icon and number shows the relative complexity of the review and the total number of lines changed in the review. [Complexity can be configured](#) by your Swarm

administrator but, by default:

- **Red:** ≥ 300 changes
- **Amber:** < 300 and > 30 changes
- **Green:** ≤ 30 changes

Tip

- Review complexity is only calculated for a review when the review is updated and the file content has changed.
- Review complexity is only stored for the current version of a review.

Hover over the complexity icon to display more detailed information:



- **Comments:** displays the number of open (non-archived) comments that are associated with the review. The icon is filled if there are comments on the review.
- **Votes:** displays the number of up votes and down votes for the review. The appropriate vote icon is filled if you have voted on the review, vote icons with only an outline show votes by others.

Note

Hover your mouse over any of the icons to see tooltips.

"Projects" on page 381 and "Groups" on page 373 have their own review lists that display reviews created by their members.

Review filters

Tip

Swarm updates the URL in your browser to reflect the filter options you have selected. This makes it easy to bookmark or share your review list filter settings with the URL.

If you share a URL with **Comment** or **Vote** filters selected, those filters are applied to the user running the URL. This means that their reviews list page will contain different reviews to your reviews list page.



The following filters are available (from left to right):

- **Branch** (only available on the project ["Reviews page" on page 385](#)): a dropdown menu that lets you filter which reviews to display based on the current project branches:

Branch 

- **Select branch:** all reviews for all of the selected branches are displayed. To select multiple branches, click the branches you want to filter by.
 - **Branch search:** an auto-complete search field that allows you to choose one of the branches within the current project. Once specified, only reviews for the selected project branch are displayed.
- **Project** dropdown: filter by the project the review is part of:

Project 

- **My projects:** displays all reviews for all of the projects you are participating in, as a member, owner, moderator, or follower.
- **Select project:** displays all reviews for all of the selected projects. To select multiple projects, click the projects you want to filter by.
- **Project search:** an auto-complete search field that allows you to choose one of the projects defined in Helix Core Server. Once specified, only reviews for the selected project are displayed.

Tip

Swarm creates a projectid based on the project name whenever a project is created. For example, a project called **Test Data** is given a projectid of **test-data**. If the project name changes, the projectid does not change.

If your search finds a match in a projectid but not in the associated project name, the project name for that projectid is returned in the search results. For example, if the project name was changed from **Test Data** to **Sample Data**, a search for *test* will return the **Sample Data** project in the search results.

- **Role** dropdown: filter reviews based on user involvement, options are:

Role 

- **All reviews:** displays all reviews.
- **Reviews I've authored:** displays reviews that you have authored.
- **Reviews I'm participating in:** displays reviews that you are a reviewer of, but not an author of.

- **Reviews I'm an individual reviewer of:** displays reviews that you are an individual reviewer of, but not a group reviewer of, or an author of.
- **Review I've authored or I'm participating In:** displays reviews that you have authored, or are a reviewer of.
- **Authored by:** an auto-complete search field that allows you to choose one of the user accounts defined in the Helix Core Server. Once specified, only reviews authored by the user are displayed.

When you select one of the available options, the list of options updates to match the currently selected filter, and the **Role** dropdown indicates the current filter: **All**, **Author**, **Participant**, **Author or Participant**, **Individual Reviewer**, or **user**.

■ **Reviewers** buttons, (**Opened** tab only):



- **Has reviewers:** displays reviews that have one or more reviewers.
- **No reviewers:** displays reviews that have no reviewers.

■ **States** buttons, (**Opened** tab):



- **Needs Review:** displays reviews that need to be reviewed.
- **Needs Revision:** displays reviews that have been reviewed, but need further revisions before the review can be accepted.
- **Approved:** displays reviews that have been approved, and should be committed.

■ **States** buttons, (**Closed** tab):



- **Approved:** displays reviews that have been approved and committed.
- **Rejected:** displays reviews that have been rejected.
- **Archived:** displays reviews that have been archived.

■ **Tests** buttons (when automated tests are enabled for the associated project):



- **Tests passed** displays reviews that have passed tests.
- **Tests failed:** displays reviews that have failed tests.

■ **Comments** buttons:



- **I have commented on:** displays reviews that you have commented on.
- **I have not commented on:** displays reviews that you are participating in but have not commented on

Note

Filters for commenting only apply to reviews which you are a participant of. Commenting on or voting on a review will automatically add you as a participant. If you leave the review after commenting on it, then this review will not be included in the list.

- **Votes** buttons:



- **I have voted up:** displays reviews that you have voted up.
- **I have voted down:** displays reviews that you have voted down.
- **I have not voted on:** displays reviews that you are participating in but have not voted on.

Note

Filters for voting only apply to reviews which you are a participant of. Commenting on or voting on a review will automatically add you as a participant. If you leave the review after voting on it, then this review will not be included in the list.

- **Clear all:** click to clear all of the filters.
- **Search for a review:** search for a partial match of review description content and an exact match for *reviewid*, *userid*, and *projectid*. The search filter boxes on the **Opened** and **Closed** tabs are independent.

Review display

Tip

The Swarm review page is described in this documentation.

The old review page is now referred to as the classic review page. To use the classic review page, use the **Classic view** toggle switch at the top of the review page.


The Swarm Classic view option will be removed in a future Swarm release.

Help for the classic Swarm review page is available in the Swarm 2021.2 documentation. See [Review display](#).

During a code review, reviewers spend most of their time using the review interface.

By default, when you open the Review page for the first time, the file list panel and the information panel are hidden. Use the **Show more context** buttons to view the file list panel and the information panel.

For more information about the show more context buttons, see ["File diff panel" on page 459](#).



Swarm supports stream specs in your workspace using the [Private editing of streams](#) feature. If a changelist or review contains a stream spec, it will be displayed first in **Files** with the prefix **stream: //**, for example: **stream: //MyStreamDepotName/MyStreamSpecLocationName**. A changelist/review can only contain one stream spec.

The review interface is very similar to the [changelist interface](#); and provides largely the same functionality, but has several notable differences that are described in the following sections.

Jump directly to a specific area of the review page using the following links:

- ["Review summary" on the next page](#)
- ["Review state" on page 436](#)
- ["Vote buttons" on page 436](#)
- ["Change state button" on page 437](#)
- ["Review actions button" on page 437](#)
 - ["Add change" on page 438](#)
 - ["Download zip" on page 439](#)
 - ["Change the review author" on page 440](#)

- "Obliterate a review" on page 440
- "Join or leave a review" on page 441
- "Mark all comments for the review as read" on page 442
- "Mark all comments for the review as unread" on page 442
- "Refresh projects" on page 442
- "Disable notifications" on page 443
- "Deployment status" on page 443
- "Review description" on page 443
 - "Edit description" on page 443
 - "Description comments" on page 444
 - "Jobs" on page 444
- "Information panel" on page 445
 - "Blocking Approval or Blocking Commit" on page 445
 - "Reviewers" on page 447
 - "Tasks" on page 450
 - "Tests" on page 451
- "Files tab" on page 453
 - "Select review versions to view" on page 454
 - "File list header" on page 455
 - "File list" on page 457
 - "File stack" on page 457
 - "File content panel header" on page 458
 - "File diff panel" on page 459
- "Comments tab" on page 463
- "Activity tab" on page 464

Review summary

The Review ID is displayed at the top of the review page. A summary of the review is displayed below the review ID, and contains:

- **Review author** avatar and name: The avatar and name of the review author. Hover over the avatar to see the name of the review author and hover over the name to see the ID of the review author. Click on the avatar or name to go to the profile of the review author. See "[Viewing another user's profile](#)" on page 371.

- **Review type:** Displays the review type, reviews can be [pre-commit](#) or [post-commit](#).
- **Review ID:** The unique number used to identify the Swarm review.
- **Review raised:** When the review was requested.
- **Version:** The version of the review being viewed and the total number of review versions.
- **Project branch:** The project branch the review files are in, click to go to the project page.
- **Last updated:** When the review was last updated.

Review state

The state of the review is displayed in the review description header. A review can be in one of the following states:

- **Needs review:** The review has started and the changes need to be reviewed.
- **Needs revision:** The changes have been reviewed and the reviewer has indicated that further revisions are required.
- **Approved:** The review has been approved. The changes may need to be committed.
- **Rejected:** The review has been completed. The changes are undesirable and should not be committed.
- **Archived:** The review has been completed for now but it is not rejected, or approved. The review has been filed away in case it is needed in the future.

For information about review states, see ["States" on page 485](#).

Vote buttons

The vote buttons are in the review description header. Click a vote button to vote the review up or down. If you have already voted on the review, the vote button you clicked is highlighted. Clicking the button again, clears your vote. If you are not a member of a review, voting on the review adds you as a reviewer.

Stale votes

When a review is updated, if the review's list of files, file content, or file-types changes, any votes cast on the review become *stale*. The vote counts are reset, and the vote indicators become muted.

If you hover your mouse over a reviewer with a stale vote, a tooltip appears displaying the *userid*, how they voted, and on which version of the review; each version is represented as a point on the ["Select review versions to view" on page 454](#).


Change state button

The **Change state** button is also used to change the state of the review. Click to select a new state from the dropdown menu. State change options are only displayed if you are authorized to make the state change:

- **Needs revision:** Select to request changes to the files in the review.
- **Needs review:** Select to request further review of the changes.
- **Approve** (only available if the voting requirements for the review are satisfied. For information on voting requirements, see ["Required reviewers" on page 478](#)): select to approve the review.
- **Commit** (only available for [pre-commit](#) reviews that have been approved): Select to commit the review.
- **Approve and commit** (only available for unapproved [pre-commit](#) reviews when the voting requirements for the review are satisfied. For information on voting requirements, see ["Required reviewers" on page 478](#).): Select to approve and commit the review in a single step. See ["Approve and commit" on page 489](#).
- **Reject:** Select to reject the review.
- **Archive:** Select to archive the review.

For information about changing the state of a review, see ["States" on page 485](#).

Review actions button


The **Review actions**  button gives you access to more actions you can perform on the review. Not all of these actions are available to you, the actions available to you depend on your role on the review, the review type, and what is enabled on your Swarm system:

- ["Add change" on the facing page](#)
- ["Download zip" on page 439](#)
- ["Change the review author" on page 440](#)
- ["Obliterate a review" on page 440](#)
- ["Join or leave a review" on page 441](#)
- ["Mark all comments for the review as read" on page 442](#)
- ["Mark all comments for the review as unread" on page 442](#)
- ["Refresh projects" on page 442](#)
- ["Disable notifications" on page 443](#)
- ["Deployment status" on page 443](#)

Add change

Once a review has been started you can add a changelist to the review. It can be useful to add changelists to an existing review. For example, if follow up changes are made to files in a review or if you need to group a number of changelists under a single review.

By default, for a pre-commit or a post-commit review, when a changelist is updated, all of the files in the review are replaced with the files in the changelist you are adding to the review. See ["Replace review with a pending changelist" on page 474](#).

The **Add change** option in the **Review actions**  button is used to add a changelist to an existing review. The option is not available if the review is in a state that is protected from change by the review workflow rules. For details of the **On update of a review in an end state** rule, see ["Workflow rules" on page 526](#).

Tip

You can also add a changelist to a review directly from the changelist description. See ["Review creation and modification outside of Swarm" on page 483](#).

The options available for the **Add change** option depend on whether the review is pre-commit or post-commit:

■ Pre-commit reviews:

- **Append pending changelist:** when you add a pending changelist to a review, the files in the changelist are appended to the existing files in the review. See ["Append a pending changelist to a review" on page 473](#).

Note

A review can only contain 1 stream spec. If you append a changelist with a stream spec to a review that already contains a stream spec, the spec in the changelist replaces the original one in the review. If it is a different spec from the original spec in the review, Swarm cannot display the diff between them and displays **File content unchanged**.

- **Replace with pending changelist:** when you add a pending changelist to a review, all of the files in the review are replaced with the files in the changelist you are adding to the review. See ["Replace review with a pending changelist" on page 474](#).
- **Replace with committed changelist:** when you add a committed changelist to a review, all of the files in the review are replaced with the files in the changelist you are adding to the review. See ["Replace review with a committed changelist" on page 475](#).

Note

If you replace a pre-commit review with a committed changelist, the new version of the review will be a post-commit review.

■ Post-commit reviews:

- **Replace with pending changelist:** when you add a pending changelist to a review, all of the files in the review are replaced with the files in the changelist you are adding to the review. See ["Replace review with a pending changelist" on page 474](#).

Note

If you replace a post-commit review with a pending changelist, the new version of the review will be a pre-commit review.

- **Replace with committed changelist:** when you add a committed changelist to a review, all of the files in the review are replaced with the files in the changelist you are adding to the review. See ["Replace review with a committed changelist" on page 475](#).

Tip

When the content of a review is changed, Swarm checks to see which branches are in the new version of the review:


■ If a new branch was added to the review:

- Default reviewers on the new branch are added to the review.
- Moderators from the added branch become moderators for the review alongside the existing moderators.
- **Only if workflow is enabled:** if the new branch is associated with a workflow, the [workflow is merged](#) with the existing workflow. The most restrictive workflow is used for the review.

■ If a branch is no longer part of the review:

- Reviewers for the review are not changed.
- Moderators from the removed branch no longer moderate the review.
- **Only if workflow is enabled:** if the branch was associated with a workflow, the branch workflow is removed from the review.

Download zip

The **Download zip** option is available from the **Review actions**  button and is used to download a ZIP archive containing all of the files in the review. The file revisions of the downloaded files are the file revisions in the most recent review version selected in the review version selector. See ["Select review versions to view" on page 454](#).

Note

The **Download zip** option is not displayed if the zip command-line tool is not installed on the Swarm server. For information about installing, and configuring the zip command-line tool, see [Zip archive](#).

When you select the **Download zip** option, Swarm performs the following steps:

1. Scans the files/folders:
 - Checks that you have permission to access their contents, according to the Helix Core Server protections.
 - Checks that the total file size is small enough to be processed by Swarm.
2. Syncs the file contents to the Swarm server from the Helix Core Server.
3. Creates the ZIP archive by compressing the file content.
4. Starts a download of the generated ZIP archive.

Note

- You might not see all of the above steps; Swarm caches the resulting ZIP archives so that repeated requests to the same files/folders can skip the sync and compress steps whenever possible.
- If an error occurs while scanning, syncing, or compressing, Swarm indicates the error.


Change the review author

Note

By default you cannot change the author of a review, this option must be enabled by your Swarm administrator. See ["Allow author change" on page 687](#) for details.

If the author of a review is no longer available, or ownership of a review is passed to a different developer, it is useful to be able to change the author of that review.

To change the author of a review:

1. Click the **Review actions**  button.
2. Select **Change author**.
3. Select the new review author from the user list.
4. Click **Save** to change author of the review.

Obliterate a review

Note

- By default, you must be a user with *admin* or *super* user rights to obliterate a review.
- **Optional:** Swarm can be configured to allow users to obliterate reviews that they have authored. This can be configured by your Swarm administrator. See ["Allow author obliterate review" on page 688](#).


Obliterate is used to permanently delete reviews that have been created by mistake. For instance, if a review is associated with the wrong changelist, or a review contains sensitive information that should not be openly available.

For information on what happens to a review when it is obliterated, see ["When you obliterate a review" on page 669](#).

Important

Obliterate must be used with care, the review and all of its associated metadata are permanently deleted. An obliterated review cannot be reinstated, not even by Perforce Support.


To obliterate a review:

1. Navigate to the review.
2. Click the **Review actions**  button and select **Obliterate Review**.
3. Click **Yes** on the confirmation dialog to complete the obliterate action.
4. The review is obliterated.

Join or leave a review

The option available depends on whether you are a reviewer on the review or not.

Join a review

1. Open the review you want to join.
2. Click the **Review actions**  button.
3. Select **Join review** from the dropdown menu. Alternatively, you can **Vote Up** or **Vote Down** the review.


Your name is added to the **Reviewers** list in the **Information panel**, and you are now a reviewer.

Leave a review

Tip

You cannot leave the review if you are a retained default reviewer on the review or you are a member of a group that is a reviewer on the review

1. Log in, if you have not already done so.
2. Open the review you want to leave.

3. Click the **Review actions**  button.
4. Select **Leave review** from the dropdown menu.


Your name is removed from the **Reviewers** list in the **Information panel**, and you are no longer a reviewer.

Mark all comments for the review as read

When you mark a comments as read, the comments are rolled up into single lines to save space. The read flag is remembered independently for each user.

Tip

Mark all comments read will mark all the comments as read including the description comments.


1. Click the **Review actions**  button.
2. Select **Mark all comments read** from the dropdown menu.
3. All of the comments in the review are rolled up into single lines to save space.

	Steve Russell commented on review 644 (version 1) //projects/mercury/dev/src/api/db.c, line 51 - 7 months ago	read 
	Allison Clayborne commented on review 644 (version 1) //projects/mercury/dev/src/api/db.c, line 105 - 7 months ago	read 
	Allison Clayborne commented on review 644 (version) - 10 days ago	read 
	Steve Russell commented on review 644 (version 2) //projects/mercury/dev/src/api/db.c, line 67 - 3 days ago (edited 3 days ago)	read 
	Jack Boone commented on review 644 (version 2) - about 15 hours ago	 read 

Mark all comments for the review as unread

Tip

Mark all comments unread will expand the content of all the comments including the description comments.

1. Click the **Review actions**  button.
2. Select **Mark all comments unread** from the dropdown menu.
3. The content of all of the comments in the review are expanded.

Refresh projects

When a review is created or updated, Swarm checks to see what projects the review files are associated with and links the review to the projects it finds. If at a later date a new project is created or a new branch is added to a project, Swarm does not check to see if existing reviews are associated with new projects or project branches.

The **Refresh projects** option is used to check if the review files are associated with any projects created or updated after the review was last updated. If any projects are found that are not already associated with the review, the review is linked to them.



Disable notifications

When you become a review participant, by joining the review or being [@mentioned](#) in a comment or in the review's description, you receive notifications for any events associated with the review. If you find that the notifications become more of a burden than benefit and you wish to continue being a review participant, you can disable notifications:

1. Click your avatar in the reviewers area to display your reviewer options.
2. Click **Disable Notifications**.

Once notifications are disabled, you no longer receive notifications. However, if you are [@mentioned](#) in a subsequent review comment, you do receive a notification for that comment; regular notifications remain disabled. This approach ensures that you don't miss anything that other reviewers or the review author deems important.

Deployment status

When [automated deployment](#) has been configured for a project, the deployment success  or failure  is indicated in the review's heading. If your deployment program can provide a URL that provides details of the deployment, the indicator becomes linked; click the indicator to see the deployment results.

Review description

The review description is automatically copied from the changelist description when the review is created. When the review description size exceeds the 25% of the page, only the initial part of the review description is displayed. Select the **Expand description** button to view the full review description and select the **Collapse description** button to hide a part of the review description.

Edit description

Click the **Edit Description**  link in the review description to update the description to reflect any updates have been made during the review.

Tip

- Markdown content is displayed in the review description, but Markdown support is limited to prevent execution of raw HTML and JavaScript content. For information about Markdown, see "[Markdown in comments and review descriptions](#)" on page 420

- If you use Markdown styles in your review description, Swarm renders them when you update the description.

Update pending changelist checkbox

Only available if you are editing the description of a pre-commit review and you are the original author of the changelist that created the review.

Select the **Update pending changelist** checkbox in the **Edit Description** dialog to also apply your review description changes to the original changelist description.

Description comments

You can add comments to the description of the review.

1. Click **Comments (n)** (where **n** is the number of comments that already exist).
2. Click **Add a comment**.
3. Add your comment in the text area.
4. Click **Post**.

Tip

To hide the description comments, click **Comments n** (where **n** is the number of description comments that exist). To display the comments again, click **Comments n**.

Jobs

A list of jobs that this change *fixes*, if any.

Adding a job to the review

Swarm does not provide the ability to create new jobs in the Helix Core Server, but jobs can be added to changelists or reviews:

1. Click **Jobs (n)** (where **n** is the number of jobs that are already part of the review).
2. Click **Add a job**.
3. Scroll through the available jobs, or enter job search criteria to search available jobs.
For more information on job search criteria, see [Jobs](#) in *Helix Core Command-Line (P4) Guide*.
4. If you find the job you want to add, click it to highlight it and then click **Select** to add it.

Note

If you attempt to add a job to a review that affects a single project, Swarm applies the project's job view filter to display only jobs that affect the project. It is not currently possible to expand the filter to include jobs outside of the project.

Unlinking a job from the review

Swarm does not provide the ability to delete jobs from the Helix Core Server, but jobs can be unlinked from reviews:

1. Click **Jobs (n)** (where **n** is the number of jobs that are already part of the review).
2. Click **X** next to the job you want to unlink.

The job is immediately unlinked from the review.

Send All Notifications

Comment notifications are delayed by default, click the **Send All Notifications (n)** link to manually send all of your delayed notifications for the review.

Where **n** is the number of delayed notifications that will be sent for the review.


For more information about comment notification delay, see ["Comment notification delay" on page 355](#).

Information panel

By default, when you open the Review page for the first time, the file list panel and the information panel are hidden. Use the **Show more context** buttons to view the file list panel and the information panel.

For more information about the show more context buttons, see ["File diff panel" on page 459](#).

Blocking Approval or Blocking Commit

Shows the total number of items that are blocking approval or commit for a review. Click  to see details about the items that are blocking approval for a review. Once a review has been approved, the **Blocking Approval** section is not displayed.

For pre-commit reviews, if a user adds a new requirement to an approved review, then that requirement is displayed in the **Blocking Commit** section. For example, if a project initially has the **Minimum up votes** requirement set to one and this requirement is satisfied, and the review is in approved state. Now, if the **Minimum up votes** requirement changes to two then this requirement is displayed in the **Blocking Commit** section. Only when this requirement is satisfied the user can do a commit.

The Blocking Approval or Blocking Commit list consists of:

- **Pending Comment Tasks:** Provides a list of pending review comments that are flagged as *tasks*. The pending comment tasks are shown only when the `disable_approve_when_tasks_open` configurable is set to true in the `SWARM_ROOT/data/config.php` file.

Note

If Swarm is configured to prevent approval of reviews with open tasks and a review has open tasks, the **Approve**, and **Approve and commit** options will not be available for the review. This option is configured by an administrator. See ["Disable approve for reviews with open tasks" on page 688](#).

To approve, or approve and commit a review with open tasks, you must address the tasks first and then set them to **Task Addressed**, or **Not a Task**. See ["Set a task to Task addressed or Not a task" on page 352](#) for details.

For more information about approving a review with pending tasks, see ["Approve a review with open tasks" on page 354](#).

- **Required Reviewers:** Provides a list of all the Individual and group reviewers that are required to vote on the review. A review can not be approved or committed until the required reviewers have voted up on the review.

A required ["Individual reviewer" on page 449](#) has a star badge over their avatar. A required ["Group reviewer" on the next page](#) either has a star badge or a star badge with a 1 over their avatar.

- **Blocking Tests:** Provides a list of workflow tests that have **Failed** or have an **In progress** status that is blocking approval for a review. A review can not be approved or committed until the review author fixes all the failed tests. For more information about a failed test, see ["Tests" on page 451](#).
- **Minimum Up votes:** Provides the number of **Minimum Up votes** required for each project that you have access to. This includes the maximum number of **Minimum Up votes** required for all branches in a project blocking approval or commit for a review. For more information about **Minimum up votes** for a project see, [General Settings tab](#). For more information about **Minimum up votes** for a project branch, see [Branches tab](#).

For [private project](#), the maximum number of **Minimum Up votes** for all private projects and their branches that is blocking approval for a review is displayed.

- **Moderator Votes:** Provides a list of branch moderators that are blocking approval or commit for a review. Depending on the configuration of `moderator_approval` configurable in the `SWARM_ROOT/data/config.php` file, the branch moderator approval is required. Ensure that the branch moderator approves the review and not just up votes it.


By default, when a review spans multiple branches that have different moderators, only one moderator from any one of the branches needs to approve the review.

Swarm can be configured to require that one moderator from each branch must approve the review, this is a global setting. If a moderator belongs to more than one of the branches spanned by the review, their approval will count for each of the branches they belong to. For instructions on how to configure moderator behavior, see ["Moderator behavior when a review spans multiple branches" on page 692](#).

Reviewers

The list of reviewers for the review is displayed in the **Information panel** to the right of the review description.

The Reviewers list consists of:

- **Edit**  button (if enabled): click to edit the reviewers for the review. See ["Edit reviewers" on page 470](#).
- **Up** vote and **Down** vote count: indicates the number of up votes and down votes the review has.
- **Groups**: lists groups that are reviewers for the review. When at least one person in the group has voted, a vote icon is displayed to the right of the group indicating whether the group, as a whole, has voted up or down. Click on the group to see who has voted, and how they have voted. See ["Group reviewer" below](#).
- **Individuals**: lists individuals that are reviewers for the review. When an individual has voted on the review, a vote icon is displayed to the right of the user indicating whether they voted up or down. For more information about individual reviewers, see ["Individual reviewer" on page 449](#).

Note

By default, reviewer group members are not displayed in the *Individuals* area of the reviews page when they interact with a review (vote, comment, update, commit, archive, etc.). This avoids overloading the *Individuals* area with individual avatars if you have large reviewer groups.

An exception to this behavior is when a member of a reviewer group is also an individual [required reviewer](#), in this case their avatar will be displayed in the *Individuals* area.

See ["Expand group reviewers" on page 691](#) for details on displaying reviewer group members when they interact with a review.

Group reviewer

If you are a member of a group that is a reviewer on the review, click a vote button in the review description header to vote up or down. Your vote is registered for the group, and is also displayed to the right of your name in the **Reviewers** list. Click on the group in the **Information panel** to see how individual members have voted on the review.

Changing group vote settings

You must be the review author, a project member, a project moderator, or a user with super privileges to change the group settings.

To change the group settings for the review:

1. Click the **Reviewers** list **Edit**  button.
2. Click the **Groups** tab.
3. Click the dropdown arrow to the right of the group you want to change.
4. Select one of the following for the group:

Optional Default:

- If any group member votes down the review, an icon is displayed to the right of the group to indicate the group has voted the review down.
- If at least one group member votes up the review and **no** members of the group vote down the review, an icon is displayed to the right of the group to indicate the group has voted the review up.

Require one: Indicated by a star badge with a 1 on the group avatar.

- If any group member votes down the review, an icon is displayed to the right of the group to indicate the group has voted the review down.
- If at least one group member votes up the review and **no** members of the group vote down the review, an icon is displayed to the right of the group to indicate the group has voted the review up. See [Required reviewers](#) for details.

Require all: Indicated by a star badge on the group avatar.

- If any group member votes down the review, an icon is displayed to the right of the group to indicate the group has voted the review down.
- If all the group members vote up the review and **no** members of the group vote down the review, an icon is displayed to the right of the group to indicate the group has voted the review up. See [Required reviewers](#) for details.

Remove from review: Removes the group from the review.

Note

If the review includes content that is part of a project or branch with [Retain default reviewers](#) enabled, the following restrictions apply to the review:

- The voting option for a retained default reviewer can only be changed to a stricter level, you cannot reduce the voting level.
- Retained default reviewers cannot be removed from the review.


Individual reviewer

When an individual reviewer has voted on a review, a vote icon indicating whether they voted up or down is displayed to the right of their name in the **Reviewers** list. [Required reviewers](#) have a star badge over their avatar.

Changing an individual 's vote settings

You must be the review author, a project member, a project moderator, or a user with super privileges to change settings for an individual reviewer.

To change an individual's settings for the review:

1. Click the **Reviewers** list **Edit**  button.
2. Click the **Users** tab if it isn't already displayed.
3. Click the dropdown arrow to the right of the user you want to change.
4. Select one of the following for the user:
 - **Optional** Default
 - **Required:** Indicated by a star badge over the user's avatar.
 - **Remove from review:** Removes the user from the review.

Note

If the review includes content that is part of a project or branch with [Retain default reviewers](#) enabled, the following restrictions apply to the review:

- The voting option for a retained default reviewer can only be changed to a stricter level, you cannot reduce the voting level.
- Retained default reviewers cannot be removed from the review.

Review approval

A review can be approved when the following requirements are met:

- All of the [Required reviewers](#) on the review have voted up.
- The **Minimum up votes** on the review has been satisfied for **each** of the [projects](#) and [branches](#) the review spans.

Use the [Change state](#) button to approve the review.

Tip

- If the review is moderated, only the moderator can approve the review. For more information about review moderation, see ["Moderators" on page 477](#).
- **Automatic approval of reviews:**
 - If the **Automatically approve reviews** workflow rule is enabled for the review's project/branch, the review is automatically approved as soon as the review requirements are satisfied.

- If the review has a moderator, the review will not be automatically approved. The moderator must approve the review manually.

For more information about the **Automatically approve reviews** rule, see "[Workflow rules](#)" on page 526.

Tasks

Comments can be flagged as tasks. For more details on working with tasks, see [Tasks](#).

Note

If Swarm is configured to prevent approval of reviews with open tasks and a review has open tasks, the **Approve**, and **Approve and commit** options will not be available for the review. This option is configured by an administrator. See "[Disable approve for reviews with open tasks](#)" on page 688.

To approve, or approve and commit a review with open tasks, you must address the tasks first and then set them to **Task Addressed**, or **Not a Task**. See "[Set a task to Task addressed or Not a task](#)" on page 352 for details.

A summary of the number and status of comments flagged as tasks is displayed in the **Information panel** to the right of the review.

Note

Archived comments that are flagged as tasks are not included in the summary or the **Tasks** dialog.



- **Red Flag:** Displays the number of open tasks on the review.
- **Green check mark:** Displays the number of addressed tasks on the review.
- **Blue double-check mark:** Displays the number of addressed and verified tasks on the review.

Show Task details  : Click to display a dialog listing all of the tasks associated with the review:

Tasks

Reporter

Status		Reporter	Description
Addressed		Allison Clayborne	I have passed this to QA.
Addressed		Allison Clayborne	We will need to update this for the next revision.
Open		Claire Brevia	This description needs more detail. I feel the lack of detail makes it very hard to review this work. When will you get some time to change this?
Verified		Claire Brevia	This makes no sense, please update

Within the **Tasks** dialog, you can filter the tasks by the **Reporter** (the userid of the user who created the task), and/or by task state using the buttons at the top of the dialog:



- Click the **Red flag** button to display only open tasks (comments that need to be addressed).
- Click the **Green check mark** button to display only addressed tasks (comments that have been addressed).
- Click the **Blue double-check mark** button to display only verified tasks (comments that have been addressed and verified).



To change the state of a task from the **Tasks** dialog, click the task state dropdown for the task and select the new task state.


To view the full comment text for a task, click the ellipses to the right of the task.

Tests










When the review has an associated [workflow](#) with [tests](#) configured on it or when [continuous integration](#) is configured for the project, Swarm displays the test results for the current review version in the information panel:

-  **Passed:** Displays the number of the tests that have passed for the review version.
-  **Failed:** Displays the number of the tests that have failed for the review version.

-  **Manual:** Displays the number of **On Demand** tests that can be run manually for the review version.
-  **In Progress:** Displays the number of tests that are running for the review version.









Show test details  button: Click to view test run information for tests associated with the review version:


Tests

Status	Title		
 Passed	main-codeSniff		Rerun test 
 Passed	main-EsLint		Rerun test 
 Failed	main-Jest		Rerun test 

Messages

Error contacting server


 Passed	main-jsLint		Rerun test 
 Manual	main-library		Run test 
 In progress	main-modules		


If your continuous integration system calls back to Swarm with a URL, the test run is linked to the URL provided. If the test result contains any messages, click the **Show more**  button to the left of the test to view them

When the content of a review is unchanged between review versions, the tests are not rerun because the files have not changed. In this case, the earlier test results are displayed and marked with the word

(Copy). You can manually run these tests if required by using the  **Run test** button.

The following actions are available from the **Tests** dialog for tests called by a workflow:

- **Rerun test:** If a test has completed you can rerun the test, click  to rerun the test for the most recent version of the review.

- **Run test:** If the review workflow has a test set to run **On Demand**, click  to manually run the test for the most recent version of the review.

Tip

- The **Run test** and **Rerun test** buttons are only available if you are logged in to Swarm, viewing the latest version of the review, and the test was called by a workflow.
- The **Rerun test** button is not available for tests that are in progress.
- When you click **Run test** or **Rerun test**, all of the buttons in the list are temporarily disabled while Swarm starts the test and checks if any other tests are in progress.
- The **Rerun test** and **Run test** buttons are only available for tests that are run by a workflow. They are not available for tests that are run because they are configured on a project.


Note

Private projects: if a test for a private project is added to your review because **Iterate tests for affected projects and branches** is selected for the test, Swarm honors the private project's permissions and displays it as **Private project** in the test list to users that do not have permission to view it.

- For information about iterating tests, see ["Iterate tests for affected projects and branches checkbox" on page 546](#).
- For information about private projects, see ["Private projects" on page 387](#).

Files tab

Use the **Files** tab to view the files in the review and to see how they have changed using the Swarm diff view.



Swarm supports stream specs in your workspace using the [Private editing of streams](#) feature. If a changelist or review contains a stream spec, it will be displayed first in **Files** with the prefix **stream: //**, for example: **stream: //MyStreamDepotName/MyStreamSpecLocationName**. A changelist/review can only contain one stream spec.

Select review versions to view



The review version selectors are used to specify which versions of a review you want to *diff*, they are located in the **Files** tab above the list of files.

- **Left dropdown selector:** the base version for the review is selected by default, base is the revision of the file that was checked out of the depot before it was changed for this review. The current review version in the depot, sometimes called Head can be selected, if there are multiple versions of the review a specific version can be selected.
- **Right dropdown selector:** the latest version of the review is selected by default. If there are multiple versions of the review, a specific version of the review can be selected.

Tip

If a review consists of one or more Swarm-managed changelists. When comparing versions of a review, Swarm is showing any differences between the selected versions, not the review author's personal changelist. See ["Internal representation" on page 424](#) for details.

File list header

The file listing header contains the following elements:

Review version and common file path summary

- Comparing the files in the latest version of the review to Base or Head:

#3: Change 753398 shelved into //depot/main/swarm/collateral/sphinx/administration

- The current version of the review.
 - Which changelist contains a shelved copy of the review's files.
 - The common path for all of the files in the review.
- Comparing two versions of the review:

#2-3: Changes between shelf 753759 and shelf 753398 into //depot/main/swarm/collateral/sphinx/administration


- Which two versions of the review are being compared.
- Which changelists contain the files being compared.
- The common path for all of the files in both versions of the review.

Total number of read files

Shows the number of files that are marked as read out of the total number of files in a review.

File change summary

The file change icons in the file listing header indicate the number and type of change for the files in the review. For example, the icons below indicate that:




- One file has been added, branched, or imported
- Six files have been edited or integrated
- Zero files have been deleted

Expand the Diff view to full screen



Select the button to expand the diffs view to full screen. Once the diffs view is expanded to full screen, select the button a second time to exit the full screen mode.

Diff actions button

The **Diff actions**  button gives you access to more actions you can perform when viewing a file diff.

- **Expand or collapse all files in a review**

Expands or collapses all the files within a review if one or more files are in the collapsed or expanded state. Depending on whether the files are expanded or collapsed, use either of the following options:

- **Expand all files** (only displayed if all or at least one file is collapsed) Select to expand all the files within a review.
- **Collapse all files** (only displayed if all files are expanded) Select to collapse all the files within a review.

- **Show or hide all in-line comments for all files in a review**

Shows the total number of comments across all files in a review. Depending on whether the inline comments are expanded or collapsed, use either of the following options:

- **Hide in-line comments**: select to collapse all of the inline file comments.
- **Show in-line comments**: select to expand all of the inline file comments.

- **Show diffs in-line or side-by-side for this file**

Depending on whether the diffs are shown in-line or side-by-side, use either of the following options:

- **Show diffs in-line**: this option only works for text files. Select to display the diffs in inline format.
- **Show diffs side-by-side**: this option only works for text files. Select to display the diffs in side-by-side format. The oldest file revision is shown in the left pane, and the newer one is shown in the right pane.

- **Show or hide whitespace and tab characters for all files in a review**

Depending on whether the diffs are shown with whitespace and tab characters or not, use either of the following options:

- **Show whitespace and tab characters** (display of CRLF line endings is only supported in the original Swarm review page): shows whitespace and tab characters for all text files. The tab characters are padded to show the correct alignment of tabs.
- **Hide whitespace and tab characters**: hides whitespace and tab characters for all text files.

- **Show or hide whitespace diffs for all files in a review**

Depending on whether the diffs are shown with whitespace or not, use either of the following options:

- **Show whitespace diffs:** whitespace changes are highlighted, makes it easier to identify changes in file types where whitespace is important.
 - **Hide whitespace diffs:** whitespace changes are not highlighted, this makes it easier to see the important changes in file types where whitespace changes are not important.
- **Navigate to the next file**
Select this option to open the next file in the Diff view or select **Alt+N** on your keyboard.
 - **Navigate to the previous file**
Select this option to open the previous file in the Diff view or select **Alt+P** on your keyboard.

File list

Important

- For a Swarm review containing large number of files, using **Ctrl+F** to search for a file does not work. You can use the file filter to search for a specific file in the review. To manually select a file from the file stack, scroll through the Review page.
- By default, when you open the Review page for the first time, the file list panel and the information panel are hidden. Use the **Show more context** buttons to view the file list panel and the information panel.

For more information about the show more context buttons, see ["File diff panel" on page 459](#).

The file list panel contains the following elements:

File filter

Filter files by file name, file extension, or folder name to find the files you are interested in. You can see the diff view for the matched search result in the File stack section.

File list resize

Change the file list width by dragging the split bar left or right. Alternatively, click the split bar to collapse or expand the file list. If you change the file list width for a review, Swarm saves the width locally for that review. The saved width is used next time you view the review.

View complete file list

Use the scroll bar to view the complete list of files.


File stack

All the files in a review are stacked together to enable you to scroll through each of the files and view their diffs. Select a file to expand the diffs view and select the file a second time to collapse the diff view.

The file stack contains the following elements for each file:

File change type

Each file is marked with a ribbon and an icon indicating whether the file was:



File content panel header



The file content panel header contains the following elements:

File name

The name and revision of the file you are currently viewing.

Total number of file comments

The total number of comments for a file. The comments count includes all inline comments and file-level comments. It does not include the archived comments.


View full context of a file

This button is only displayed for text file diffs. Toggles between displaying only the portions of the file that have changed and the full file.

Open content in a new tab

Click the button to open the selected file in a new browser tab.

Mark file as read

Beside each file in a review is a **Mark file as read** button , which help you keep track of which files you have reviewed. The read flag is remembered independently for each user. If the content of a file is changed in an update to the review, the read flag automatically clears. This is particularly useful when a code review consists of many files.

When clicked, color is added to the button background and the read by badge is added to the file in the **File list**.

If a file has been marked as read, click the button a second time to reset the status to unread.

File diff panel

By default, all files larger than 1 MB in size are truncated. When a file is truncated a message is displayed. Use the "[max_size](#)" on [page 622](#) configurable in the `SWARM_ROOT/data/config.php` file to increase or decrease this limit.


■ Text file diff:




When you view a diff, the changes are highlighted:



- Red indicates lines that have been removed.
- Blue indicates lines that have been modified.
- Green indicates lines that have been added.

For more information on the supported extensions for syntax highlighting in the Review page, see "[Supported syntax highlighting in the Review page](#)" on [page 331](#).

Example inline diff view:

 ./db.c#1


4
 






 42 lines hidden above

```



43      43  * perferendis consequuntur qui quidem. Et et aut est ex.
44      44  */
45      45  void * initVoluptatem() {
46      46      int loop = 1;
47      47
48      48  // Abbatia hodie conubia litora nunc.
49      49  dolor = flag / 1;
50      50  printf("Femina sem bis non dolor.");
51      51  printf("Wisi, mando gladius.");


```





Steve Russell commented on review 644 (version 1) - about a year ago

✓
 
 ...

@allison.clayborne Is this really necessary to spam the console with this?

 Add a comment




 34 lines hidden

```

87      92  popFugiat(aperiam, quo);
88      93  addBeataeDevice(iste, ut);
89      94  }
90      95
91      96
92      97  /**
93      98  * Herbam class tortor justo, augue justo congue. Volutpat a magnis
94      99  * cursus pupula. Aliquet diam peristo ullamcorper cultura. Suspendisse
95     100  * nostra.
96     101  */

```

Example side-by-side diff view:

./db.c#1

4

42 lines hidden above

43

* perferendis consequuntur qui quidem. Et et aut es

t ex.

44

*/

45

void * initVoluptatem() {

46

int loop = 1;

47

48

// Harum quaerat in distinctio vitae adipisci impe

dit distinctio.

49

printf("Dicta voluptas mollitia est placeat au

t.\n");

50

et = tempora + 1;

51

voluptas = atoi(qui)?quas:"eos";

52

runOccaecati(ut, sed);

34 lines hidden

87

popFugiat(aperiam, quo);

88

addBeataeDevice(iste, ut);

89

}

90

91

42 lines hidden above

43

* perferendis consequuntur qui quidem. Et et aut es

t ex.

44

*/

45

void * initVoluptatem() {

46

int loop = 1;

47

48

// Abbatia hodie conubia litora nunc.

49

dolor = flag / 1;

50

printf("Femina sem bis non dolor.");

51

printf("Wisi, mando gladius.");

Steve Russell

commented on review 644

(version 1) - about a year ago

@allison.clayborne Is this really necessary to spam the console with this?

Add a comment

52

// Harum quaerat in distinctio vitae adipisci impe

dit distinctio.

53

printf("Dicta voluptas mollitia est placeat au

t.\n");

54

et = tempora + 1;

55

voluptas = atoi(qui)?quas:"eos";

56

runOccaecati(ut, sed);

57

34 lines hidden

92

popFugiat(aperiam, quo);

93

addBeataeDevice(iste, ut);

94

}

95

96

■ Image file diff:

Image types that are natively supported by your browser are displayed with a vertical slider on the image.

Drag the slider to the left to reveal more of the newer image and to the right to reveal more of the earlier one. Alternatively, click on the image to move the slider.

Example image diff:

461










■ **Show more context buttons:**

Sometimes, the concise diff view needs to be expanded to fully understand the context of the change, use the **Show More** and **Show All** buttons to display extra lines around the change:

Tip

The following buttons are not available for stream specs.

- **Show the file list panel**  button (only displayed on the vertical bar next to the file list panel): click to view the file list panel.
- **Show the information panel**  button (only displayed at the bottom of the information panel): click to view the detailed information panel.
- **Show All Lines to Start of File**  button (only displayed for the first change in the file): click to show all of the lines up to the start of the file.
- **Show More Lines for the Code Below**  button: click to show 10 more lines above the change, the extra lines are displayed in the pane below the button.
- **Show Entire Section**  button: click to show all of the lines between the changes that are above and below the button, the two changes and the lines between them are displayed in a single pane.
- **Show More Lines for the Code Above**  button: click to show 10 more lines below the change, the extra lines are displayed in the pane above the button.
- **Show All Lines to End of File**  button (only displayed for the last change in the file): click to show all of the lines down to the end of the file.

Comments tab

The **Comments** tab is used to view all of the comments in the review.

Files **6** Comments **4** Activity Send All Notifications (0)

5 archived comments

Steve Russell commented on review 644 (version 1) //projects/mercury/dev/src/api/db.c, line 51 - 11 months ago ✓ read

Allison Clayborne commented on review 644 (version 1) //projects/mercury/dev/src/api/db.c, line 105 - 10 months ago 👁️ ⋮

```
+ */
+int TalisAliquam(long quandoScelerisqueSem, float count, double litora) {
+ // Peristo velit parco.
+ parco = volutpat / 1;
+ printf("Aquam arcu vehicula faucibus umbra magna.");
```

@steve.russell what do you think of this?

👍

Steve Russell commented 4 months ago 👁️ ⋮

This looks good to me. I am happy for it to be checked in.

👍

Allison Clayborne commented 4 months ago 👁️ ⋮

Thanks for the confirmation.

👍

Add a comment

Drag and drop files to attach them or [Choose your files](#)

Post ☐ Flag as Task ☐ Post and notify (0)

For information about working with comments, see "Comments" on page 348.

Activity tab

The **Activity** tab presents a list of the events on this review.

Files **6** Comments **4** Activity Send All Notifications (1)

Allison Clayborne cleared an issue on review 644 for mercury:dev 6 minutes ago ⋮

This needs another review please

Allison Clayborne archived comment on review 644 for mercury:dev 8 minutes ago ⋮

This description needs more detail.

Allison Clayborne replied to a comment on review 644 for mercury:dev 8 minutes ago ⋮

Thanks for the confirmation.

Steve Russell replied to a comment on review 644 for mercury:dev 9 minutes ago ⋮

This looks good to me. I am happy for it to be checked in.

Events in the activity tab include:

- When the review was started
- When a new reviewer joins the review
- When the review's state changes
- When the review's files are updated
- When a reviewer votes on the review
- When someone comments on the review, or one of its files
- When tests pass or fail, provided continuous integration is configured

Activities

This section describes the major activities that affect code reviews, including starting a review, updating a review, and fetching a review's files.

Start a review

Important

If your Helix Core Server is configured as a commit-edge deployment, and your normal connection is to an edge server, Swarm refuses to start reviews for shelved changes that have not been promoted to the commit server.

Within Swarm, this means that the **Request Review** button does not appear for unpromoted shelved changes. Outside of Swarm, attempts to start reviews for unpromoted shelved changelists appear to do nothing. Ask your Helix Core Server administrator for assistance if you cannot start a review.

An administrator of the Helix Core Server can automatically promote shelved changes to the commit server by setting the configurable `dm.shelve.promote` to 1.

Tip

If your changelist only contains a stream spec and its location in the Helix Core Server is not associated with a Swarm project, the review that you create will not have any default reviewers or workflow rules. To associate the review with a project so that it has default reviewers and obeys the project workflow rules, include a file change from the project path in the changelist when you create the review.

1. Use the Helix Core command-line (P4) or a client to create the shelved or committed changelist.
2. Start a code review by using one of the following approaches:

Use Swarm:

1. Use Swarm to view a shelved or submitted changelist.

Tip

To view a shelved or submitted changelist, use a [Quick URL](#). For example, if your change is **54321**, visit the URL: `https://myswarm.url/54321`.

2. Click the **Request Review** button to request a review of that changelist.

Requesting a review on a shelved changelist uses the [pre-commit](#) model and requesting a review on a submitted changelist uses the [post-commit](#) model.

Use Helix Core command-line (P4):

When you are about to shelve or submit files:

1. Include `#review` within your changelist (separated from other text with whitespace, or on a separate line).

Once the review begins, Swarm replaces `#review` with `#review-12345`, where **12345** is the review's identifier.

Note

The `#review` keyword is customizable. For details, see ["Review keyword" on page 680](#).

2. At this time, you can add reviewers to the code review by using [@mention](#) for users, and [@@mention](#) for groups in the changelist description for each desired reviewer.

If your `@mention` or `@@mention` includes an asterisk (*) before the *userid* or *groupid*, for example `@*userid`, that user or all of the group members become required reviewers. If your `@@mention` includes an exclamation mark (!) before the *groupid*, for example `@@!groupid`, the members of that group become required reviewers but only one member of the group is required to vote. See ["Required reviewers" on page 478](#) for details.

3. Complete your shelve or submit operation.

Warning

If you shelve a changelist and subsequently edit the description to include `#review`, a review is **not started**. You must re-shelve the files after adding `#review`.

Tip

You can also start a Swarm review with P4V, P4VS, and P4Eclipse. See below for details:

- **P4V:** see the [Swarm integration features](#) section of the *Helix Core Visual Client (P4V) Guide*.
- **P4VS:** see the [Managing files](#) chapter of the *P4VS User Guide*.
- **P4Eclipse:** see the [Reviewing changes](#) chapter of the *P4Eclipse User Guide*.

Note

If you are using P4V and its Swarm integration, and you encounter the error **Host requires authentication**, ask your Helix Core Server administrator for assistance. See ["P4V Authentication" on page 588](#) for details.

Update a review

To update a code review, use one of the following approaches:

- For a [pre-commit](#) review that you authored:
 1. Edit the files
 2. Shelf the files

You can repeat these steps as many times as necessary.

Tip

If you want to stop Swarm from updating the review until you have completed your changes, add the work-in-progress tag (**#wip** by default) to the changelist description. For information on using the work-in-progress tag, see ["Work-in-progress tag" on page 393](#).

- For a [post-commit](#) review, or a review where you are not the author:
 1. [Fetch the review's files](#) into a new changelist
 2. Edit the files
 3. Update the changelist's description to include **#review-12345** (separated from other text with whitespace, or on a separate line)
 4. Shelf the changelist's files

Once these steps are complete, further updates involve editing the files, and then shelving the changelist's files.

Warning

If you use an invalid review identifier, it will appear that nothing happens. Swarm is currently unable to notify you of this situation. If the review has not been correctly updated, use the **Add Change** button in the Swarm review heading to add the changelist to the review, see ["Add a changelist to a review" on page 471](#).

Fetch a review's files

First, determine the changelist containing the review's files:

1. Visit the review's page.
2. The current review version's changelist appears in the file list heading:

#3: Change 697707 shelved into //depot/main/swarm

In this example, the changelist is **697707**. You use the identified changelist in place of shelved changelist below.

3. Decide whether you will use [Helix Core command-line \(P4\)](#) or [P4V](#) to fetch the files, and follow the instructions in the appropriate section below.

Using Helix Core command-line (P4)

1. For a **shelved changelist**, use a command-line shell and type:

```
$ p4 unshelve -s shelved changelist
```

2. For a committed changelist, use a command-line shell and type:

```
$ p4 sync @committed changelist
```

Note

Your client's view mappings need to include the changelist's path.

Using P4V

For a shelved changelist:

1. Select **Search > Go To**.
2. Change the select box to **Pending Changelogist**.
3. Type in the *shelved changelist* number and click **OK**.
4. Select the files in the **Shelved Files** area.
5. Right-click and select **Unshelve**.
6. Click **Unshelve**.

For a committed changelist:

1. Select **Search > Go To**.
2. Change the select box to **Submitted Changelogist**.
3. Type in the *submitted changelist* number and click **OK**.

4. Select the files in the **Files** area.
5. Right-click and select **Get this Revision**.
6. Click **Close**.

Deleting shelves

If you have a number of unused shelves, you can delete them if you want to tidy up your workspace.

Important

By default, when you delete files from a shelved changelist, the files are not removed from the associated review.

However, Swarm can be configured to remove files from a review when they are deleted from an associated shelf, see ["Process shelf file delete when" on page 693](#).

Do not use the Swarm user that is configured in the [Swarm configuration file](#) when deleting shelves, or deleting files from shelves. The Swarm logic processes the *shelve-delete* trigger event, if the event is invoked by the Swarm user it is rejected. The delete operations will fail.

To delete a shelf from a changelist without removing the shelved files from the associated review, use [Helix Core command-line \(P4\)](#) or [P4V](#) to delete the entire shelf in one operation as described below. This enables you to delete a shelved changelist in your workspace without affecting the associated Swarm review even if Swarm is configured to remove files from a review when they are deleted from a shelf.

Using Helix Core command-line (P4)

To delete a shelved changelist without removing the files from the associated review:

1. Use a command-line shell and type:

```
$ p4 shelve -d -c changelist
```

Important

Do not delete the files from the shelf individually. Deleting the entire shelf ensures that files are not removed from an associated review even if Swarm is configured to do so.

2. The shelved files are deleted from the pending changelist.

Using P4V

To delete a shelved changelist without removing the files from the associated review:


1. In P4V, right-click on the pending changelist.
2. Select **Delete Shelved Files**.

Important

Do not delete the files from the shelf individually. Deleting the entire shelf ensures that files are not removed from an associated review even if Swarm is configured to do so.

3. The shelved files are deleted from the pending changelist.

Edit reviewers

A review author can edit the reviewers for a review by using the Reviewers **Edit**  button in the review page **Information panel**. Reviewers are able to join or leave reviews, or to change whether their vote is required or optional.

Note


If the review includes content that is part of a project or branch with [Retain default reviewers](#) enabled, the following restrictions apply to the review:

- The voting option for a retained default reviewer can only be changed to a stricter level, you cannot reduce the voting level.
- Retained default reviewers cannot be removed from the review.

Additionally, the following individuals can edit reviewers:

- Users with *admin* or *super* privileges.
- If the review is [moderated](#), the moderators.
- If the review is part of a project, but not moderated, all project members.
- If the review is not part of a project, any authenticated user.

To edit reviewers on a review:


1. Go to the review.
2. From the review **Information panel**, click the **Reviewers** list **Edit**  button.
The **Manage reviewers** dialog is displayed.
3. Click the **Users** tab to edit a user or the **Groups** tab edit a group. This field auto-suggests users, and groups within Helix Core Server as you type (up to a combined limit of 25 entries).
4. **Change the voting requirements of a user or group:**
 - a. Click the dropdown arrow to the right of the user or group you want to change.
 - b. Select one of the following:

- **Users:** select whether their vote is **Required** or **Optional**. A solid star means that their vote is required to approve a review, whereas the outlined star means that their vote is optional.
 - **Groups:** select whether the group vote is **Require all**, **Require one**, or **Optional**. A solid star means that all group member votes are required to approve a review, a solid star with a 1 inside means at least one group member must vote up and no group members vote down to approve a review, and the outlined star means that the group vote is optional.
5. **Add a user or group to a review:**
 - a. Click the **Add user** or **Add group** dropdown icon to display a list of users or groups:
 - b. Start typing the user or group name, the list is filtered as you type.
 - c. Click on the user or group to add to the review.
 - d. Click **+** to add them as a reviewer.
 - e. Set the reviewer type from the dropdown to the right of the user or group.
 6. **Remove a user or group from a review:**
 - a. Click the dropdown icon to the right of the user or group to be removed.
 - b. Select **Remove from the review** from the dropdown menu.
 7. To close the **Manage reviewers** dialog, click **X**.

Add a changelist to a review

Once a review has been started you can add a changelist to the review. It can be useful to add changelists to an existing review. For example, if follow up changes are made to files in a review or if you need to group a number of changelists under a single review.

By default, for a pre-commit or a post-commit review, when a changelist is updated, all of the files in the review are replaced with the files in the changelist you are adding to the review. See ["Replace review with a pending changelist" on page 474](#).

The **Add change** option in the **Review actions**  button is used to add a changelist to an existing review. The option is not available if the review is in a state that is protected from change by the review workflow rules. For details of the **On update of a review in an end state** rule, see ["Workflow rules" on page 526](#).

Tip

You can also add a changelist to a review directly from the changelist description. See ["Review creation and modification outside of Swarm" on page 483](#).

The options available for the **Add change** option depend on whether the review is pre-commit or post-commit:

■ Pre-commit reviews:

- **Append pending changelist:** when you add a pending changelist to a review, the files in the changelist are appended to the existing files in the review. See ["Append a pending changelist to a review" on the next page](#).

Note

A review can only contain 1 stream spec. If you append a changelist with a stream spec to a review that already contains a stream spec, the spec in the changelist replaces the original one in the review. If it is a different spec from the original spec in the review, Swarm cannot display the diff between them and displays **File content unchanged**.

- **Replace with pending changelist:** when you add a pending changelist to a review, all of the files in the review are replaced with the files in the changelist you are adding to the review. See ["Replace review with a pending changelist" on page 474](#).
- **Replace with committed changelist:** when you add a committed changelist to a review, all of the files in the review are replaced with the files in the changelist you are adding to the review. See ["Replace review with a committed changelist" on page 475](#).

Note

If you replace a pre-commit review with a committed changelist, the new version of the review will be a post-commit review.

■ Post-commit reviews:

- **Replace with pending changelist:** when you add a pending changelist to a review, all of the files in the review are replaced with the files in the changelist you are adding to the review. See ["Replace review with a pending changelist" on page 474](#).

Note

If you replace a post-commit review with a pending changelist, the new version of the review will be a pre-commit review.

- **Replace with committed changelist:** when you add a committed changelist to a review, all of the files in the review are replaced with the files in the changelist you are adding to the review. See ["Replace review with a committed changelist" on page 475](#).

Tip

When the content of a review is changed, Swarm checks to see which branches are in the new version of the review:

- **If a new branch was added to the review:**
 - Default reviewers on the new branch are added to the review.
 - Moderators from the added branch become moderators for the review alongside the existing moderators.
 - **Only if workflow is enabled:** if the new branch is associated with a workflow, the [workflow is merged](#) with the existing workflow. The most restrictive workflow is used for the review.
- **If a branch is no longer part of the review:**
 - Reviewers for the review are not changed.
 - Moderators from the removed branch no longer moderate the review.
 - **Only if workflow is enabled:** if the branch was associated with a workflow, the branch workflow is removed from the review.

Append a pending changelist to a review

The **Append pending changelist** option is used to add a pending changelist to a pre-commit review.

Example:

- Original review contains the following files: A^{#1}, B^{#2}, C^{#1}, D^{#1}, and E^{#4}
- Changelist contains the following files: A^{#2}, C^{#3}, E^{#4} (marked for delete), and F^{#1}
- Appending the changelist to the original review generates a new version of the review that contains the following files: A^{#2}, B^{#2}, C^{#3}, D^{#1}, E^{#4} (marked for delete), and F^{#1}

Important

Committing a review with Swarm (recommended): Swarm automatically commits the files in the approved version of the review.

Committing a review outside of Swarm:

Before you commit the review:

1. Unshelve the review into the pending changelist associated with the review.
2. Reshelve the files in the pending changelist.
3. Commit the pending changelist.
4. This process ensures that all of the files in the approved version of the review are committed.


Workflow feature [enabled](#) (default):

Swarm can be configured to automatically check that the files being committed match the files in the approved version of the review by using the **On commit with a review** workflow rule. For information about the **On commit with a review** rule, see ["Workflow rules" on page 526](#).

Workflow feature disabled:

Swarm can be configured to automatically check that the files being committed match the files in the approved version of the review by using the [strict](#) trigger option.

To append a changelist to a review:

1. Click the **Review actions**  button.
2. Select **Append pending changelist** from the dropdown menu and the append review dialog is displayed with a list of recent pending changelists.
3. Enter the pending changelist number directly in the **@ Change** box, filter by depot path and user, or select a changelist from the list of recent changelists.

Note

The changelist must not be part of another review, if it is Swarm will reject it.

4. Click **Select** to append the changelist to the review.
5. The changelist is appended to the review and the review version is updated.

Replace review with a pending changelist

The **Replace with a pending changelist** option is used to replace the files in a review with the files in a pending changelist.

Note

If you replace a post-commit review with a pending changelist, the new version of the review will be a pre-commit review.


Example:

- Original review contains the following files: A^{#1}, B^{#2}, C^{#1}, D^{#1}, and E^{#4}
- Changelist contains the following files: A^{#2}, C^{#3}, E^{#4} (marked for delete), and F^{#1}
- Replacing the original review with the changelist generates a new version of the review that contains the following files: A^{#2}, C^{#3}, E^{#4} (marked for delete), and F^{#1}

Tip

When you replace a review with a changelist, the base revisions of the files in the new version of the review are the base revisions of the files in the replacement changelist.

To replace the files in a review with the files in a changelist:

1. Click the **Review actions**  button.
2. Select **Replace with a pending changelist** from the dropdown menu and the replace review dialog is displayed with a list of recent pending changelists:
3. Enter the pending changelist number directly in the **@ Change** box, filter by depot path and user, or select a changelist from the list of recent changelists.

Note

The changelist must not be part of another review, if it is Swarm will reject it.

4. Click **Select** to replace the review with the changelist.
5. The changelist files replace the files in the review and the review version is updated.

Replace review with a committed changelist

The **Replace with a committed changelist** option is used to replace the files in a review with the files in a committed changelist.

Note

If you replace a pre-commit review with a committed changelist, the new version of the review will be a post-commit review.


Example:

- Original review contains the following files: A^{#1}, B^{#2}, C^{#1}, D^{#1}, and E^{#4}
- Changelist contains the following files: A^{#2}, C^{#3}, E^{#4} (marked for delete), and F^{#1}
- Replacing the original review with the changelist generates a new version of the review that contains the following files: A^{#2}, C^{#3}, E^{#4} (marked for delete), and F^{#1}

Tip

When you replace a review with a changelist, the base revisions of the files in the new version of the review are the base revisions of the files in the replacement changelist.

To replace the files in a review with the files in a changelist:

1. Click the **Review actions**  button.
2. Select **Replace with a committed changelist** from the dropdown menu and the replace review dialog is displayed with a list of recent committed changelists:
3. Enter the committed changelist number directly in the **@ Change** box, filter by depot path and user, or select a changelist from the list of recent changelists.

Note


The changelist must not be part of another review, if it is Swarm will reject it.

4. Click **Select** to replace the review with the changelist.
5. The changelist files replace the files in the review and the review version is updated.

Responsibility

Initially, code reviews have no reviewers.

User and group default reviewers can be set for individual projects and project branches. Each time a new review is created in the project or project branch, the default reviewers will be added to the review. See [default reviewers](#) for details.

Users can show their interest in participating in a code review by clicking the **Review actions**  button and selecting **Join review**, by voting on a review, or by commenting on a review or one of its files. Once a user shows such interest, they are added to the review's list of reviewers and share in the responsibility of performing the code review. Later on, reviewers can change whether their vote is required or optional, or leave a review (perhaps to prevent further notifications).

Looking at a [review list](#) can help you determine which reviews have likely not been started, using the **No reviewers** filter. Once a review has reviewers, it is considered to be active and appears in the review list with the state **Has Reviewers**.


Review participation is *advisory* by default, and is used to inform your team that a code review is being conducted. The disposition of the review is reflected in the review's current state, the badges that may appear over each reviewer's avatar, and any comments reviewers might add.

Authors

The code review author is the user that created the review from a changelist.

As a code review author, you can:

- Add users as reviewers by including an [@mention](#) for each optional reviewer and an [@*mention](#) for each required reviewer in the changelist description. See "[Links in descriptions and comments](#)" on page 409.

- Add groups as reviewers by including [@@mention](#) for each optional group reviewer, [@@*mention](#) for each [required reviewer](#) group (all members required), and [@@!mention](#) for each [required reviewer](#) group (one vote required) in the changelist description. See ["Links in descriptions and comments" on page 409](#).
- Start a review by using P4V, Swarm, or from the command line of P4D. See ["Start a code review" on page 69](#).
- Edit users and groups as reviewers by using the **Edit Reviewers**  button on the review page to add, remove, and edit existing reviewers. See ["Edit reviewers" on page 470](#).
- Add the work-in-progress tag ([#wip](#) by default) to a pending changelist so that you can update files in your changelist without updating the review or creating a new review until you are ready. See ["Work-in-progress tag" on page 393](#)
- Add a changelist to your review, click the **Add change** button on the review and select the changelist you want to add to your review. See ["Add a changelist to a review" on page 471](#).
- Commit your review when it has been approved, click the **Review state** button on the review page and select **Commit**. See ["Change review state" on page 488](#).
- Delete your shelf when the review has been committed. See ["Deleting shelves" on page 469](#).

Moderators

A project moderator is a user assigned to moderate reviews for a specific branch associated with a project.

When you add a moderator to a project branch, only the moderators can approve or reject reviews restriction is enabled for a project branch. See ["Project settings" on page 499](#) for details on adding moderators to project branches.

Changing the state of any review associated with a moderated branch is restricted as follows:

- Only moderators can approve or reject the review. Moderators can also transition a review to any other state.

Important

Moderators prevent the automatic approval of reviews, for more information about automatically approving reviews using workflow rules see ["Workflow rules" on page 526](#).

Note

By default, when a review spans multiple branches that have different moderators, only one moderator from any one of the branches needs to approve the review.

Swarm can be configured to require that one moderator from each branch must approve the review, this is a global setting. If a moderator belongs to more than one of the branches spanned by the review, their approval will count for each of the branches they belong to. For instructions on how to configure moderator behavior, see ["Moderator behavior when a review spans multiple branches" on page 692](#).

- The review's author, when not a moderator, can change the review's state to **Needs review**, **Needs revision**, **Archived**, and can attach committed changelists.

Normally, the review's author cannot change the review's state to **Approved** or **Rejected** on moderated branches. However, authors that are also moderators have moderator privileges, and may approve or reject their own review.

When `disable_self_approve` is enabled, authors who are moderators (or even users with admin privileges) cannot approve their own reviews.

- Project members can change the review's state to **Needs review** or **Needs revision**, and can attach committed changelists. Project members cannot change the review's state to **Approved**, **Rejected**, or **Archived**.
- Users that are not project members, moderators, or the review's author cannot transition the review's state.
- For the review's author and project members, if a review is not in one of their permitted states, for example if the review's state is **Rejected**, they cannot transition the review to another state.

These restrictions have no effect on who can start a review.

Required reviewers

Reviews can optionally have required reviewers. When a review has required reviewers, the review cannot be approved until all required reviewers and required reviewer groups have up-voted the review. If the review is associated with a project that has assigned moderators, even the moderators cannot approve the review without up-votes from all required reviewers (but they can reject the review).

When a group is a required reviewer, it can be set to operate in one of two ways:

- **Require all:** all members of the group must up-vote the review to allow the review to be approved.
- **Require one:** at least one member of the group must up-vote the review to allow the review to be approved. If any member of the group down-votes the review, the review cannot be approved.

Required reviewers are expected to take greater care while performing a review than non-required reviewers, as their votes affect whether a review can be approved or not.


To edit the reviewers for a review and to change whether a reviewer is required or not, see ["Edit reviewers" on page 470](#).

Note

If a review involves a branch with assigned moderators, only a moderator can approve the review, even if all required reviewers have up-voted the review.

See ["Project settings" on page 499](#) for details on adding moderators to project branches.

Add yourself as a reviewer


1. Open the review you want to join.
2. Click the **Review actions**  button.
3. Select **Join review** from the dropdown menu. Alternatively, you can **Vote Up** or **Vote Down** the review.

Your name is added to the **Reviewers** list in the **Information panel**, and you are now a reviewer.

Remove yourself as a reviewer

Tip

You cannot leave the review if you are a retained default reviewer on the review or you are a member of a group that is a reviewer on the review

1. Log in, if you have not already done so.
2. Open the review you want to leave.
3. Click the **Review actions**  button.
4. Select **Leave review** from the dropdown menu.

Your name is removed from the **Reviewers** list in the **Information panel**, and you are no longer a reviewer.

Review workflow

There are many possible code review workflows. The section describes typical scenarios for a code review that Swarm can handle.

Basic review workflow

This section describes a basic workflow when reviewing code with Swarm.

Another developer reviews your code

1. You request a code review with a shelved change.
2. Another developer, Bill, sees the email notification from Swarm, clicks the review link in the email, and begins looking at the diffs in the files belonging to the review. Curious about an implementation detail, Bill clicks the line he's curious about and adds his query in a Swarm comment.
3. You receive an email notification regarding Bill's query. His query prompts you to clarify the code, say by renaming some variables and adding some better descriptive text in the surrounding code comments. You then update the review with your changes.
4. Bill sees the email notification that you have updated the review. He checks out the change, likes what he sees, and marks the review **Approved**.
5. You see the email notification that Bill has approved your review, so you commit your code.

You review the code from another developer

1. Another developer, Charlie, requests a code review with a shelved change.
2. You receive an email notification from Swarm, click the review link in the email, and begin looking at the diffs in the files in the review. You don't like what you see, as Charlie has tried to fix a bug using a technique you have already tried previously and know to be incorrect. You add comments to the code that needs attention, flag your comments as **tasks**, and mark the review **Needs revision**.
3. Charlie receives an email notification regarding your review, but disagrees with you, and adds his own comments justifying his implementation.
4. You receive an email notification regarding Charlie's comment. The technique is somewhat complicated, so rather than attempt to describe how it is incorrect, you unshelve the review's code to your own workspace, change Charlie's code, and shelve your changes. Swarm updates the review with your new code.
5. Charlie receives an email notification regarding your updates to the review. He's still unconvinced, but he unshelves your changes to try them in his local workspace. He finds that your implementation works better, but sees a couple of areas where there could be improvements. He reshelves his latest work to update the review.
6. You receive an email notification regarding Charlie's updates, check out his changes, and realize that Charlie's work is now moot because the customer has revised his plans. You add a comment to the review reporting that fact, and **Reject** the review.

Additional workflow tasks

This section describes additional workflow tasks that can be used with the [basic workflow](#) described above.

Remove a file from a review

Important

Swarm must be configured to use this function, see ["Process shelf file delete when"](#) on page 693.

When configured, Swarm will automatically remove files from a review when they have been deleted from an associated shelved changelist and the review is in a specified state. Typically Swarm is configured to remove files from a review when the review state is **Needs review** and **Needs revision** but not when it is any other state.

To remove a file from a review using P4V:

1. Right-click the shelved file you want to remove from the review.
2. Select **Delete**, and click **Yes** when prompted to confirm the deletion. The file is removed from the associated review if the review is in one of the states specified in the Swarm configuration.

Tip

If you want to tidy up your workspace by deleting an unused shelf without removing the files from an associated review, see ["Deleting shelves"](#) on page 469.

Revert the content of a review to match an earlier review version

If the content of the latest version of a review contains changes you want to rollback, or is created by mistake, you can revert the review content to match an earlier version of the review. This section describes the process for reverting the content of a review.

Tip

- Reverting the content of a review increments the review version, it does not change the review version to the version you are reverting to.
- It is good practice to update the review description, or add a comment, so that users know that you have reverted the review with files from an earlier version.

To revert the content of a review:

1. Find the changelist number for the review version you want to revert to.

Tip

The changelist numbers in a review can be found using the review version selector on the review page, see ["Select review versions to view" on page 454](#).

2. Unshelve the changes from the changelist you found in the previous step.
3. Create a new changelist and shelve the changes into the new changelist.
4. Use Swarm to navigate to the review you want to revert.
5. Select the **Add change** option, and select **Replace with pending changelist** from the dropdown list. The **Select a pending changelist to replace the review** dialog is displayed.
6. Select the new changelist you shelved the files in and click **Select**. The new changelist files replace the files in the review and the review version is updated.

Fix a review if it has been replaced with the wrong changelist

If you or another user has replaced the files in your review with the wrong changelist, you can fix the mistake so that the review contains the correct files.

For example:

1. Create a pending changelist with files A^{#1}, and B^{#1}, this is changelist 250.
2. Request a review of changelist 250 using Swarm by clicking the **Request Review** button on the changelist page.
 - Version 1 of review 251 is created containing files A^{#1}, and B^{#1}.
3. Go to review 251 in Swarm, select the **Add change** option, select **Replace with committed changelist**, and select changelist 241. Changelist 241 contains files X^{#1}, and Y^{#1}.
 - This replaces the files in the review with the files in changelist 241.
 - Version 2 of review 251 is created, containing only files X^{#1}, and Y^{#1}.
4. You realize that you have made a mistake, you meant to append changelist 242 to review 251.

Fix the mistake:

Review-251 should contain files A^{#1}, and B^{#1} from changelist 250 and files C^{#1}, and D^{#1} from changelist 242.

Tip

It is good practice to update the review description, or add a comment, so that users know why you have changed the files in the review.

1. Go to review 251 in Swarm, select the **Add change** option, select **Replace with pending changelist**, and select changelist 242. Changelist 242 contains files C^{#1}, and D^{#1}.
 - This replaces the files in the review with the files in changelist 242.
 - Version 3 of review 251 is created, containing files C^{#1}, and D^{#1}.
2. Edit the description of changelist 250 to change `#review-251` to `#append-251`, this changes the add mode for review-251 to append.
3. Shelve your files in pending changelist 250. Changelist 250 contains files A^{#1}, and B^{#1}. This appends the files to the review because that is the default add mode for review-251.
 - Version 4 of review 251 is created, containing files A^{#1}, B^{#1}, C^{#1}, and D^{#1}.

Review creation and modification outside of Swarm

You can create a review and add a changelist to a review from outside of Swarm by adding keywords to the changelist description. The keyword is processed by Swarm when the changelist is shelved for review, when files in a pending changelist that is associated with a review are reshelved, and when a changelist is committed. For details on using a keyword to create a review and add a changelist to a review, see ["Create a review" on page 682](#) and ["Add a changelist to a review" on page 682](#).

Note

Swarm acts on the first valid keyword it finds in the changelist description, Swarm then ignores any further valid keywords it finds in the description.

Tip

When the content of a review is changed, Swarm checks to see which branches are in the new version of the review:

- **If a new branch was added to the review:**
 - Default reviewers on the new branch are added to the review.
 - Moderators from the added branch become moderators for the review alongside the existing moderators.
 - **Only if workflow is enabled:** if the new branch is associated with a workflow, the [workflow is merged](#) with the existing workflow. The most restrictive workflow is used for the review.
- **If a branch is no longer part of the review:**
 - Reviewers for the review are not changed.
 - Moderators from the removed branch no longer moderate the review.
 - **Only if workflow is enabled:** if the branch was associated with a workflow, the branch workflow is removed from the review.

The following examples show how adding keywords to the changelist descriptions works in practice. In each case, you are collaborating with other users to produce a single review in Swarm:

Basic workflow using #review and #append in the changelist description

1. You create a review with files A^{#1}, B^{#1}, and C^{#1} by using **#review** in the changelist description.
2. Review-xxxx is created and **#review** is automatically replaced with **#review-xxxx** in the changelist description.
3. User-2 appends their changelist that contains files C^{#2}, D^{#1}, and E^{#1} to review-xxxx by using **#append-xxxx** in the changelist description.
 - Review-xxxx now contains files A^{#1}, B^{#1}, C^{#2}, D^{#1}, and E^{#1}.
 - The default add mode for review-xxxx is now append because User-2 added the changelist to the review using **#append-xxxx**.
4. You edit files A^{#1} and B^{#1} in your changelist so they become A^{#2} and B^{#2}. The changelist description still contains **#review-xxxx**.
5. You shelve your pending changelist which has **#review-xxxx** in the changelist description, this appends the files to the review because that is the default add mode for review-xxxx.
 - Review-xxxx now contains files A^{#2}, B^{#2}, C^{#2}, D^{#1}, and E^{#1}.

Advanced workflow using #review and #append in the changelist description

1. You create a review with files A^{#1} and B^{#1} by using **#review** in the changelist description.
2. Review-xxxx is created and **#review** is automatically replaced with **#review-xxxx** in the changelist description.
3. User-2 appends their changelist that contains file C^{#1} to review-xxxx by using **#append-xxxx** in the changelist description.
 - Review-xxxx now contains files A^{#1}, B^{#1}, C^{#1}.
 - The default add mode for review-xxxx is now append because User-2 added the changelist to the review using **#append-xxxx**.
4. User-3 adds their changelist that contains files A^{#2} and D^{#1} to review-xxxx by using **#review-xxxx** in the changelist description. This appends the files to the review because that is the default add mode for review-xxxx.
 - Review-xxxx now contains files A^{#2}, B^{#1}, C^{#1}, and D^{#1}.
5. You edit files A^{#2} and B^{#1} in your changelist so they become A^{#3} and B^{#2}, and you add file E^{#1}. The changelist description still contains **#review-xxxx**.

6. You shelve your pending changelist which has `#review-xxxx` in the changelist description, this appends the files to the review because that is the default add mode for review-xxxx.
 - Review-xxxx now contains files A^{#3}, B^{#2}, C^{#1}, D^{#1}, and E^{#1}.

Advanced workflow using `#review`, `#append`, and `#replace` in the changelist description

1. You create a review with files A^{#1} and B^{#1} by using `#review` in the changelist description.
2. Review-xxxx is created and `#review` is automatically replaced with `#review-xxxx` in the changelist description.
3. User-2 appends their changelist that contains file C^{#1} to review-xxxx by using `#append-xxxx` in the changelist description.
 - Review-xxxx now contains files A^{#1}, B^{#1}, C^{#1}.
 - The default add mode for review-xxxx is now append because User-2 added the changelist to the review using `#append-xxxx`.
4. User-3 replaces review-xxxx with their changelist that contains files A^{#2} and D^{#1} to review-xxxx by using `#replace-xxxx` in the changelist description.
 - Review-xxxx now contains files A^{#2} and D^{#1}.
 - The default add mode for review-xxxx is now replace because User-3 added the changelist to the review using `#replace-xxxx`.
5. You edit files A^{#2} and B^{#1} in your changelist so they become A^{#3} and B^{#2}, and you add file E^{#1}. The changelist description still contains `#review-xxxx`.
6. You shelve your pending changelist which has `#review-xxxx` in the changelist description, this replaces the files in the review because that is the default add mode for review-xxxx.
 - Review-xxxx now contains files A^{#3}, B^{#2}, and E^{#1}.

States

Reviews can be in one of several states. The biggest differentiator is whether the review's files have any outstanding, uncommitted changes or not.

Whenever a review state changes, an email notification is sent to all review participants, including:

- The review author
- Any user who comments on the review or its files
- Any user who has changed the review's state previously
- Any user who is [@mentioned](#), or a member of a group that is [@@mentioned](#) in the review's description or comments.

The review state is indicated by the **Review state** icon on the [Review page](#). The **Change state** button is used to change the state of a review, see ["Change review state" on page 488](#).

Code reviews can be in one of the following states:

- **Needs review:** The review has started and the changes need to be reviewed.
- **Needs revisions:** The changes have been reviewed and the reviewer has indicated that further revisions are required.
- **Approved:** The review has completed. The changes may need to be committed. If the changes have been committed then this review will be Approved and closed, otherwise it will be Approved and open. See the note [below](#).
- **Rejected:** The review has completed. The changes are undesirable and should not be committed.
- **Archived:** The review has completed for now. However, it is neither rejected nor approved; it is simply put aside in case it is needed in the future.

Note

By default, when an **Approved** review is committed or updated, Swarm changes the state to **Needs Review** if the files have been modified since the review was approved. Files are considered modified if the list of involved files changes, or if the file content or file-type changes.

If one or more files in a review has the filetype +k (**ktext**), this behavior is undesirable because the files will appear to be modified when the Helix Core Server replaces RCS keywords with their current values. See ["Unapprove modified reviews" on page 742](#) to see how to disable this behavior.

Self-approval by review authors

By default, review authors can approve their own reviews. This behavior is based on Swarm's [advisory nature](#).

Self-approval by authors can be prohibited on a project-by-project basis by specifying moderators for project branches (see ["State change restrictions with moderation" below](#)). However, authors who are moderators can self-approve their own reviews.

Administrators can configure Swarm to prevent all self-approval by review authors. See ["Disable self-approval of reviews by authors" on page 689](#).

State change restrictions with moderation

Typically, any authenticated user can change the state of a review (remember that the review state is merely advisory in most cases).

When you add a moderator to a project branch, only the moderators can approve or reject reviews restriction is enabled for a project branch. See ["Project settings" on page 499](#) for details on adding moderators to project branches.

Changing the state of any review associated with a moderated branch is restricted as follows:

- Only moderators can approve or reject the review. Moderators can also transition a review to any other state.

Important

Moderators prevent the automatic approval of reviews, for more information about automatically approving reviews using workflow rules see ["Workflow rules" on page 526](#).

Note

By default, when a review spans multiple branches that have different moderators, only one moderator from any one of the branches needs to approve the review.

Swarm can be configured to require that one moderator from each branch must approve the review, this is a global setting. If a moderator belongs to more than one of the branches spanned by the review, their approval will count for each of the branches they belong to. For instructions on how to configure moderator behavior, see ["Moderator behavior when a review spans multiple branches" on page 692](#).

- The review's author, when not a moderator, can change the review's state to **Needs review**, **Needs revision**, **Archived**, and can attach committed changelists.

Normally, the review's author cannot change the review's state to **Approved** or **Rejected** on moderated branches. However, authors that are also moderators have moderator privileges, and may approve or reject their own review.

When `disable_self_approve` is enabled, authors who are moderators (or even users with admin privileges) cannot approve their own reviews.

- Project members can change the review's state to **Needs review** or **Needs revision**, and can attach committed changelists. Project members cannot change the review's state to **Approved**, **Rejected**, or **Archived**.
- Users that are not project members, moderators, or the review's author cannot transition the review's state.
- For the review's author and project members, if a review is not in one of their permitted states, for example if the review's state is **Rejected**, they cannot transition the review to another state.

These restrictions have no effect on who can start a review.

State change restrictions without moderation

Changing the state of any review associated with a project that has no moderation is restricted as follows:

- Only members and the review author can change the state of a review or attach commits.
- If a user or group is added in ["Members or groups who can ignore ALL workflow rules" on page 758](#) list in the **Global Workflow** setting, they can transition the review state.

Required reviewers

Reviews can optionally have required reviewers. When a review has required reviewers, the review cannot be approved until all required reviewers and required reviewer groups have up-voted the review. If the review is associated with a project that has assigned moderators, even the moderators cannot approve the review without up-votes from all required reviewers (but they can reject the review).

When a group is a required reviewer, it can be set to operate in one of two ways:

- **Require all:** all members of the group must up-vote the review to allow the review to be approved.
- **Require one:** at least one member of the group must up-vote the review to allow the review to be approved. If any member of the group down-votes the review, the review cannot be approved.

Required reviewers are expected to take greater care while performing a review than non-required reviewers, as their votes affect whether a review can be approved or not.

To edit the reviewers for a review and to change whether a reviewer is required or not, see ["Edit reviewers" on page 470](#).

Note

If a review involves a branch with assigned moderators, only a moderator can approve the review, even if all required reviewers have up-voted the review.

See ["Project settings" on page 499](#) for details on adding moderators to project branches.

Change review state

The **Change state** button on the ["Review display" on page 433](#) page is used to manually change the state of a Swarm review.

To change the state of the review:

1. Navigate to the review you want to update.
2. Click the **Change State** button.
3. Select the new state from the dropdown menu.
The options available depend on the current state of the review and your user permissions:
 - **Needs revision:** Select to request changes to the files in the review.
 - **Needs review:** Select to request further review of the changes.
 - **Approve** (only available if the voting requirements for the review are satisfied. For information on voting requirements, see ["Required reviewers" on page 478](#)): select to approve the review.

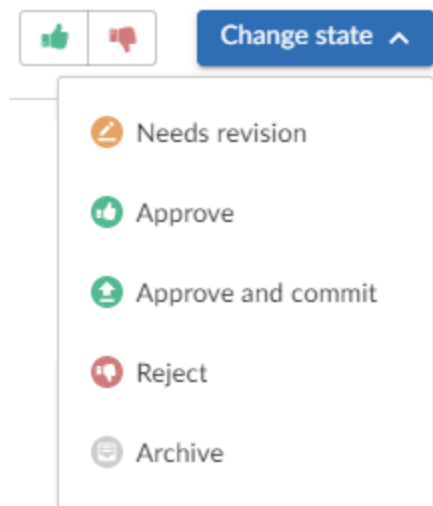
- **Commit** (only available for [pre-commit](#) reviews that have been approved): Select to commit the review.
- **Approve and commit** (only available for unapproved [pre-commit](#) reviews when the voting requirements for the review are satisfied. For information on voting requirements, see ["Required reviewers" on page 478.](#)): Select to approve and commit the review in a single step. See ["Approve and commit" below.](#)
- **Reject**: Select to reject the review.
- **Archive**: Select to archive the review.

Approve and commit

Approve and commit is available in the dropdown menu for unapproved pre-commit reviews, this enables you to approve and commit a review in a single step if required.

Note

- By default, any authorized user can commit a review. Swarm can be configured to restrict this behavior so that only the review author can commit the review, see ["Disable commit" on page 702](#) for details.
- **Approve and commit** is only available for unapproved [pre-commit](#) reviews when the voting requirements for the review are satisfied. For information on voting requirements, see ["Required reviewers" on page 478.](#)



Note

If Swarm is configured to prevent approval of reviews with open tasks and a review has open tasks, the **Approve**, and **Approve and commit** options will not be available for the review. This option is configured by an administrator. See ["Disable approve for reviews with open tasks" on page 688.](#)

To approve, or approve and commit a review with open tasks, you must address the tasks first and then set them to **Task Addressed**, or **Not a Task**. See ["Set a task to Task addressed or Not a task" on page 352](#) for details.

1. Select **Approve and commit** from the dropdown menu.
2. The **Commit Review** dialog is displayed.

Commit Review

Use the new P4 Core package names for specifying dependencies.

Also update the copyright and release information on the meta files.

☒ **job083648 fixed** Need to rename the packages to be helix-insights rather than performe-...

Job Status on Commit: Fixed ▾ ☐ Remove pending changelists **Approve and Commit** Cancel

3. Edit the review description if required.
4. Select which jobs should be associated with the review, and specify the job status on commit.
5. Select **Remove pending changelists**, Swarm will attempt to automatically clean up any changelists left behind after the review has been committed, including removing any shelved files. This option can be removed by an administrator, see ["Review cleanup" on page 678](#) for details.
6. Click **Approve and commit** to approve the review and commit the associated files.

Note

By default, Swarm records that you committed the review on behalf of the review's author. This can be configured by an administrator to only credit the committer and not the review author, see ["Commit credit" on page 586](#) for details.

6 | Groups

Groups are a feature of Helix Core Server that makes it easier to manage permissions for users. Swarm can use groups to coordinate review activities and responsibilities. This chapter covers how to manage groups in Swarm, including how to:

- "Add a group" below
- "Edit a group" on page 497
- "Delete a group" on page 497

See "Groups" on page 373 for an introduction to Swarm groups.

Add a group


Important


You must have *super* privileges in Helix Core Server (**p4d**), or have *admin* privileges in **p4d** version 2012.1 or later, to create a group. If you do not have sufficient permissions, Swarm does not display the **Add Group** button.

1. To view a list of groups, click **Groups** in the menu.
2. Click the **+ Add Group** button.

The Add Group **Settings** tab is displayed:

Add Group

 Settings


 Notifications

Name

Description


Description

Owners



 Add an Owner

Members



 Add a User or Subgroup

Save

Cancel

3. Provide a name for the group.
4. **Optional:** provide a description.
5. **Optional:** specify an owner. This field auto-suggests users within Helix Core Server as you type.

Group owners are not members of their groups. A group owner can be added as a member of a group.

Once specified, modifying the group's definition is restricted to group owners and users with *super* privileges in Helix Core Server.

If you do not specify an owner, you must specify at least one member (below).

6. **Optional:** specify group members. This field auto-suggests projects, groups, and users within Helix Core Server as you type (up to a combined limit of 20 entries).

If you specify a project, the project's members become members of the group. If you specify a group, that group becomes a sub-group of your new group, and all of its members (and members of any of its sub-groups) become members of your new group.

If you do not specify any members, you must specify at least one owner (above).

7. Choose one the following options:
 - Click **Save** to finish adding the group. By default group members will be emailed when a new review is requested.

Note

The **Save** button is disabled if any required fields are empty.

or

- **Optional:** configure email notifications for the group in the Add Group, Notifications tab. See the next step for details.

8. **Optional:** click the **Notifications** tab to configure group email notifications.

Add Group

[Settings](#)
[Notifications](#)

☐ Use mailing list instead of notifying by individual group member's emails (must add email address)

Group mailing list address

Email Notifications ☒ Email members when a new review is requested

☐ Email members when a change is committed

Adjust when notifications are sent to group members (use mailing list address to enable).

Email group members when:

[Reset to default](#)

Check all ☐

The group is a member of a project, and

a review is started in the project ☒

a review or change is committed ☒

The group is a reviewer on a review, and

files in the review are updated ☒

tests on the review have finished ☒

a vote is cast on a review ☒

the state of the review changes ☒

someone joins or leaves the review ☒

a review or change is committed ☒

a comment is made on the review ☒

a comment on the review is updated ☒

The group is a moderator on a project, and

files in the review are updated ☒

tests on the review have finished ☒

a review or change is committed ☒

Save

Cancel

9. **Optional:** add a group mailing list address by selecting **Use mailing list instead of notifying by individual group member's emails (must add email address)** and entering a valid group email address.

Note

- **Group mailing list enabled:** notifications are sent to the group email address.

When a group mailing list is enabled it overrides the group member's preference set for the `notify_self` configurable in the `SWARM_ROOT/data/config.php` file. So even when the `notify_self` configurable is set to false the group member will receive an email notification.

There is a caveat that if a user is a group member and the same user is added individually to the review and has set the `notify_self` configurable to true, the user receives two email notifications.

- **Group mailing list disabled:** notifications are sent to the group members individual email addresses.

The format of the email address is validated as you type.



10. Group members can be notified when a member of the group starts a review. Group members can be notified when a change is committed by, or on behalf of, a changelist owner who is also a member of this group. These settings are always available even if the group mailing list is not enabled.

Select which actions send a notification to the group:

- **Email members when a review is requested:** When any member of this group creates a review, this group will be notified.
- **Email members when a change is committed:** When a change is committed into Perforce, if the owner of the changelist is a member of this group, this group will be notified.

Tip

When a user commits a changelist in Swarm, it is committed on behalf of the changelist owner. If the changelist owner is a member of this group, this group will be notified.

Note

Members of your group may receive notification emails even if group notifications are disabled as they may be members of a project, or follow a project or user, or Swarm's review daemon functionality may be enabled. See [Notifications](#) for details.

11. Group notification settings allow you to configure which notifications are sent to the group mailing list when events occur within Swarm (the group mailing list must be enabled). This allows you to limit the number of emails sent to the group mailing list. These settings apply across all projects.

Note

The following group notification settings are only available if the group mailing list address is configured.

Note

Defaults for these options are configured by the Swarm administrator, and they may force some of these options to on or off. See ["Global settings" on page 664](#) for how this is configured.

Adjust when notifications are sent to group members.

Email group members when:

Reset to default

Check all ☐

The group is a member of a project, and	
a review is started in the project	<input checked="" type="checkbox"/>
a review or change is committed	<input checked="" type="checkbox"/>
The group is a reviewer on a review, and	
files in the review are updated	<input checked="" type="checkbox"/>
tests on the review have finished	<input checked="" type="checkbox"/>
a vote is cast on a review	<input checked="" type="checkbox"/>
the state of the review changes	<input checked="" type="checkbox"/>
someone joins or leaves the review	<input checked="" type="checkbox"/>
a review or change is committed	<input checked="" type="checkbox"/>
a comment is made on the review	<input checked="" type="checkbox"/>
a comment on the review is updated	<input checked="" type="checkbox"/>
The group is a moderator on a project, and	
files in the review are updated	<input checked="" type="checkbox"/>
tests on the review have finished	<input checked="" type="checkbox"/>
a review or change is committed	<input checked="" type="checkbox"/>

SaveCancel

Toggle notifications for each event on or off to control whether the group receives an email when that event occurs.

Clicking **Reset to default** resets the options back to system defaults.

- Click **Save**.

Note

The **Save** button is disabled if any required fields are empty.

Edit a group

Important

You must be an owner of the group, or be a user with *super* privileges in Helix Core Server, to edit the group.

1. Visit the group page you want to edit.
2. Click the **Settings** tab for the group to edit group details. If you do not have permission to edit a group, this tab does not appear.
3. Edit group details as required. See ["Add a group" on page 491](#) for more details.
4. Click the **Notifications** tab for the group to edit group notifications. If you do not have permission to edit a group, this tab does not appear.
5. Edit group notifications as required. See ["Add a group" on page 491](#) for more details.
6. Click **Save**.

Note

When a group is edited, the list of associated reviews is not immediately updated to match any changes in membership. The association with a review does not change until the review is updated.

Delete a group

Important

You must be an owner of the group, or be a user with *super* privileges in Helix Core Server, to delete the group.

1. Visit the group page you want to delete.
2. Click the **Settings** tab for the group. If you do not have permission to edit a group, this tab does not appear.
3. Click **Delete**.
A tooltip is displayed to confirm that you want to delete this group.
4. Click **Delete** to confirm.

7 | Projects

A Swarm project is a group of Helix Core Server users who are working together on a specific codebase, defined by one or more branches of code, along with a job filter, and automated test integration. This chapter covers Swarm's project management capabilities, including:

- "Add a project" below
- "Project settings" on the next page
- "Edit a project" on page 514
- "Membership" on page 514
- "Delete a project" on page 521

See "Projects" on page 381 for an introduction to Swarm projects.

Add a project

1. On the Swarm **Projects** page, click the **+ Add Project** button.

The **Add Project** page dialog is displayed.

Note

- The ability to add projects can be [limited to administrators only](#), or [limited to members of specific groups](#). When limited, users who are not administrators, or a member of the specified group, will not see the **+ Add Project** button.
- By default, any member of a project can edit the project's configuration. Administrators can configure Swarm to prevent changes to the project's name and branch definition(s).

2. Provide a name for the project and click the **Create** button.

A project summary with Participants tab, Branches tab, Integrations tab, and General Settings tab is displayed.

HelixSwarm Projects/acme/settings

ACME
Project ID acme

Follow

28 Members 1 Followers 2 Branches

This is a sample project.

Participants Branches Integrations General Settings

Members(28) Manage members

Owners(6) Manage owners

Default reviewers(10) Retain default reviewers Manage default reviewers

Followers(1)

☒ Email members and moderators when a review is requested

☒ Email members, moderators, and followers when a change is committed

For more information on how to use the **Settings** page, see ["Project settings" below](#).

Project settings

This section describes how to use the new project settings page to view and edit project settings for a project. For more information about how to add a project, see ["Add a project" on the previous page](#).

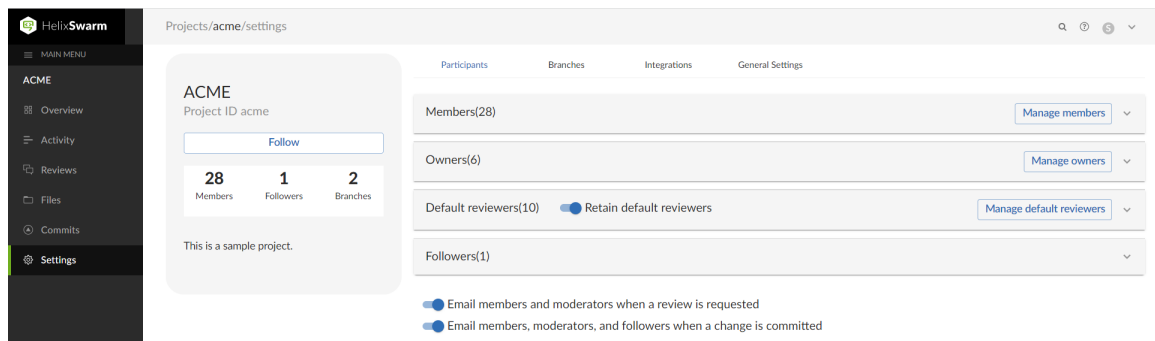
Note

Modifying the project's definition is restricted to project owners and administrators (users with *admin*-level or *super*-level privileges in Helix Core Server). See ["Membership" on page 514](#) for more details.

To view or edit project settings for a specific project:

1. Click **Projects** in the menu.
2. Click on a project name displayed in **My Projects** list.
3. Click **Settings** in the menu.

A project summary with Participants tab, Branches tab, Integrations tab, and General Settings tab is displayed.



The following information is displayed in the project summary:

- The name of the project.
- The unique associated project identifier.
- A **Follow** button that you can click to follow the project. This button is not available if you are a member of the project.
- Total number of members, followers, and branches associated with the project.
- A short description about the project.

Participants tab











The **Participants** tab shows the following details:

Members

Displays a list of all the project members. This includes users, groups, and projects.

[Participants](#) [Branches](#) [Integrations](#) [General Settings](#)

Members(28) [Manage members](#) ^

Name	Type	
 Adena Waters	User	
 Adolfo Whipple	User	
 Adrienne Fitzpatrick	User	
 Albina Bona	User	
 Becky Hoesly	User	

To add a project member:

1. Click **Manage members** button.
2. Select the **Users**, **Groups**, or **Projects** tab that you wish to add to the project members list.
3. Type in to search the name of a user, group, or project and click **Add** button.

To delete a project member, click the delete icon against their name.

Owners

Displays a list of all the project owners.













Note

Only users can be added as project owners.

[Participants](#) [Branches](#) [Integrations](#) [General Settings](#)

Members(28) [Manage members](#) ▾

Owners(6) [Manage owners](#) ▲

Name	
 Aaron Gleber	
 Abigail Germann	
 Chassidy Shorts	
 Domonique Pantoja	
 Suzie Swansbrough	
 Vernon Renno	

Rows per page: 25 ▾ 1–6 of 6 < >

To add a project owner, click **Manage owners** button. Type in the name of the user you wish to add to the project owners list. The field auto-suggests users within Helix Core Server as you type. Click the **Add** button to add a user to the project owners list.

To delete a project owner, click the delete icon against their name.

Default reviewers

















Displays a list of all the default reviewers. A default reviewer can be a user or a group.

Participants Branches Integrations General Settings

Members(28) Manage members ▾

Owners(6) Manage owners ▾

Default reviewers(9) ☐ Retain default reviewers Manage default reviewers ▴

Name	Vote Requirement ↑	Type
 Abbie Korewdit	Required 	User
 Ada Garrison	Optional 	User
 Avril Cottrill	Optional 	User
 Bennett Duriga	Optional 	User
 Caitlin Dom	Required 	User
 Suzie Swansbrough	Optional 	User
 Small	Require one 	Group
 Small1	Optional 	Group

To add a new default reviewer:

1. Click **Manage default reviewers** button.
2. Select the **Users** or **Groups** tab that you wish to add to the default reviewers list.
3. Type in to search the name and click **Add** button.
4. Click the drop-down box against the user or group name and select your desired option.

To edit your default reviewer preferences for an existing reviewer, click the drop-down box against the user name and select your desired option.

For a user, you can set the vote requirement to **Optional** or **Required**. When a group is a default reviewer, it can be set to operate in one of three ways:

■ **Optional:**

No members of the group are required to approve the review.

■ **Require one** indicated by a star badge with a 1:

At least one member of the group must up-vote the review to allow the review to be approved. If any member of the group down-votes the review, the review cannot be approved.

- **Required** indicated by a star badge:

All members of the group must up-vote the review to allow the review to be approved.

To delete a default reviewer, click the delete icon against their name.

Important

When a review is part of multiple projects/project branches:

- The default reviewer lists for all of the projects and project branches the review is part of are combined and added to the review.
- If a default reviewer has different reviewer options set on projects and project branches that the review is part of, the strictest reviewer option is used for the review.

Example: A review is created and it is part of **Project A**, **Project B**, and **Project Branch b**.

Project A: default reviewer X is an Optional reviewer

Project B: default reviewer X is an Optional reviewer

Project Branch b: default reviewer X is a Required reviewer

Result: default reviewer X is added to the review as a Required reviewer

Note

If users or groups are [@mentioned](#) in a new changelist description that includes **#review**, they will be added to the review as reviewers. If any of these reviewers are already specified as default reviewers they will not be added to the review again, the reviewer's most restrictive reviewer option is used for the review.

Note

If a default reviewer is deleted from Helix Core Server they will not be added to new reviews.

To retain default reviewers, toggle the **Retain default reviewers** switch. For more details about retaining default reviewers, see ["Membership" on page 514](#).

Followers

Displays a list of all the followers.

Participants Branches Integrations General Settings

Members(28) [Manage members](#)

Owners(6) [Manage owners](#)

Default reviewers(10) ☒ Retain default reviewers [Manage default reviewers](#)

Followers(1)

Name
Yael Larez

Rows per page: 25 1-1 of 1

☒ Email members, moderators, and followers when a review is requested

☒ Email members, moderators, and followers when a change is committed

By default, project members and moderators are notified when a new review is started. Project members, moderators, and followers are notified when a change is committed.

Toggle the following switches at the bottom of the **Participants** tab to ensure which actions send a notification:

- **Email members, moderators, and followers when a review is requested:** When a new review is requested for the project, all project members and moderators of the project are added to the email notification list.
- **Email members, moderators, and followers when a change is committed:** When a change is committed for the project, all project members, project moderators and project followers of this project are added to the email notification list.

Note

- When a group is a project member or project moderator, all of the members of that group are notified using the same logic as for individual project members and moderators.
- Any "[Links in descriptions and comments](#)" on [page 409](#) users and groups, or users and groups who are explicitly added to a review or changelist, will receive notifications even if new review/committed review notifications are disabled.

Branches tab

The Branches tab enables you to add a branch and configure the following options:

- Associate a workflow with the project branch
- Set minimum up votes
- Manage moderators
- Manage branch specific default reviewers

Add a branch

To add a branch:

1. Click the **Add new branch** button.
The Add new branch dialog displays.
2. Enter a **Name** for your branch.
3. Enter one or more branch paths in the **Depot paths** field using depot syntax, one path per line.

Tip

- Branch paths, and files can be excluded by putting a minus symbol – at the start of the path.
- Branch paths are processed in order, starting with the first file path in the list.

For example:

```
//depot/main/swarm/...  
-//depot/main/swarm/test/...  
//depot/main/swarm/test/ResultSummary.html
```

The first path includes all of the directories and files under `//depot/main/swarm/` in the project branch.

The second path excludes all of the files in `-//depot/main/swarm/test/` from the project branch.

The third path includes the `ResultSummary.html` file from the previously excluded `//depot/main/swarm/test/` directory.

Note

- Wildcards should not be used; the only exception is that the branch path can end with the Helix Core Server wildcard `...`
- Branch paths are case sensitive.
- Branch paths, and files are not checked to see if they are valid when you save the branch:

- If you enter an invalid path ending in the wildcard `...`, the path will not be displayed in the [project file browser](#) until the path is created. This allows you to specify a path before it has been created.
- If you enter a path that ends with a file that has not been committed, or a non-existent file, Swarm displays a 404 error when you navigate the path to the file with the [project file browser](#).

Note

The [Project Commits tab](#) can fail to show some Helix Core Server commits in the top level view. Individual branch views display the commits correctly:

The **Project Commits** tab top level client view is made up of all of the branches of the project.

For example, Project Alpha:

Branch: QA:

`//depot/alpha/dev/QA/...`

Branch: Dev :

`//depot/alpha/dev/...`

`-//depot/alpha/dev/QA/...`

The project commits tab view is generated by processing the branches in the order that they were created in and from top to bottom for the paths in each of those branches.

For the project Alpha example above:

The first path includes all of the directories and files under

`//depot/alpha/dev/QA/` in the project commits tab top level view.

The second path includes all of the directories and files under

`//depot/alpha/dev/` in the project commits tab top level view.

The third path excludes all of the directories and files in

`//depot/alpha/dev/QA/` from the project commits tab top level view.

Result: commits made to `//depot/alpha/dev/QA/` that should be shown for the **QA** branch are not displayed in the **Project Commits** tab top level view.

When you have a simple branch structure this can be avoided by considering this issue when you create your branches. In the example above, creating the **Dev** branch first and then creating the **QA** branch avoids the problem because `//depot/alpha/dev/QA/` is not excluded from the **Project Commits** tab top level view.

4. Click the **Confirm** button.

A new branch is created.

Note

For multiple branches the most recently created branch is placed at the top.

Edit a branch

Note

All settings described in this section are optional settings for a branch.

Expand the branch pane to view and further edit the branch settings as follows:

To associate a workflow, select the workflow from the **Workflow** dropdown list, or enter the workflow name in the search field. The field auto-suggests workflows as you type. Only workflows that you own and shared workflows are shown in the dropdown list.

Tip

- When a workflow is associated with a project, the workflow is used for all of the branches in that project.
- When a project branch is associated with a workflow, the workflow of the parent project is ignored and the branch workflow is used.

For more information about workflows and how project workflows interact with branch

1. workflows, see ["Workflow basics" on page 529](#).

To use the parent project workflow, select **Inherit from project** from the **Workflow** dropdown list. This is the default when you create a new branch. If the parent project is set to **No workflow**, the branch will use the global workflow rules.

To remove the existing workflow without replacing it with another workflow, select **No Workflow** from the **Workflow** dropdown list. This is the default when you create a new project.

2. Set the **Minimum Upvotes** required for reviews associated with this branch.

When **Minimum up votes** is set on a project, the project setting is used on the project branches unless **Minimum up votes** is set to **1** or higher on a branch. In this case the branch setting is used and the project setting is ignored.

A review cannot be approved until all of the [Required reviewers](#) have voted up the review and the **Minimum up votes** specified has been satisfied.

- If a review spans projects/branches, the Minimum up votes for **each** of the projects and branches must be satisfied before you can approve the review.
- Required reviewers are included when up votes are counted.
- When **Count votes up from** is set to **Members** for a workflow associated with a project/branch, only the up votes of members of the project contribute to satisfying the **Minimum up votes** for a project/branch. For more information about the **Count votes up from** rule, see ["Workflow rules" on page 526](#).

Important

If the Workflow feature is disabled, all votes are counted not just votes from project members.

Important

If the **Minimum up votes** required is set higher than the number of reviewers that exist for a review, approval will be blocked for that review. This is true even if all the reviewers on the review have voted up the review.

3. Add moderators for a branch from the **Branch moderators** pane.

Note

This ensures that only moderators can approve or reject reviews.

To add a moderator:

- a. Click **Manage moderators** button.
- b. Select the **Users** or **Groups** tab that you wish to add to the moderators list.
If a group is specified as a moderator, all of the members of that group have the same moderator privileges for that project branch as if they were added individually.
- c. Type in to search the name of the user or group. The field auto-suggests groups and users within Helix Core Server as you type. Click **Add** button.

When you add a moderator to a project branch, only the moderators can approve or reject reviews restriction is enabled for a project branch. Changing the state of any review associated with this moderated branch is restricted as follows:

- Only moderators can approve or reject the review. Moderators can also transition a review to any other state.

Important

Moderators prevent the automatic approval of reviews, for more information about automatically approving reviews using workflow rules see ["Workflow rules" on page 526](#).

Note

By default, when a review spans multiple branches that have different moderators, only one moderator from any one of the branches needs to approve the review.

Swarm can be configured to require that one moderator from each branch must approve the review, this is a global setting. If a moderator belongs to more than one of the branches spanned by the review, their approval will count for each of the branches they belong to. For instructions on how to configure moderator behavior, see ["Moderator behavior when a review spans multiple branches" on page 692](#).

- The review's author, when not a moderator, can change the review's state to **Needs review**, **Needs revision**, **Archived**, and can attach committed changelists.
Normally, the review's author cannot change the review's state to **Approved** or **Rejected** on moderated branches. However, authors that are also moderators have moderator privileges, and may approve or reject their own review.
When `disable_self_approve` is enabled, authors who are moderators (or even users with admin privileges) cannot approve their own reviews.
- Project members can change the review's state to **Needs review** or **Needs revision**, and can attach committed changelists. Project members cannot change the review's state to **Approved**, **Rejected**, or **Archived**.
- Users that are not project members, moderators, or the review's author cannot transition the review's state.
- For the review's author and project members, if a review is not in one of their permitted states, for example if the review's state is **Rejected**, they cannot transition the review to another state.

These restrictions have no effect on who can start a review.

To delete a moderator, click the delete icon against their name.

4. Add default reviewers for a branch from the **Default reviewers** pane.

Important**When a review is part of multiple projects/project branches:**

- The default reviewer lists for all of the projects and project branches the review is part of are combined and added to the review.
- If a default reviewer has different reviewer options set on projects and project branches that the review is part of, the strictest reviewer option is used for the review.

Example: A review is created and it is part of **Project A**, **Project B**, and **Project Branch b**.

Project A: default reviewer X is an Optional reviewer

Project B: default reviewer X is an Optional reviewer

Project Branch b: default reviewer X is a Required reviewer

Result: default reviewer X is added to the review as a Required reviewer

Note

If users or groups are [@mentioned](#) in a new changelist description that includes **#review**, they will be added to the review as reviewers. If any of these reviewers are already specified as default reviewers they will not be added to the review again, the reviewer's most restrictive reviewer option is used for the review.

Note

If a default reviewer is deleted from Helix Core Server they will not be added to new reviews.

To add a default reviewer:

- Click **Manage default reviewers** button.
- Select the **Users** or **Groups** tab that you wish to add to the default reviewers list.
- Type in to search the name. The field auto-suggests groups and users within Helix Core Server as you type. Click **Add** button.

Toggle the **Retain default reviewers** switch to prevent default reviewers being removed from reviews associated with this branch. For more information about retaining default reviewers, see [Retain default reviewers](#).

To delete a default reviewer, click the delete icon against their name.

Integrations tab

Important

The project level test and deploy code features will be deprecated in a later Swarm release. We recommend you use test integration to automatically deploy code within a review. For more information, see ["Add a test" on page 544](#).

Tip

The project **Automated Tests** and **Automated Deployment** details are hidden from project members unless they are an owner or an administrator. This enables project members to check the project settings but not change them.

In the Integrations tab you can configure **Automated tests** and **Automated deployment** options.

Note

You can edit and save all options for automated tests and automated deployment when each of their toggles is disabled. However, the new updates take effect only when the toggles are enabled.

Automated tests

To enable or disable running automated tests for a project when reviews are created or updated, slide **Automated tests** toggle. When you wish to enable automated tests, enter a URL that triggers a test suite execution. Use the special arguments described in the dialog to help compose a URL that informs your test suite about important details.

Select the format of the POST parameters, either **URL Encoded** or **JSON Encoded**.

- **URL Encoded:** POST parameters are parsed into name=value pairs.
- **JSON Encoded:** parameters are passed raw in the POST body.

For more details about automated tests, see ["Automated testing for reviews" on page 567](#).

Automated deployment

To enable or disable deploying a project when reviews are created or updated, slide the **Automated deployment** toggle. When you wish to enable automated deployment, enter a URL that triggers a deployment of the project's code. Use the special arguments described in the dialog to help compose a URL that informs your deployment program with important details. For more details about automated deployment, see ["Automatically deploy code within a review" on page 94](#).

General Settings tab

In the General Settings tab, you can configure the following options:

1. Specify a project name.
2. Add a short description about the project.
3. Select a **Workflow** to associate with the project. Your selection is auto saved.

Note

The **Workflow** option is not displayed if the workflow feature is disabled.

To remove the existing workflow without replacing it with another workflow, select **No Workflow** from the **Workflow** dropdown list. This is the default when you create a new project.

To associate a workflow, select the workflow from the **Workflow** dropdown list, or enter the workflow name in the search field. The field auto-suggests workflows as you type. Only workflows that you own and shared workflows are shown in the dropdown list.

Tip

- When a workflow is associated with a project, the workflow is used for all of the branches in that project.
- When a project branch is associated with a workflow, the workflow of the parent project is ignored and the branch workflow is used.

For more information about workflows and how project workflows interact with branch workflows, see ["Workflow basics" on page 529](#).

4. Set the **Minimum up votes** for a project. The project setting is used on the project branches unless **Minimum up votes** is set to **1** or higher on a branch. In this case the branch setting is used and the project setting is ignored.

Enter the **Minimum up votes** or use the up and down arrow keys to increase or decrease the number of **Minimum up votes**. To save your changes click the green tick icon. Click the cross icon to delete your changes.

A review cannot be approved until all of the [Required reviewers](#) have voted up the review and the **Minimum up votes** specified has been satisfied.

- If a review spans projects/branches, the Minimum up votes for **each** of the projects and branches must be satisfied before you can approve the review.
- Required reviewers are included when up votes are counted.
- When **Count votes up from** is set to **Members** for a workflow associated with a project/branch, only the up votes of members of the project contribute to satisfying the **Minimum up votes** for a project/branch. For more information about the **Count votes up from** rule, see ["Workflow rules" on page 526](#).

Important

If the Workflow feature is disabled, all votes are counted not just votes from project members.

Important

If the **Minimum up votes** required is set higher than the number of reviewers that exist for a review, approval will be blocked for that review. This is true even if all the reviewers on the review have voted up the review.

5. Specify a **Job filter**.

The job filter allows you to specify criteria that are used to associate jobs with projects. For example, entering Subsystem=ProjectA associates jobs whose subsystem field is set to ProjectA with the current project.

Note

This job filter is simpler than the filters available in other Helix Core Server clients. The filter must be expressed as **field=value** pairs; bare keywords are not permitted. The asterisk for wildcard matching is permitted, but no other filter expression syntax is permitted.

6. Slide the **Private project** toggle to make this project private. Private projects and their associated reviews are only visible to project owners, moderators, and members, plus users with super privileges in Helix Core Server.

For more information, see ["Private projects" on page 387](#).

7. Delete a project as follows:

Note

For more information about deleting a project, see [Delete a project](#).

a. Click **Delete Project**.

A confirmation dialog appears to confirm whether you want to delete this project.

b. Click **Confirm** to delete the project.**Important**

It is possible to create a project that you cannot edit. This can happen if you have specified owners but not yourself as an owner, or if you have not specified yourself as a member. Swarm can detect some (but not all) such situations when you save a project; when it does detect such a situation, a warning dialog is displayed.

If you see this dialog, click **Continue** to save the project without your ownership/membership, or click **Cancel** within the dialog to continue editing the project. The project page's **Save** and **Cancel** buttons are disabled while this dialog is visible.

Edit a project

1. Visit the project page you want to edit.
2. Click **Settings** in the project menu.

A project summary with Participants tab, Branches tab, Integrations tab, and General Settings tab is displayed.

3. Edit the project's details as required.

For more information on how to use the Settings page, see ["Project settings" on page 499](#).

Tip

If sharing is switched off for a workflow, any projects or branches that were associated with the workflow while it was shared will remain associated with it.

If you edit a project or branch associated with that unshared workflow, you will still see the name of the workflow in the **Workflow** field, even if you don't own that workflow. If you remove that workflow from the project/branch you will not be able to see the workflow in the **Workflow** dropdown list unless you own it.

Note

By default, any member of a project can edit the project's configuration. Administrators can configure Swarm to prevent changes to the project's name and branch definition(s). See ["Projects" on page 670](#) for details.

Membership

Membership in a Swarm project identifies users as belonging to the project, making them part of the team.

There are only a few notable differences between project members and non-members:

Difference	Member	Non-Member	Description
"Notifications" on page 398	✓	✗	Members receive project notifications; non-members do not.
"Avatars" on page 413	✓	✗	The project's home page features member's avatars.
"States" on page 485	✓	✗	Members can transition code review states; non-members cannot.

There are two ways to become a member of a project in Swarm:

1. Add a project and make yourself a member.
2. Ask a member of an existing project to add you as a member.

Note

If the project has any [owners](#) specified, you need to ask a project owner to add you as a member.

Note


Users with *super* privileges in Helix Core Server can always adjust the settings for any project, including adjusting membership.

Add a member

If you are an owner of a project, or a member of a project without specified owners:

1. Visit the project page that needs the new member.
2. Click **Settings** in the project's toolbar.
3. The **Members** pane displays a list of all the project members.
4. Click **Manage members** button and specify a Helix Core Server user, Helix Core Server group, or Swarm project to add to the members for this project. The field auto-suggests *projectids*, *groupids*, and *userid*s by matching what you have typed so far against the list of users in the Helix Core Server.

When you specify a project or group, all of the members of that project or group become members of this project. Swarm does not display all of the individual users, but it does provide a visual separation: project or group names are displayed first, with a darker blue background.



5. Click **Add**.

Remove a member

If you are an owner of a project, or a member of a project without specified owners:

1. Visit the project page that has a member you want to remove.
2. Click **Settings** in the project's toolbar.
Known members of the project are displayed beneath the **Members** text field, with a medium blue button representing projects or groups and a light blue button representing individual users.
3. Click the **X** next to the project id, group id, or userid you want to remove.
4. Click **Save**.

Warning

You are able to remove your own membership or ownership. Doing so could prevent you from managing the project.

Owners

A project *owner* is a Helix Core Server user that controls the configuration for a project. An owner does not need to be a member of a project, but only an owner or user with *super* privileges in Helix Core Server can edit any project settings.

Tip

The project **Automated Tests** and **Automated Deployment** details are hidden from project members unless they are an owner or an administrator. This enables project members to check the project settings but not change them.

Moderators

A project *moderator* is a user assigned to moderate reviews for a specific branch associated with a project. See how to [specify moderators](#).

Once the branch specification is complete and the project has been saved, changing the state of any review associated with the moderated branch is restricted as follows:

- Only moderators can approve or reject the review. Moderators can also transition a review to any other state.

Important

Moderators prevent the automatic approval of reviews, for more information about automatically approving reviews using workflow rules see ["Workflow rules" on page 526](#).

Note

By default, when a review spans multiple branches that have different moderators, only one moderator from any one of the branches needs to approve the review.

Swarm can be configured to require that one moderator from each branch must approve the review, this is a global setting. If a moderator belongs to more than one of the branches spanned by the review, their approval will count for each of the branches they belong to. For instructions on how to configure moderator behavior, see ["Moderator behavior when a review spans multiple branches" on page 692](#).

- The review's author, when not a moderator, can change the review's state to **Needs review**, **Needs revision**, **Archived**, and can attach committed changelists.

Normally, the review's author cannot change the review's state to **Approved** or **Rejected** on moderated branches. However, authors that are also moderators have moderator privileges, and may approve or reject their own review.

When `disable_self_approve` is enabled, authors who are moderators (or even users with admin privileges) cannot approve their own reviews.

- Project members can change the review's state to **Needs review** or **Needs revision**, and can attach committed changelists. Project members cannot change the review's state to **Approved**, **Rejected**, or **Archived**.
- Users that are not project members, moderators, or the review's author cannot transition the review's state.
- For the review's author and project members, if a review is not in one of their permitted states, for example if the review's state is **Rejected**, they cannot transition the review to another state.

These restrictions have no effect on who can start a review.

Default reviewers

User and group default reviewers can be set for individual projects and project branches. Each time a new review is created in the project or project branch, the default reviewers will be added to the review. See [projects](#) and [project branches](#) for adding default reviewers.

- A user can be set as a [required reviewer](#) or an optional reviewer.
- A group can be set as a [required reviewer \(one vote\)](#), a [required reviewer \(all votes\)](#), or an optional reviewer.

Important

When a review is part of multiple projects/project branches:

- The default reviewer lists for all of the projects and project branches the review is part of are combined and added to the review.
- If a default reviewer has different reviewer options set on projects and project branches that the review is part of, the strictest reviewer option is used for the review.

Example: A review is created and it is part of **Project A**, **Project B**, and **Project Branch b**.

Project A: default reviewer X is an Optional reviewer

Project B: default reviewer X is an Optional reviewer

Project Branch b: default reviewer X is a Required reviewer

Result: default reviewer X is added to the review as a Required reviewer

Note

If users or groups are [@mentioned](#) in a new changelist description that includes **#review**, they will be added to the review as reviewers. If any of these reviewers are already specified as default reviewers they will not be added to the review again, the reviewer's most restrictive reviewer option is used for the review.

Note

If a default reviewer is deleted from Helix Core Server they will not be added to new reviews.

Retain default reviewers

By default, default reviewers can be removed from an individual review by using the [edit reviewers](#) button on the review display page. Individual projects and branches can be configured to prevent default reviewers from being removed from individual reviews. For instructions on how to enable **Retain default reviewers** for a project or branch, see [project](#) or [project branch](#).

Retain default reviewers basics

- Retained default reviewers can be added or removed by editing the project or project branch. Each time a review is updated in the project or branch, the list of default reviewers is checked.
 - Any new default reviewers are added to the review.
 - If a default reviewer has been removed from the project/branch, they will remain on the review. They are no longer a retained default reviewer and can be removed from the review if required.
 - If a review is updated and it is no longer associated with the project/branch, the default reviewers for that project/branch will remain on the review. They are no longer retained default reviewers and can be removed from the review if required.

A review is checked when:

- The [review state](#) changes
 - A new version of the review is created
 - A [comment task state](#) is changed
 - A user votes on the review, or clears their vote
 - [Reviewers are edited](#) on the review
 - The [review author is changed](#)
- Retained default reviewers cannot be removed from a review by [editing reviewers](#) in the review.
 - The retained default reviewer voting option cannot be reduced in a review by [editing reviewers](#) in the review.

Example: if a user is a [Required reviewer](#), the voting option cannot be reduced to a less strict option such as Optional reviewer.

- The retained default reviewer voting option can be made stricter for a review by [editing reviewers](#) in the review.

Example: if a group is an Optional reviewer, the voting option can be increased to a stricter option such as Required reviewer (one vote).

Reviews spanning multiple projects and branches

Related projects and branches

When a review spans related projects and branches and a user/group is a retained default reviewer in one but a standard default reviewer in the other, the branch voting option is used for the user/group reviewer in the review.

Example:

- **Review 5678:** spans **Project A**, and **Branch A-1**.
- **Project A:**
 - **User-X:** Required reviewer
 - **Retain default reviewers:** Enabled.
 - **Branch A-1:** a child of **Project A:**
 - **User-X:** Optional reviewer
 - **Retain default reviewers:** Disabled.
- **Review:**
 - **User-X:** Optional reviewer, not a retained default reviewer.

The project voting option is ignored and the branch voting option is used.

There is no restriction on changing the default reviewer voting option.

The default reviewer can be removed from the review.

Unrelated projects and branches

When an individual user/group default reviewer is retained for some of the projects and branches the review spans but not for others, the strictest voting option is used for a default reviewer on the review. The strictest retained default reviewer voting option is the minimum voting option for the user/group default reviewer in the review.

Example:

- **Review 1234:** spans **Project A**, **Project C**, and **Branch F-1**.
- **Project A:**
 - **Group-D:** Optional reviewer.
 - **Retain default reviewers:** Enabled.
- **Project C:**
 - **Group-D:** Required reviewer (all votes).
 - **Retain default reviewers:** Disabled.
- **Branch F-1:** not related to **Project A** or **Project C**:
 - **Group-D:** Required reviewer (one Vote)
 - **Retain default reviewer:** Enabled.
- **Review:**
 - **Group-D:** Required reviewer (All votes), retained default reviewer.
The strictest voting option is used for the default reviewer.
You can edit the voting option on the review. The minimum voting option is the strictest of the retained reviewer voting options. In this case that is, Required reviewer (one vote).
The reviewer cannot be removed from the review.

Delete a project

Note

- Users with *super* or *admin* privileges in Helix Core Server can delete projects.
- Project owners can delete projects that they own.
- If a project has no owners, any member of the project can delete the project.

When you delete a project:

- The deleted project is removed from the Swarm project list on the dashboard, and from project searches.
- The deleted project is removed from the profile page of the project owners, members, moderators, and followers.
- The project name is removed from the **Project** column on the ["Reviews list" on page 428](#) page.
- Reviews that belong to the deleted project are not changed. The open and closed reviews remain accessible, their [review states](#), [comments](#), and [tasks](#) can be modified as normal.
- Reviews that belong to the deleted project, the project branch name link in the [review heading](#) is replaced with a link to the common depot location that contains the files included in the review.

Note

The deleted project name cannot be reused for a new project. This behavior is not case sensitive, this means that if you delete a project called **Project B** you cannot create a new project called **project b**.

Use the following steps to delete a project:

1. Visit the project page you want to remove.
2. Click **Settings** in the project menu and navigate to the **General Settings** tab.
3. Click **Delete**.
A confirmation dialog appears to confirm whether you want to delete the project.
4. Click **Confirm** to delete.

8 | Workflows

This chapter covers Swarm workflow management, including:

- ["Why should I use Swarm workflow?" below](#)
- ["Workflow overview" on page 526](#)
- ["Add a workflow" on page 537](#)
- ["Edit a workflow" on page 542](#)
- ["Delete a workflow" on page 542](#)

Related information:

- For instructions on how to list, search, and view workflows, see ["Workflows" on page 388](#).
- For instructions on how to add a workflow to a project, see [Add a workflow to a project](#).
- For instructions on how to add a workflow to a branch, see [Add a workflow to a project branch](#).
- For a step-by-step walk-through of setting up a workflow, adding it to a project, and using it to progress through a review, see the [Helix Swarm User Workflow Guide](#).
- **Administrators:** For instructions on how to configure the global workflow rules, see [Workflow global rules](#).

Tip

For Swarm 2019.2 and later, the workflow feature is enabled by default. If the workflow feature is not working, check that the workflow prerequisites have been met, see ["Workflow prerequisites" on page 111](#).

Why should I use Swarm workflow?

By default Swarm's review process is basically advisory in nature with very few restrictions imposed on the review process. Team members can work in whatever way they want, but as projects grow in size and complexity, teams may prefer a stricter set of workflow rules over flexibility in order to improve productivity and quality. However, relying on people to stick to the agreed rules is not always enough.

The quality and speed of delivery of a product is heavily dependent on having a standardized and controlled workflow that follows best practice. It is also beneficial to automate the workflow where possible because the benefits of automating the workflow increases exponentially with scale. Teams within your organization will have varied workflow needs, it is important to maintain a balance between a company-wide policy and project specific needs. The Swarm workflow features have been designed to help you achieve this balance.

Tip

Additional Workflow rules and enhancements are on the road map for future Swarm releases.

The Swarm workflow feature enables you to enforce your rules and enables you to define specific rules for different projects and branches, providing you with flexibility and control over your review workflow. Workflows can be shared and they can be used on multiple projects/branches. Changing the workflow rules updates the workflow for all of the projects and branches it is associated with. This makes it easy to make changes in one place and change the workflow of multiple projects/branches. Specific workflow rules can be applied globally, ensuring basic company policies are followed by every project.

As a project owner you want to setup a project using the [Mainline model](#) with 'development', 'main' and 'release' branches. Each of these branches will have workflow rules that are tailored to its branch type.

Typically, you want your development branch to allow quick changes to the project content so that it can be populated quickly. Breaking test is acceptable at this point in the project. It is usually acceptable for a review to be raised after the content has been committed, this is the post-review commit model. Often you would only require one reviewer to vote up a review to trigger automatic approval. This minimal workflow configuration helps the development line to be built up quickly.

Tip

If your developers are pair programming, you might even allow them to commit changes without any review to further speed up the process. In the example shown below this is not the case and reviews are enforced by the workflow rules.

For your main branch you probably want to enforce pre-commit reviews, require more reviewers to vote up a review to trigger automatic approval, and restrict the up votes to project members only. This all helps to increase the stability of the main branch.

For a release branch you want even more protection for the content of the branch by only allowing critical changes to be committed to the branch. This is controlled by adding senior project members as moderators to the branch. Moderators are the only users that can approve or reject a review so they act as the gatekeepers for the branch. Moderators are added to the branch from the [project branch settings](#) pane.

Development branch example:

A Development branch needs to be fast moving, allowing new features, collateral, and technologies to be added very quickly with issues fixed as they occur. In essence this is rapid development allowing a lot of work to be completed quickly but the branch is unstable. The workflow for your development branch would be configured to allow for this style of working, typically:

- The branch rules allow post-commit and pre-commit reviews as required:
 - Changes can be committed and a review can be requested later, this is a post-commit review. This approach is particularly useful if you are making large changes and you just want to get content into the depot and then check it.

- Changes can be shelved with a review, this is a pre-commit review. This is useful when a user is adding standalone content or a feature that needs to be approved before it is committed to the depot.
- Only one reviewer needs to vote up a review to trigger automatic approval.

Main branch example:

A main branch needs to be more tightly controlled than the development branch as it is more stable as you move towards release. In essence, the main branch is used to fix any lingering issues and firm up the content and any tests you have in place. The workflow for your main branch would be configured to allow for this style of working, typically:

- Changelists must be shelved with a review, this is a pre-commit review.
- Up votes for reviews are only counted from project members.
- Two reviewers from the project are required to vote up a review to trigger automatic approval.
- The changelist can only be committed if the content of the changelist is identical to the content of the approved review.

Release branch example:

A Release branch typically needs to be very tightly controlled to protect the codeline with only a select set of users empowered to approve and commit reviews to the project before the product is released. In essence, the branch should be stable and you should not be making changes unless absolutely necessary. The workflow for your release branch would be configured to allow for this style of working, typically:

- Changelists must be shelved with a review, this is a pre-commit review.
- Senior members of the project are required to vote up the review before it can be approved.
- Moderators act as the final gatekeepers for review approval, only they have the power to approve a review once the required reviewers have voted the review up.
- The changelist can only be committed if the content of the changelist is identical to the content of the approved review.
- Approval is blocked if specific tests fail.

These are three very different workflows, it is easy for team members to forget the rules especially when they are new to the project or team. Swarm workflows enable you to enforce your rules and control your branches more effectively. No more last minute changes breaking your release build and no more post-commit reviews in your pre-commit world.

Workflow overview

Note

- For Swarm 2019.2 and later, the workflow feature is enabled by default. If the workflow feature is not working, ask your Swarm administrator to check that the workflow prerequisites have been met, see ["Workflow prerequisites" on page 111](#).
- The Swarm administrator can give individual groups and users special permission to ignore all of the Swarm workflow rules. Typically, these permissions are given to a small set of trusted groups and users, for example project owners and administrators. For more information about excluding users and groups from being bound by the workflow rules, see ["Members or groups who can ignore ALL workflow rules" on page 758](#).

Tip

To understand the benefits of using the Swarm Workflow feature, see ["Why should I use Swarm workflow?" on page 523](#)

A workflow can be applied to a project or project branch to ensure that changelists and reviews in the project/branch follow the rules specified in that workflow. The ["Workflow basics" on page 529](#) section gives an overview of global workflow rules, workflow rules and how they interact with each other.

Sections in this chapter:

- ["Workflow rules" below](#).
- ["Workflow basics" on page 529](#).
- ["Example workflows" on page 530](#).
- ["Merging multiple workflows" on page 532](#).

Workflow rules

A workflow is made up of the following workflow rules:

On commit without a review:

This rule is applied when a changelist without an associated review is submitted from outside of Swarm.

Select one of the following options:

- **Allow:** the changelist is not checked, the changelist is submitted without a review.
- **Create a review:** the changelist is submitted and a review is created automatically for the changelist.
- **Reject:** the changelist submit is rejected.

Tip

The selected rule is also applied when a changelist is submitted with **#review** in the description. For more information about creating a review by including a keyword in the changelist description, see ["Create a review" on page 682](#).

On commit with a review: This rule is applied when:

- A Swarm review is committed.
- A changelist with an associated review is submitted from outside of Swarm.

Select one of the following options:

- **Allow:** the changelist review state is not checked. The changelist is committed even if its associated review is not approved.
- **Reject unless approved:** the changelist is only submitted if its associated review is approved and the content of the changelist is identical to the content of the approved review.

Tip

The selected rule is also applied when a changelist is submitted with **#review-nnnnn**, **#replace-nnnnn**, or **#append-nnnnn** in the description (**nnnnn** = review ID). For more information about adding a changelist to a review by including a keyword in the changelist description, see ["Add a changelist to a review" on page 682](#).

On update of a review in an end state:

Used to stop review content being automatically changed for reviews that are in specific states. By default, the protected end states are **Archived**, **Rejected**, and **Approve:Commit**. The end states are set by the Swarm administrator, see [end_states](#).

Tip

When a review is in a protected end state, it can still be updated manually by a user from the Swarm UI.

This rule is applied when a changelist is added to a review.

Select one of the following options

- **Allow:** the review is not checked. The changelist is added to the review no matter what the review state is. This is the default setting.
- **Reject:** if the review is in one of the states specified in the **end_state** configurable:
 - The **Add Change** button is disabled for the review.
 - The changelist is rejected if it is added outside of Swarm using **#review-nnnnn**, **#replace-nnnnn**, or **#append-nnnnn** in the description (**nnnnn** = review ID).

Count votes up from:

By default, all of the up votes on a review are counted for the **Minimum up votes** value set on the project/branch the review is associated with. Limit the up votes that are counted to just the members of the project the review is associated with by using this rule.

This rule is applied when a user votes on a review.

Select one of the following options:

- **Anyone:** votes are counted for all reviewers on a review. This is the default setting.
- **Members:** only the up votes of members of the project the review is associated with are counted for the **Minimum up votes** set on projects/branches.

Tip

For instructions on how to set **Minimum up votes** for projects and branches, see [Project minimum up votes](#) and [Branch minimum up votes](#).

Automatically approve reviews:

By default, reviews must be manually approved. Enable automatic approval of reviews with this rule.

This rule is applied when a user votes on a review.

Select one of the following options:

- **Never:** reviews are not automatically approved. This is the default setting.
- **Based on vote count:** reviews are automatically approved if:
 - There are no down votes on the review.
 - There are no moderators on the review. If a review has moderators it cannot be automatically approved.
 - All of the [Required reviewers](#) on the review have voted up.
 - The **Minimum up votes** on the review has been satisfied for **each** of the projects and branches the review spans.
 - All of the tests configured to block approval have passed.

Important

Moderators prevent the automatic approval of reviews. For more information about moderators, see ["Moderators" on page 477](#).

Tip

After a review has been automatically approved it needs to be manually committed.

For information about creating a new workflow, see ["Add a workflow" on page 537](#).

Workflow basics

- **Workflows:** can be associated with projects and/or project branches.
 - A workflow can be created by any Swarm user.
 - If a workflow is associated with any projects or project branches, that workflow cannot be deleted.
 - A workflow can be shared by the workflow owner.
 - A shared workflow can be viewed by any Swarm user.
 - A shared workflow can be applied to a project or project branch by any Swarm user that is authorized to edit the project.
 - If sharing is switched off for a workflow, any projects or branches associated with that workflow will remain associated with it.
- **Projects and branches:**
 - When a workflow is associated with a project, the workflow is used for all of the branches in that project.
 - When a project branch is associated with a workflow, the workflow of the parent project is ignored and the branch workflow is used. The project workflow is replaced by the branch workflow, the workflows are not merged. For examples of why you might use this, see ["Restrictive project workflow with a less restrictive development branch workflow" on the next page](#), and ["Project workflow with a more restrictive release branch workflow" on page 531](#).
 - When a branch is associated with a workflow and a global workflow rule is set to **Enforce on** by the administrator, the workflow of the parent project is ignored and the global workflow rule is merged with the branch workflow. The most restrictive rule is used.
- **Modifying a workflow:**
 - If a global workflow rule is set to **Enforce off** and that rule is modified, projects and branches that do not have an associated workflow will use the modified rule.
 - If a global workflow rule is set to **Enforce on** and the rule is modified, the entire Helix Core Server will use the modified rule.
 - If a workflow is modified, any projects or branches associated with that workflow will use the modified workflow.
- **Global workflow rules:** configured by the Swarm administrator in [Global workflow rules](#), each rule can be set in one of two modes:
 - **Enforce off:** the global workflow rule is used for projects and branches that do not have an associated workflow.
 - If a project or branch has an associated workflow, the global rule is ignored for that project or branch.

- If projects are not used, the global workflow rule applies to the entire Helix Core Server.
- **Enforce on:** the global workflow rule is used for the entire Helix Core Server. This allows a minimum workflow to be set for the Helix Core Server.
 - The global workflow rule is merged with the workflow associated with a project or branch. The most restrictive rule is used.
- **Tests:**
 - If a test is associated with a workflow, it can be configured to run when a review associated with that workflow is either created/updated or submitted, or run manually from the review.
 - If a test is associated with the global workflow, it can be configured to run when a review is either created/updated or submitted, or run manually from the review.
 - If a test is associated with a workflow or a global workflow, it can be configured to block approval if it fails.

Example workflows

This section contains examples of typical workflows that can be configured in Swarm.

Enforce a pre-commit review workflow

To enforce a [pre-commit](#) review workflow set the rules to:

- **On commit without a review:** set to **Reject**, ensures that changelist cannot be submitted without an approved review.
- **On commit with a review:** set to **Reject unless approved**, ensures that if a changelist already has an associated review, that review is approved and the content of the changelist is identical to the content of the approved review.

Allow a post-commit review workflow

To allow a [post-commit](#) review workflow set the rules to:

- **On commit without a review:** set to **Create a review**, ensures that all submitted changelists have a review.
- **On commit with a review:** set to **Allow**, enables changelists that already have an associated review to be submitted.

Restrictive project workflow with a less restrictive development branch workflow

The release and main branches need a restrictive workflow to protect the codelines. All submissions must have an associated approved review and the content of the changelist must be identical to the content of the approved review.

The development branch can be less restrictive because you want to get features in quickly and review them after they have been committed. In this scenario you would probably want to ensure that all submitted changelists have an associated review.

Project workflow:

- **On commit without a review:** set to **Reject**, ensures that changelist cannot be submitted without an approved review.
- **On commit with a review:** set to **Reject unless approved**, ensures that if a changelist already has an associated review, that review is approved and the content of the changelist is identical to the content of the approved review.
- **On update of a review in an end state:** set to **Reject**, ensures that the content of reviews that are considered complete cannot be changed. These end states are set by the Swarm administrator, see ["Protected end states" on page 689](#).
- **Count votes up from:** set to **Members**, ensures that only the up votes of members of the project the review is associated with are counted for the **Minimum up votes** set on projects/branches.
- **Automatically approve reviews:** set to **Never**, ensures that reviews are not automatically approved.

Development branch workflow:

- **On commit without a review:** set to **Create a review**, ensures that all submitted changelists have a review.
- **On commit with a review:** set to **Allow**, enables changelists that already have an associated review to be submitted.
- **On update of a review in an end state:** set to **Reject**, ensures that the content of reviews that are considered complete cannot be changed. These end states are set by the Swarm administrator, see ["Protected end states" on page 689](#).
- **Count votes up from:** set to **Anyone**, ensures that votes of all reviewers on a review are counted for the **Minimum up votes** set on projects/branches.
- **Automatically approve reviews:** set to **Based on vote count**, ensures that reviews are automatically approved when all of the review requirements have been satisfied.

Project workflow with a more restrictive release branch workflow

The project is set up with workflow that allows post commit reviews, this allows you to get features in quickly and review them after they have been committed. In this scenario it is a good idea make all submitted changelists have an associated review.

The release branch needs a more restrictive workflow to protect the codeline. All submissions must have an associated approved review and the content of the changelist must be identical to the content of the approved review.

Project workflow: defines the workflow for the project and its branches.

- **On commit without a review:** set to **Create a review**, ensures that all submitted changelists have a review.
- **On commit with a review:** set to **Allow**, enables changelists that already have an associated review to be submitted.
- **On update of a review in an end state:** set to **Allow**, allows the content of reviews that are considered complete to be changed.
- **Count votes up from:** set to **Anyone**, ensures that votes of all reviewers on a review are counted for the **Minimum up votes** set on projects/branches.
- **Automatically approve reviews:** set to **Based on vote count**, ensures that reviews are automatically approved when all of the review requirements have been satisfied.

Release branch workflow: defines the workflow for the release branch. The workflow of the parent project is ignored and the development branch workflow is used.

- **On commit without a review:** set to **Reject**, ensures that changelist cannot be submitted without an approved review.
- **On commit with a review:** set to **Reject unless approved**, ensures that if a changelist already has an associated review, that review is approved and the content of the changelist is identical to the content of the approved review.
- **On update of a review in an end state:** set to **Reject**, ensures that the content of reviews that are considered complete cannot be changed. These end states are set by the Swarm administrator, see ["Protected end states" on page 689](#).
- **Count votes up from:** set to **Members**, ensures that only the up votes of members of the project the review is associated with are counted for the **Minimum up votes** set on projects/branches.
- **Automatically approve reviews:** set to **Never**, ensures that reviews are not automatically approved.

Merging multiple workflows

If a changelist spans multiple projects that have different workflows, the workflows are merged and the most restrictive workflow rules are used for the changelist. All of the tests associated with the merged workflows are run. If there are any duplicate tests they are only run once.

Example 1: Changelist 1234 spans Project A, Project B, and Project C

- **Project A:** `//depot/jam/lib/...`
- **Project B:** `//depot/jam/docs/...`
- **Project C:** `//depot/common/src/...`

■ **Changelist 1234** contains the following files:

- //depot/jam/lib/tests/checkLoad.sh
- //depot/jam/docs/basics.workflow.html
- //depot/common/src/index.html

The files span all three locations so the changelist is subject to the combined workflow rules for the three locations. The most restrictive workflow is used for the changelist:

	Changelist without a review	Changelist with a review	On update of a review in an end state	Count votes up from	Automatically approve reviews	Workflow Tests
Project A workflow	Allow	Reject unless approved	Allow	Anyone	Based on vote count	Smoke Test A
Project B workflow	Allow	Reject unless approved	Reject	Members	Never	Smoke Test B
Project C workflow	Create a review	Allow	Allow	Anyone	Based on vote count	Smoke Test B Smoke Test C
Changelist 1234 workflow	Create a review	Reject unless approved	Reject	Members	Never	Smoke Test A Smoke Test B Smoke Test C

Tip

The three workflows are combined and the most restrictive workflow rules are used for the changelist. All of the tests associated with the three workflows are run but **Smoke Test B** on **Project A** and **Project B** is only run once.

Example 2: Changelist 6789 spans Project M, Project N, and Project branch N-1

- **Project M:** //depot/thirdparty/lib/...
- **Project N:** //depot/thirdparty/lib/tests/...
- **Project branch N-1:**
 - //depot/projectN/...
 - //depot/thirdparty/lib/...
- **Changelist 6789** contains the following files:
 - //depot/thirdparty/lib/tests/checkReturns.sh
 - //depot/projectN/src/index.html
 - //depot/projectN/tests/pageLoadTest.php

The changelist files span all three locations so the changelist is subject to the combined workflow rules for the three locations. The most restrictive workflow is used for the changelist:

	Changelist without a review	Changelist with a review	On update of a review in an end state	Count votes up from	Automatically approve reviews	Workflow Tests
Project M workflow	Create a review	Reject unless approved	Allow	Anyone	Never	Smoke Test M
Project N workflow	Reject	Reject unless approved	Reject	Members	Never	Smoke Test N
Project branch N-1 workflow	Create a review	Allow	Allow	Anyone	Based on vote count	Smoke Test N-1
Changelist 6789 workflow	Create a review	Reject unless approved	Allow	Anyone	Never	Smoke Test M Smoke Test N-1

Tip

It is important to understand that the rules for **Project N** are ignored because **Project branch N-1** is a branch of **Project N** and it has its own workflow rules. This is why the result for **Changelist without a review** is **Create review**, the result for **On update of a review in an end state** is **Allow**, and the result for **Count votes up from** is **Anyone**.

Example 3: Changelist 3456 spans Project X, and Project Y

- **Project X:** //depot/main/product/...
- **Project Y:** //depot/main/product/docs/...
- **Changelist 3456** contains the following files:
 - //depot/main/product/tests/fileSelector.sh
 - //depot/main/product/docs/admin.html

The changelist files span both locations so the changelist is subject to the combined workflow rules for the two locations, and the [Global workflow rules](#) set by the administrator. The most restrictive workflow is used for the changelist:

	Changelist without a review	Changelist with a review	On update of a review in an end state	Count votes up from	Automatically approve reviews	Workflow Tests
Project X workflow	Create a review	Allow	Allow	Anyone	Based on vote count	Smoke Test X
Project Y workflow	Reject	Allow	Reject	Anyone	Based on vote count	Smoke Test Y
Global workflow rules	Create a review (Enforce on)	Reject unless approved (Enforce off)	Allow (Enforce on)	Members (Enforce on)	Never (Enforce off)	Global Smoke Test Global Full Test

	Changeli st without a review	Changeli st with a review	On update of a review in an end state	Count votes up from	Automatic ally approve reviews	Workflo w Tests
Changeli st 3456 workflow	Reject	Allow	Reject	Member s	Based on vote count	Smoke Test X Smoke Test Y Global Smoke Test Global Full Test

Tip

- Global workflow rules:

- When a global workflow rule is set to **Enforce** off, the rule is only used if the project/branch does not have an associated workflow. If the project/branch has an associated workflow, the global rule setting is ignored. This is why the result for the **Changelist with a review** rule is **Allow**, and the **Automatically approve reviews** is **Based on vote count**.
- When a global workflow rule is set to **Enforce** on, the rule is merged with project/branch workflow and the most restrictive rule setting is used. This is why the result for the **Changelist without a review** rule is **Reject**, the result for the **On Update of a review in an end state** rule is **Reject**, and the result for **Count votes up from** is **Members**.

The global workflow rule mode determines how the global rule interacts with the workflow rules. For more information about global rule modes, see ["Workflow basics" on page 529](#).

- If one or more of the project branches has a moderator, the moderator prevents automatic approval of the review. For more information about moderators, see ["Moderators" on page 477](#).

Add a workflow

Note

- Swarm workflows can be created by any Swarm user.
- Shared Swarm workflows can be viewed by any Swarm user.
- Shared Swarm workflows can be applied to a project or project branch by any Swarm user that is authorized to edit the project.
- The global workflow can be viewed by any Swarm user but can only be edited by a Swarm user with *super* or *admin* privileges, or users that have been made an owner.

To add a workflow:

- On the Swarm **Workflows** page, click the **+ Add Workflow** button.

The **Add Workflow** page is displayed:

The screenshot shows the 'Add Workflow' form. It has a title bar 'Add Workflow' with a close button 'x'. The form is divided into several sections:

- Name:** A text input field with a red border.
- Description:** A large text area.
- Owners:** A section with a user icon and a text input 'Add an Owner'.
- Individuals:** A section with a blue pill 'swarm' and a close button 'x'.
- Shared with others:** A checkbox that is currently unchecked.
- Rules:** A section with a title 'Rules' and a help icon. It contains five rows, each with a label and a dropdown menu:
 - On commit without a review: Allow
 - On commit with a review: Allow
 - On update of a review in an end state: Allow
 - Count votes up from: Anyone
 - Automatically approve reviews: Never
- Tests:** A section with a title 'Tests' and a help icon. It contains a blue link '+Add Test' and a message 'There are no tests associated with this workflow'.
- When:** A section with a title 'When' and a message 'There are no tests associated with this workflow'.
- Blocks:** A section with a title 'Blocks'.
- Footer:** A section with the text 'No projects use this workflow' and two buttons: 'Save' and 'Cancel'.

- Enter a name for the workflow.
- Optional:** provide a description for the workflow.
- You are automatically added as the owner of the workflow.

Optional: add more owners if required. This field auto-suggests groups, and users within Helix Core Server as you type (up to a combined limit of 20 entries).

Important

- A workflow must have at least one owner.
- If you remove yourself as an owner, you cannot edit this workflow configuration later unless you have *super* user rights.

5. **Optional:** if you want other Swarm users to be able to use this workflow, select the **Shared with others** checkbox.

Tip

Leave the checkbox unselected until you have proved that the workflow rules work as expected. This keeps the workflow private and stops other Swarm users using the workflow until you are happy with it. Once you are happy with the workflow, select the checkbox to share the workflow with other users.

6. **Rules:**

On commit without a review:

This rule is applied when a changelist without an associated review is submitted from outside of Swarm.

Select one of the following options:

- **Allow:** the changelist is not checked, the changelist is submitted without a review.
- **Create a review:** the changelist is submitted and a review is created automatically for the changelist.
- **Reject:** the changelist submit is rejected.

Tip

The selected rule is also applied when a changelist is submitted with **#review** in the description. For more information about creating a review by including a keyword in the changelist description, see ["Create a review" on page 682](#).

On commit with a review: This rule is applied when:

- A Swarm review is committed.
- A changelist with an associated review is submitted from outside of Swarm.

Select one of the following options:

- **Allow:** the changelist review state is not checked. The changelist is committed even if its associated review is not approved.
- **Reject unless approved:** the changelist is only submitted if its associated review is approved and the content of the changelist is identical to the content of the approved review.

Tip

The selected rule is also applied when a changelist is submitted with `#review-nnnnn`, `#replace-nnnnn`, or `#append-nnnnn` in the description (**nnnnn** = review ID). For more information about adding a changelist to a review by including a keyword in the changelist description, see ["Add a changelist to a review" on page 682](#).

On update of a review in an end state:

Used to stop review content being automatically changed for reviews that are in specific states. By default, the protected end states are **Archived**, **Rejected**, and **Approve:Commit**. The end states are set by the Swarm administrator, see [end_states](#).

Tip

When a review is in a protected end state, it can still be updated manually by a user from the Swarm UI.

This rule is applied when a changelist is added to a review.

Select one of the following options

- **Allow:** the review is not checked. The changelist is added to the review no matter what the review state is. This is the default setting.
- **Reject:** if the review is in one of the states specified in the `end_state` configurable:
 - The **Add Change** button is disabled for the review.
 - The changelist is rejected if it is added outside of Swarm using `#review-nnnnn`, `#replace-nnnnn`, or `#append-nnnnn` in the description (**nnnnn** = review ID).

Count votes up from:

By default, all of the up votes on a review are counted for the **Minimum up votes** value set on the project/branch the review is associated with. Limit the up votes that are counted to just the members of the project the review is associated with by using this rule.

This rule is applied when a user votes on a review.

Select one of the following options:

- **Anyone:** votes are counted for all reviewers on a review. This is the default setting.
- **Members:** only the up votes of members of the project the review is associated with are counted for the **Minimum up votes** set on projects/branches.

Tip

For instructions on how to set **Minimum up votes** for projects and branches, see [Project minimum up votes](#) and [Branch minimum up votes](#).

Automatically approve reviews:

By default, reviews must be manually approved. Enable automatic approval of reviews with this rule.

This rule is applied when a user votes on a review.

Select one of the following options:

- **Never:** reviews are not automatically approved. This is the default setting.
- **Based on vote count:** reviews are automatically approved if:
 - There are no down votes on the review.
 - There are no moderators on the review. If a review has moderators it cannot be automatically approved.
 - All of the [Required reviewers](#) on the review have voted up.
 - The **Minimum up votes** on the review has been satisfied for **each** of the projects and branches the review spans.
 - All of the tests configured to block approval have passed.

Important

Moderators prevent the automatic approval of reviews. For more information about moderators, see ["Moderators" on page 477](#).

Tip

After a review has been automatically approved it needs to be manually committed.

7. Tests

Optional: add tests to the workflow, the tests are run when a review associated with the workflow is either created/updated or submitted. Selected from the **When** dropdown for the test.

Tip

- Tests are only saved on the workflow when you click the **Save** button.
- If a review is associated with multiple workflows, all of the tests on all of those workflows are run for the review. If the same test is on more than one of the associated workflows, it is only run once.

Add a test:

- a. Click **+Add Test** in the Tests area.
- b. Select the test to run from the **Tests** dropdown list.
- c. Select when you want the test to run from the **When** dropdown list.

Tip

- Review content change is when the files or content of files in a review change. A change to the review description does not trigger a test.
- **On Demand** tests are unaffected by the review content change check.


- **On Update** the test will run when:

- a review is created.
- a review is submitted and the review content has changed compared to the previous review version.
- a review is updated and the review content has changed since the previous review version.



Tip

If the review content has not changed since a test last successfully reported a pass, it will not be re-run when a review is versioned or submitted. If a test is still running or has not updated Swarm with a pass state, it will be re-run, even on a submit.

- **On Submit** the test will run when a pre-commit review is submitted or a post-commit review is created.
- **On Demand** the test will not run automatically, but you can run it manually from the ["Tests" on page 451](#) dialog on the review page.

- d. Select whether Swarm blocks approval if the test fails from the **Blocks** dropdown list:
 - **Nothing** if the test fails it will not block approval.
 - **Approve** if the test fails it will prevent the review from being approved.
- e. Click the **Accept**  button to add the test to the workflow.

Edit a test on the workflow:

- a. Click the **Edit**  button next to the test you want to edit.
- b. The **Test**, **When**, and **Blocks** dropdown lists are displayed for the test.
- c. Make changes as required.
- d. Click the **Accept**  button to confirm your changes.

Remove a test from the workflow:

Click the **Delete**  button next to the test to remove it from the workflow.

The test is removed immediately, no confirmation is requested. You must save the workflow to complete the test removal.

8. Click **Save**.

Note

The **Save** button is disabled if any required fields are empty.

Edit a workflow

Important

Changes made to a workflow will change the workflow for projects and project branches that use that workflow.

Note

- A workflow can only be edited by the owner of the workflow or a user with *super* user rights.
- To edit the **Global Workflow**, see "[Workflow global rules](#)" on page 752.

1. Visit the **Workflow** page you want to edit.
2. Edit the workflow details as required, see "[Add a workflow](#)" on page 537 for more details.
3. Click **Save**.

Delete a workflow

Note

- A workflow can only be deleted by the owner of the workflow, or a user with super user rights.
- The **Global Workflow** cannot be deleted.

To delete a workflow:

1. Navigate to the workflow you want to delete.
2. Make sure that there are no projects, or branches associated with the workflow.

If projects, or branches are associated with the workflow, the **Delete** button is muted, and the workflow cannot be deleted.

For instructions on how to remove the workflow from a project, or branch, see [Associate a workflow with the project](#) and [Associate a workflow with the project branch](#).

Note

Due to bug in Swarm 2019.3 and earlier, a workflow might falsely report that it is linked to a project when it is not. To remove the false link, navigate to the [Settings page](#) of the project reported as linked and click the **Save** button.

3. Click **Delete** to delete the workflow.

9 | Tests

Important

The **Tests** page is only available if Workflow is enabled, workflow is enabled by default.

This chapter covers Swarm test management, including:

- ["Add a test" below](#)
- ["Edit a test" on page 551](#)
- ["Delete a test" on page 551](#)
- To add a test to a workflow, see ["Add a workflow" on page 537](#).

Related information:

- For instructions on how to list, search, and view tests, see ["Tests" on page 395](#).
- **Administrators:** For instructions on how to add tests the global workflow, see ["Global workflow" on page 753](#).

Add a test

Note

- Swarm tests can be created by any Swarm user.
- Shared Swarm tests can be viewed by any Swarm user.
- Shared Swarm tests can be added to a workflow by any Swarm user that is authorized to edit the workflow.
- Shared Swarm tests can be added to the global workflow by any Swarm user that is authorized to edit the global workflow.

Note

In earlier versions of Swarm, global tests were set in the Swarm `config.php` file. From Swarm 2020.1, tests are added to the global workflow on the Swarm **Workflows** page so that they operate as global tests. If you upgrade from an earlier version of Swarm, your global tests are automatically migrated. Each global test is migrated as follows: the owner is set as the Swarm admin user (not shared), the name is set to the original test name, the description is set to the original test title, and it is added to the **Global Workflow**.

Associate a test with a [workflow](#) to ensure that the test is run when a review associated with that workflow is either created/updated or submitted. Associate a test with the [global workflow](#) to ensure that the test is run whenever a Swarm review is either created/updated or submitted. This ensures that the global tests are enforced for all changes even if they are not part of a project.

To add a test:

1. On the Swarm **Tests** page, click the **+ Add Test Definition** button.

The **Add Test Definition** page is displayed:

2. Enter a name for the test. This should be a unique name that is 1 to 32 characters in length.
3. **Optional:** provide a description for the test.
4. You are automatically added as the owner of the workflow.

Optional: add more owners if required. This field auto-suggests groups, and users within Helix Core Server as you type (up to a combined limit of 20 entries).

Important

- A test must have at least one owner.
- If you remove yourself as an owner, you cannot edit this test configuration later unless you have *super* user rights.

5. **Optional:** if you want other Swarm users to be able to use this test, select the **Shared with others** checkbox.

Tip

Leave the checkbox unselected until you have proved that the test works as expected. This keeps the test private and stops other Swarm users using the test until you are happy with it. Once you are happy with the test, select the checkbox to share the test with other users.

6. Enter the test request **URL** that will trigger your test suite in the URL text box.

Special arguments are also available to pass details from Swarm to your test suite, see ["Pass special arguments to your test suite" on page 548](#).

Note

Helix Plugin for Jenkins 1.10.11 and later: Swarm must send the parameters for the build to Jenkins as a POST request. To do this, enter all the **ParameterName=ParameterValue** pairs in the **Body** and select **URL Encoded**. Separate each **ParameterName=ParameterValue** pair with an ampersand & sign. For more information about the supported parameters, see [Building jobs](#) in the *Helix Plugin for Jenkins Guide*.

7. **Optional:** set a **Timeout (in seconds)** as an integer. The timeout is used when Swarm connects to your test suite. If a timeout is not set, the `http_client_options timeout` setting is used. By default, this is 10 seconds.
8. **Optional:** specify parameters that your test suite requires in the request **Body**, select the encoding method using the radio buttons:
 - **URL Encoded:** POST parameters are parsed into name=value pairs. This is the default.
 - **JSON Encoded:** parameters are passed raw in the POST body.
 - **XML Encoded:** parameters are passed in XML format in the POST body.

Special arguments are also available to pass details from Swarm to your test suite, see ["Pass special arguments to your test suite" on page 548](#).

9. **Iterate tests for affected projects and branches** checkbox

Optional: select to create a separate test run for each branch and project the review is associated with. The details generated by the `{projects}`, `{branches}`, and `{branch}` arguments in the **URL** or **Body** are used in the test name displayed in the ["Tests" on page 451](#) dropdown on the review page.

For example: if you have the following `{projects}` and `{branches}` arguments in your test **URL** or **Body**:

```
'{projects}' => 'foo,bar',
'{branches}' => 'foo:dev,foo:main,bar:dev,bar:candidate',
```

Swarm will iterate through the arguments and create a test run for each of these project branches:

- Project foo dev branch
- Project foo main branch
- Project bar dev branch
- Project bar candidate branch

If **Iterate tests for affected projects and branches** is not selected, only one test will be run for the review, even if multiple branches and projects are associated with the review.

Note

Private projects: if a test for a private project is added to a review because **Iterate tests for affected projects and branches** is selected for the test, Swarm honors the private project's permissions and displays it as **Private project** in the review's test list to users that do not have permission to view it.

- For information about the reviews test list, see ["Tests" on page 451](#).
- For information about private projects, see ["Private projects" on page 387](#).

10. Headers

Optional: enables you to specify name=value pairs to pass to the test suite.

Tip

Headers are only saved on the test when you click the **Save** button.

Add a header name=value pair:

- a. Click [+Add header](#) in the headers area.
- b. Enter the name=value pair in the empty **Header** and **Value** boxes.
- c. The new header is saved when you click **Save** on the test page.

Edit a header name=value pair:

- a. Edit the name=value pair in the **Header** and **Value** boxes.
- b. The changes to the header are saved when you click **Save** on the test page.

Remove a header name=value pair from the test:

Click the **Delete**  button next to the header to remove it from the test.

The header is removed immediately, no confirmation is requested. You must save the test to complete the header removal.

11. Click **Save**.

Note

The **Save** button is disabled if any required fields are empty.

Pass special arguments to your test suite

You can include special arguments in the request **URL** and **Body** to pass Swarm information about the change that triggered the test run to your test suite. Swarm automatically replaces the arguments with the relevant Swarm information when it calls the test:

Note

The curly braces { } are part of the arguments and must be used.

- { **change** } the change number
- { **status** } the status of the shelved change, shelved or committed
- { **review** } the review identifier
- { **version** } the version of the review
- { **reviewStatus** } the Swarm status of the review: **needsReview**, **needsRevision**, **archived**, **rejected**, **approved**, **approved:commit**
- { **description** } the change description of the change used to generate this update. **{description}** cannot be used in the **URL**, it can only be used in the request **Body**.
- { **test** } the name of the test
- { **testRunId** } the test run id
- { **projects** } the project identifiers of projects that are part of the review, null if the review is not part of a project. Comma separated if more than one project is involved in the review.
- { **branches** } the branch identifiers of branches that are part of the review. Branch identifiers are prefixed by a project identifier (with a colon : separator by default) if the branch is part of a project, null if the branch is not part of a project. Comma separated if more than one branch is involved in the review.

Tip

Some CI systems, such as Jenkins, escape the default `:` character resulting in a failed test request. If your CI system needs a different separator character, your Swarm *admin* can configure Swarm to use that character.

Be careful when selecting your separator character, some characters can cause issues when calling tests. For example, a `#` character in a URL call means the branch is treated as an anchor and a `%` character in a JSON body encoding is escaped. Both of these will result in a failed test request.

For instructions on configuring the separator character, see ["project_and_branch_separator" on page 734](#).

- **{branch}** the branch identifiers impacted by the review, comma-separated.

Note

This argument is intended for use in [project tests](#) and is supported by workflow tests for backward compatibility. However, we recommend you use the **{branches}** argument for workflow tests going forwards.

- **{update}** the update callback URL. You can include any or all of the following when calling the update url to update the test run: status, messages, and a url in the body that links to the CI system for that run. They should be formatted in JSON in the body of the POST request. **Status:** valid status values are **running**, **pass**, and **fail**. **Messages:** you can pass a maximum 10 messages, if you provide more than 10 messages only the first 10 are saved. Each message can contain a maximum of 80 characters, any messages with more than 80 characters will be automatically truncated. **{update}** is the preferred option for Swarm 2019.3 and later.

Important

If your test system cannot POST to Swarm, you cannot use **update** and you must use **pass** and **fail** instead.

Note**When using {update}:**

- If your build script has access to any messages related to the test execution, pass the messages to Swarm using the **{update}** URL. Swarm uses the provided message(s) to add to the test results.
- If your build script has access to the results of test execution, include a POST parameter called **url** when calling the update URL. Swarm uses the provided url to link reviews to the test results. Valid test status values are **running**, **pass**, and **fail**.
- The **{update}** callback url accepts a JSON body where you can specify any or all of the following: messages, url, and status.

```
{
  "messages" : ["My First Message", "My Second Message"],
  "url" : "http://jenkins_host:8080/link_to_run",
  "status": "pass"
}
```

- **{pass}** the tests pass callback URL. From Swarm 2019.3 and later, **{update}** is preferred. For more details, see the note below.
- **{fail}** the tests fail callback URL. From Swarm 2019.3, **{update}** is preferred. For more details, see the note below.

Note

- **{update}**, **{pass}**, and **{fail}** are composed automatically by Swarm, they include Swarm's own per-review authentication tokens.
- **When using {pass} and {fail}**: if your build script has access to the results of test execution, include a GET or POST parameter called **url** when calling the pass or fail URLs. Swarm uses the provided url to link reviews to the test results.

Tip

Swarm 2019.3 and later still supports **{pass}** and **{fail}**, however **{update}** is preferred because you can also include a message with the test status.

- **{projectNames}** the project names of projects that are part of the review, null if the review is not part of a project. Comma separated if more than one project is involved in the review.

Note

Not iterated when the **Iterate tests for affected projects and branches** checkbox is selected for the test. If you want to iterate tests, use the `{projects}` argument.

- `{branchName}` the branch name(s) impacted by the review, comma-separated.

Note

Not iterated when the **Iterate tests for affected projects and branches** checkbox is selected for the test. If you want to iterate tests, use the `{branches}` argument.

Edit a test

Important

Changes made to a test will change the test for workflows that use it.

Note

A test can only be edited by the owner of the test or a user with *super* user rights.

1. Navigate to the **Test** page you want to edit.
2. Edit the test details as required, see ["Add a test" on page 544](#) for more details.
3. Click **Save**.

Delete a test

Note

A test can only be deleted by the owner of the test or a user with super user rights.

To delete a test:

1. Navigate to the test you want to delete.
2. Make sure that there are no workflows associated with the test.
If workflows are associated with the test, the **Delete** button is disabled, and the test cannot be deleted.
For instructions on how to remove the test from a workflow, see ["Add a workflow" on page 537](#).
3. Click **Delete** to delete the test.

10 | Integrations

Swarm integrates with a variety of other applications and processes to provide important functionality, such as automated testing and deployment of code in reviews, issue tracking, file previewing, and downloading archives. This chapter describes each of the integrations available with Swarm.

Included integrations:

- [JIRA](#)
- [LibreOffice](#)
- [Zip archive](#)
- [Automated deployment for reviews](#)
- [Automated testing for reviews](#)
- [P4V](#)
- [Slack integration](#)
- [TeamCity](#)

Jira

Swarm's Jira integration allows code reviews and committed changes to be associated with Jira issues, making it easy to reference associated issues, and see the state of a code review or committed change within Jira. To associate a code review with a Jira issue, include a Jira issue identifier in the review's description, for example **SW-1234**. Swarm links to the Jira issue and creates a link within the Jira issue back to the code review in Swarm. Multiple Jira issues can be included in the changelist description. As a code review progresses, Swarm updates each associated Jira issue with the review's current status.

For instructions on how to configure basic Jira integration, see ["Jira integration" on the next page](#).

Additionally, if you have the Helix Defect Tracking Gateway (P4DTG) configured to integrate Perforce jobs with Jira issues. Swarm can be configured to add a link to a Perforce job within a Jira issue. Perforce job links are created when a new Jira or job is created or updated. This gives you direct access to the Perforce job from a link in the **Issue Links** section of the Jira issue.

For instructions on how to configure Jira integration and Perforce job links, see ["Jira integration and Perforce job links" on page 556](#).

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

Jira integration

Basic Jira integration enables:

- **In Swarm:** Hyperlinks are created to Jira issues referenced in changelist descriptions and review descriptions. Jira links in comments are only supported in the original Swarm review page. For example, if Jira issue **SW-1234** is found in the text, the text is linked to the Jira issue in Swarm.

For the legacy Jira linkify, the links for Jira issues in changelist descriptions, review descriptions, and comments open in a new tab in the web browser.
- **In the Issue Links section of the Jira issue:** Creates hyperlinks to the associated Swarm changelist (commit) and review pages.

Note

Swarm fetches Jira project identifiers using a worker, once per the worker process' lifetime. See ["Worker configuration" on page 750](#). New Jira projects will not auto-link to Jira until the project identifiers have been updated. By default, project identifiers are updated once every ten minutes.

Prerequisites for basic Jira integration

- The Swarm workers must be working, see ["Worker status" on page 750](#).
- The Jira module must be enabled, see ["Enabling the Jira module" below](#).
- The Jira user specified in the [Jira module](#) configuration must have permission to view issues, update issues, and add remote links to issues in Jira projects where the Jira integration is required.

To troubleshoot the permissions:

1. Log in to Jira as the Swarm Jira user.
2. Retrieve an issue.
3. Try to create a link.

Enabling the Jira module

By default, the Jira module is not enabled. To enable the Jira module, add the following configuration block to your `SWARM_ROOT/data/config.php` file:

```
<?php
// this block should be a peer of 'p4'
'jira' => array(
    'host'          => '', // URL for your installed Jira web
interface (start with https:// or http://)
    'api_host'      => '', // URL for Jira API access, 'host' is
```

```

used for Jira API access if 'api_host' is not set
    'user'          => '', // Jira Cloud: the username or email
                        address used to connect to your Atlassian account
                        // Jira on-premises: the username
required for Jira API access
    'password'      => '', // Jira Cloud: a special API token
                        obtained from https://id.atlassian.com/manage/api-tokens
                        // Jira on-premises: the password
required for Jira API access
    'relationship' => '', // Jira subsection name links are added
                        to defaults to empty, links added to the "links to" subsection
),

```

- **host**: URL for your Jira web interface starting with **https://** or **http://**
For example, for a Jira cloud instance it might be: **https://my-company-jira.atlassian.net**
- **api_host** (optional): URL for Jira API access, the **host** URL is used for Jira API access if **api_host** is not set.
- **user**: The user must have permission to view issues, update issues, and add remote links to issues in Jira projects where the Jira integration is required.
 - **Jira Cloud**: Username or email address used to connect to your Atlassian account.
 - **Jira on-premises**: Username for Jira API access.
- **password**:
 - **Jira Software Cloud**: A special API token. To obtain your API token:
 - a. Go to <https://id.atlassian.com/manage/api-tokens>
 - b. Log in with your Atlassian administrator credentials.
 - c. Click **Create** and name your API token something recognizable, for example **swarmjira**.
 - d. The API token is displayed in a dialog.
 - e. Copy the token and paste it into the **password** configurable. It is a good idea to save the API token somewhere safe in case you need it in the future.
 - **Jira on-premises**: Password for Jira API access.
- **relationship**: sets which subsection of the Jira **Issue Links** section Swarm commit, Swarm review, and Perforce job links are added to. P4DTG is required for Perforce job links, see ["Jira integration and Perforce job links" on page 556](#).

- '' empty: links are added to the **links to** subsection of **Issue Links** section of Jira issues. This is the default value.
- `<JobLinksSubsectionName>`: creates a new subsection in the **Issue Links** section of Jira issues. Links are added to this subsection.

Note

Existing links are not moved to the subsection defined by the **relationship** configurable. However, if a link is updated or a new link is added, the link is added to the new subsection. The link is not deleted from its original location.

Jira SSL client configuration

Tip

If your Jira site certificate is from one of the more common Certificate Authorities (such as the CA used by Jira Software Cloud), the following configuration changes are not usually required.

If your Jira site certificate is not from one of the more common Certificate Authorities, you might need to configure the Swarm [http_client_options](#) configuration block in your `SWARM_ROOT/data/config.php` file.

For example, if your root and intermediate certificates are stored in a **pem** format file in `/path/to/certs/jira.pem`, point to the file with:

```
<?php
// this block should be a peer of 'p4'
'http_client_options' => array(
    'hosts' => array(
        'jira.example.com' => array(
            'sslcafile' => '/path/to/certs/jira.pem',
            'sslpassphrase' => 'keep my JIRA secure',
            'timeout' => 15,
        ),
    ),
),
),
)
```

For more information about configuring your Jira HTTP client options, see ["HTTP client options" on page 706](#).

If you need to troubleshoot Swarm's connection with your Jira instance, see [Swarm to JIRA Connection Troubleshooting](#) in the Perforce Knowledgebase.

Jira integration and Perforce job links

Jira integration and Perforce job links enable:

- **In Swarm:**
 - Creates hyperlinks to Jira issues referenced in changelist descriptions and review descriptions. Jira links in comments are only supported in the original Swarm review page. For example, if Jira issue **SW-1234** is found in the text, the text is linked to the Jira issue in Swarm.
 - Creates hyperlink to associated Jira issues on the Swarm job page.
- **In the Issue Links section of the Jira issue:**
 - Creates hyperlinks to the associated Swarm changelist (commit) and review pages.
 - When a Perforce job is created, creates a hyperlink to the Swarm job page.
 - If a job is added to a review, creates a hyperlink to the Swarm review page.
 - For fixes (**p4 fix**) to a job, when the `job_field` is set, creates a hyperlink to the Swarm changelist (commit) page.
 - When P4DTG replication detects a change to an issue, creates a job link for the associated Jira issue if it does not already exist.

For a summary of Jira integration and Perforce job link workflow, see "[Jira integration and Perforce job link workflow](#)" on page 559.

Prerequisites for Jira integration and Perforce job links

- The Helix Defect Tracking Gateway (P4DTG) must be configured for Swarm, Jira, and P4D, see the [Helix Defect Tracking Gateway Guide](#).
- The Swarm workers must be working, see "[Worker status](#)" on page 750.
- The Jira module must be enabled, see "[Enabling the Jira module](#)" below.
- To add Perforce job links to Jira issues, job links must be enabled and the Jira `job_field` must be set, see "[Enabling the Jira module](#)" below.

Enabling the Jira module

Important

If Swarm fails to add a Perforce job link to a Jira issue because that Jira issue does not exist, the failure is logged by Swarm. Swarm will not try to add the job link to that Jira issue again.

Note

If a Perforce job is updated and the `DTG_DTISUE` field value is changed, a job link will be added to the new Jira issue. The Perforce job link will not be removed from the original Jira issue.

By default, the Jira module is not enabled. To enable Perforce job links in Jira issues, add the following configuration block to your `SWARM_ROOT/data/config.php` file:

```
<?php
// this block should be a peer of 'p4'
'jira' => array(
    'host'          => '', // URL for your installed Jira web
                        // interface (start with https:// or http://)
    'api_host'      => '', // URL for Jira API access, 'host' is
                        // used for Jira API access if 'api_host' is not set
    'user'          => '', // Jira Cloud: the username or email
                        // address used to connect to your Atlassian account
                        // Jira on-premises: the username
                        // required for Jira API access
    'password'      => '', // Jira Cloud: a special API token
                        // obtained from https://id.atlassian.com/manage/api-tokens
                        // Jira on-premises: the password
                        // required for Jira API access
    'job_field'     => '', // if P4DTG is replicating Jira issue
                        // IDs to a job field, list that field here
    'link_to_jobs'  => false, // set to true to enable Perforce
                        // job links in Jira, P4DTG and job_field required
    'delay_job_links' => 60, // delay in seconds, defaults to 60
                        // seconds
    'relationship'  => '', // Jira subsection name links are
                        // added to defaults to empty, links added to the "links to" subsection
),
```

- **host**: URL for your Jira web interface starting with **https://** or **http://**
For example, for a Jira cloud instance it might be: **https://my-company-jira.atlassian.net**
- **api_host** (optional): URL for Jira API access, the **host** URL is used for Jira API access if **api_host** is not set.
- **user**: The user must have permission to view issues, update issues, and add remote links to issues in Jira projects where the Jira integration is required.
 - **Jira Cloud**: Username or email address used to connect to your Atlassian account.
 - **Jira on-premises**: Username for Jira API access.

- **password:**
 - **Jira Software Cloud:** A special API token. To obtain your API token:
 - a. Go to <https://id.atlassian.com/manage/api-tokens>
 - b. Log in with your Atlassian administrator credentials.
 - c. Click **Create** and name your API token something recognizable, for example **swarmjira**.
 - d. The API token is displayed in a dialog.
 - e. Copy the token and paste it into the **password** configurable. It is a good idea to save the API token somewhere safe in case you need it in the future.
 - **Jira on-premises:** Password for Jira API access.
- **job_field:** If P4DTG is replicating Jira issue IDs to a job field, list the job field here.
- **link_to_jobs:**
 - **false:** disable Perforce job links in Jira. The default setting is **false**.
 - **true:** enable Perforce job links in Jira.
- **delay_job_links:** delay set in seconds. Sets the delay between a new Jira or job being created or updated, and Swarm adding the job link to the Jira issue. This configurable is used to avoid a race condition between Jira and the Helix Core Server. By default **delay_job_links** is set to 60 seconds, if a race condition is experienced increase the delay time.
- **relationship:** sets which subsection of the Jira **Issue Links** section Swarm commit, Swarm review, and Perforce job links are added to. P4DTG is required for Perforce job links.
 - **' '** empty: links are added to the **links to** subsection of **Issue Links** section of Jira issues. This is the default value.
 - **<JobLinksSubsectionName>**: creates a new subsection in the **Issue Links** section of Jira issues. Links are added to this subsection.

Note

Existing links are not moved to the subsection defined by the **relationship** configurable. However, if a link is updated or a new link is added, the link is added to the new subsection. The link is not deleted from its original location.

Jira SSL client configuration

Tip

If your Jira site certificate is from one of the more common Certificate Authorities (such as the CA used by Jira Software Cloud), the following configuration changes are not usually required.

If your Jira site certificate is not from one of the more common Certificate Authorities, you might need to configure the Swarm [http_client_options](#) configuration block in your [SWARM_ROOT/data/config.php](#) file.

For example, if your root and intermediate certificates are stored in a **pem** format file in **/path/to/certs/jira.pem**, point to the file with:

```
<?php
// this block should be a peer of 'p4'
'http_client_options' => array(
    'hosts' => array(
        'jira.example.com' => array(
            'sslcafile' => '/path/to/certs/jira.pem',
            'sslpassphrase' => 'keep my JIRA secure',
            'timeout' => 15,
        ),
    ),
),
)
```

For more information about configuring your Jira HTTP client options, see ["HTTP client options" on page 706](#).

If you need to troubleshoot Swarm's connection with your Jira instance, see [Swarm to JIRA Connection Troubleshooting](#) in the Perforce Knowledgebase.

Jira integration and Perforce job link workflow

This section describes the workflow for Jira integration and job links in Jira issues when you request a new Swarm review from P4V.

Tip

The process for adding a Perforce job and creating a review differs for the other clients but the Swarm worker and Helix Defect Tracking Gateway (P4DTG) workflow is the same.

1. Create a new Jira issue in your Jira system. P4DTG detects the new Jira issue and creates a new Perforce job for that Jira issue.
2. Add the new Perforce job to a pending changelist in P4V. See [Add a job to a pending changelist](#) in the [Helix Core Visual Client \(P4V\) Guide](#).
3. Request a new Swarm review for the pending changelist in P4V. See [Request a review](#) in the [Helix Core Visual Client \(P4V\) Guide](#).

P4V passes the review request to the Swarm worker queue. The review request is processed by the next available Swarm worker, the new review is created. Swarm checks its Jira configuration. If the `job_field` exists in the Perforce job, Swarm adds a link to the new Perforce job in the **Issue Links** section of the Jira issue.

To check that these steps have been completed and to view the links:

1. Open the new Swarm review. The associated Perforce jobs are displayed on the review page below the review description.
2. Click the Perforce job link to open the job. The Swarm review link is displayed on the job page below the job description.
3. Click the **Details** tab on the Perforce job page to view the **DTG_DTISSUE** field.
4. Click the Jira link in the **DTG_DTISSUE** field to view the Jira issue. The Swarm review and Perforce job links are displayed in the **Issue Links** section of the Jira issue page.

LibreOffice

LibreOffice is a free power-packed open source personal productivity suite. When LibreOffice is installed on the server hosting Swarm, Swarm automatically detects its presence and uses LibreOffice to prepare PDF previews of a variety of file types, including:

- Word documents (`.doc`, `.docx`)
- PowerPoint presentations (`.ppt`, `.pptx`)
- Excel spreadsheets (`.xls`, `.xlsx`)
- Visio diagrams (`.vsd`)
- Rich-text files (`.rtf`).



Depending on your server's platform and distribution, LibreOffice may be provided as multiple packages and not all packages may be installed by default. If certain filetypes do not preview as expected in Swarm, you may need to install these optional packages to include all of the file handling capabilities of LibreOffice.

For more information on LibreOffice, see <https://www.libreoffice.org>.

Limitations

The LibreOffice integration has several limitations:

- Document previews in Swarm may appear different than on a desktop system if the document's fonts are not installed on the server hosting Swarm and LibreOffice. In addition, LibreOffice has some limitations rendering Microsoft Office file types, so LibreOffice-generated previews may differ from what you see using Microsoft Office.
- Large files may require a notable amount of time to preview. Very large files may exhaust the resources of the server hosting Swarm, causing the preview to fail and temporarily impacting Swarm performance.
- Swarm is currently not able to detect or show differences in LibreOffice-supported file types.

Installation

We recommend that you install LibreOffice from your OS distribution, via `apt-get`, `yum`, etc.

The minimal packages, and their transitive dependencies required for Swarm are:

- `libreoffice-calc`
- `libreoffice-draw`
- `libreoffice-impress`
- `libreoffice-writer`
- `libreoffice-headless` (CentOS/RHEL only)

Zip archive

When the `zip` command-line tool is installed on the Swarm server, Swarm allows users to download ZIP archives of files and folders. You can download a Zip archive using the **Download zip** option from the following Swarm pages:

- ["Review display" on page 433](#) page.
- ["Changelists" on page 338](#) page.
- ["Files" on page 322](#) page.

Installation

Install the zip command-line tool from your operating system distribution using `apt-get`, `yum`, etc.

Configuration

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

Configure the archiving feature with the following configuration block in the [SWARM_ROOT/data/config.php](#) file:

```
<?php
// this block should be a peer of 'p4'
'archives' => array(
    'max_input_size'      => 512 * 1024 * 1024, // 512M (in bytes)
    'archive_timeout'     => 1800,              // 30 minutes
    'compression_level'   => 1,                 // 0-9
    'cache_lifetime'      => 60 * 60 * 24,      // 1 day
),
```

The `max_input_size` key specifies the maximum file/folder content size that can be processed into a ZIP archive. The default value permits up to 512 megabytes of content to be compressed. Smaller values limit the amount of file/folder content but provide faster downloads; larger values can allow increased scanning, syncing, compressing, and downloading times.

The `archive_timeout` key specifies the amount of time, in seconds, to allow Swarm to prepare the ZIP archive for downloading. Shorter times can limit the practical size of a ZIP archive, depending on the performance of your network and the filesystem hosting Swarm; even with a generous `max_input_size` setting, if `archive_timeout` seconds have elapsed, the archive operation is terminated.

The `compression_level` key specifies the compression level to use, and must be within the range 0 to 9. 0 means no compression, 9 means maximum compression. As this value is increased, smaller ZIP archives may result, but may require greater compression time. Swarm uses the default of 1, which provides a reasonable tradeoff of fast compression times with light compression that can still result in an archive notably smaller than the original file/folder content.

The `cache_lifetime` key specifies the desired maximum age of cached ZIP archives. Increasing the value increases the amount of time that ZIP archives exist in the cache, which can improve the user experience for frequently downloaded files. However, ZIP archives can be quite large (depending on the size of your depot within the Helix Core Server) and can require significant disk storage. Decreasing the value can mitigate the amount of disk space required for the cache; the tradeoff is that frequently accessed ZIP archives may need to be generated more frequently, which can have an impact on CPU and disk resources.

Download files as a Zip archive

To download the zip archive navigate to the review, changelist, file or folder:

Note

The **Download zip** option is not displayed if the zip command-line tool is not installed on the Swarm server. For information about installing, and configuring the zip command-line tool, see [Zip archive](#).

When you select the **Download zip** option, Swarm performs the following steps:

1. Scans the files/folders:
 - Checks that you have permission to access their contents, according to the Helix Core Server protections.
 - Checks that the total file size is small enough to be processed by Swarm.
2. Syncs the file contents to the Swarm server from the Helix Core Server.
3. Creates the ZIP archive by compressing the file content.
4. Starts a download of the generated ZIP archive.

Note

- You might not see all of the above steps; Swarm caches the resulting ZIP archives so that repeated requests to the same files/folders can skip the sync and compress steps whenever possible.
- If an error occurs while scanning, syncing, or compressing, Swarm indicates the error.

TeamCity

Swarm's TeamCity integration allows you to configure TeamCity to post build statuses as comments to your code reviews of shelved files.

For instructions on how to configure Swarm with TeamCity, see [Integrating TeamCity with Perforce](#).

11 | Administration

This section covers administration and configuration of Swarm.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the "Reload Configuration button" on page 732.

Automated deployment for reviews	566
Automated testing for reviews	567
Avatars	575
Backups	577
Changelist files limit	577
Client integration	578
Comment attachments	580
Comment mentions	581
Comments	585
Commit credit	586
Commit-edge deployment	587
Commit timeout	589
Configuration overview	589
Diff configuration	602
Email configuration	602
Emoji	608
Environment	609
Excluding Users from Activity Streams	613
Extensions Management	614
Files configuration	622
Groups configuration	624
Ignored users for reviews	625
Job keywords	625
License	628
Localization & UTF8 character display	629
Logging	631
Mainline branch identification	636
Markdown	639
Menu helpers	640
Multiple Helix Core Server instances	646
Notifications	662
Obliterate Review	668
P4TRUST	669

Projects	670
Redis server	673
Review cleanup	678
Review keyword	680
Reviews filter	686
Reviews	687
Search	696
Security	698
Short links	711
Single Sign-On PHP configuration	712
Skip activity	718
Slack integration	719
swarm_root	724
System Information	725
Tag processor	732
Test definitions	734
Trigger options	735
Unapprove modified reviews	742
Uninstall Swarm	742
Upgrade index	745
Users	746
Workers	749
Workflow global rules	752

Automated deployment for reviews

Important

The project level test and deploy code features will be deprecated in a later Swarm release. We recommend you use test integration to automatically deploy code within a review. For more information, see ["Add a test" on page 544](#).

Deploying code in a code review automatically involves enabling **Automated Deployment** in your project's configuration and providing a *trigger URL*. When the *trigger URL* is requested, Swarm expects a deployment program to be executed.

When the deployment processing ends, Swarm expects either a *success callback URL* or *failure callback URL* to be requested by your deployment program. These callback URLs should include a url parameter (either via GET or POST); when a valid-looking URL is included, clicking the deployment status indicator directs the user to the specified URL. This is intended to facilitate easy viewing of the successfully deployed review, or a report indicating why the deployment failed. The url parameter is mandatory for successful deployments, but is optional for failures.

1. Navigate to the project page.
2. Click the project **Settings** tab to display the **Project Settings** page.

3. **Automated Deployment** checkbox: select **Enable** to display the configuration field:

Automated Deployment ☒ Enable

`http://deploy-server/deploy?change={change}`

A URL that will trigger a deployment when reviews are created or updated. Some special [arguments](#) are supported. [See help for more details.](#)

4. Provide a URL that triggers your deployment execution.

Special arguments are available to inform your deployment program of various details from Swarm:

Note

Helix Plugin for Jenkins 1.10.11 and later: Swarm must send the parameters for the build to Jenkins as a POST request. To do this, enter the parameters in the **Post Body** and select **URL Encoded**.

{change}
The change number

{status}
Status of the change, *shelved* or *submitted*

{review}
The review's identifier

{project}
The project's identifier

{projectName}
The project's name

{branch}
The branch identifier(s), comma-separated

{branchName}
The branch name(s), comma-separated

{success}
Deployment successful callback URL

{fail}
Deployment failure callback URL

Automated testing for reviews

Important

The project level test and deploy code features will be deprecated in a later Swarm release. We recommend you use test integration to automatically deploy code within a review. For more information, see ["Add a test" on page 544](#).

Tip

From Swarm 2020.1, the preferred method for defining tests is on the Swarm ["Tests" on page 544](#) page. This enables you to:

- Associate a test with a [workflow](#) to ensure that the test is run when a review associated with that workflow is either created/updated or submitted.
- Associate a test with the [global workflow](#) to ensure that the test is run whenever a Swarm review is either created/updated or submitted. This ensures that the global tests are enforced for all changes even if they are not part of a project.

Integrating Helix Swarm with a test suite involves enabling **Automated Tests** in your project's configuration and providing a *trigger URL*. When the *trigger URL* is requested, Swarm expects your test suite to be executed. When the tests complete, Swarm expects an *update callback URL*, a *pass callback URL*, or a *fail callback URL* to be requested by your test suite.

1. Navigate to the **Project** page.
2. Click the project **Settings** tab to display the **Project Settings** page.
3. Ensure that paths in each named branch configured for the project do not overlap with paths in other named branches.
4. **Automated tests** checkbox: select **Enable** to display the configuration fields:

Automated Tests ☒ Enable

`http://test-server/build?change={change}`

A URL that will trigger automated tests to run when reviews are created or updated.

Some special [arguments](#) are supported. [See help for more details.](#)

POST Body

`foo=bar&baz=buzz`

URL Encoded ▼

Optional data to POST to the above URL. The special arguments supported for URLs can also be used here.

5. Provide a URL that triggers your test suite execution.

Special arguments are available to inform your test suite of various details from Swarm. Swarm automatically replaces the arguments with the relevant Swarm information when it calls the test:

Note

Helix Plugin for Jenkins 1.10.11 and later: Swarm must send the parameters for the build to Jenkins as a POST request. To do this, enter the parameters in the **Post Body** and select **URL Encoded**.

```
{test}
    The name of the test
{testRunId}
    The test run id
{change}
    The change number
```

Tip

If your CI system supports **change=now**, you can use this instead of the **change={change}** to make sure your review is always tested against the latest version of the submitted code.

```
{status}
    Status of the change, shelved or submitted
{review}
    The review's identifier
{version}
    The version of the review
{description}
    The change description of the change used to generate this update. {description}
    cannot be used in the URL, it can only be used in the POST Body.
{project}
    The project's identifier
{projectName}
    The project's name
{branch}
    The branch identifier(s) impacted by the review, comma-separated
{branchName}
    The branch name(s) impacted by the review, comma-separated
{update}
    The update callback URL. You can include any or all of the following when calling the
    update url to update the test run: status, messages, and a url in the body that links to the
    CI system for that run. They should be formatted in JSON in the body of the POST
    request. Status: valid status values are running, pass, and fail. Messages: you
    can pass a maximum 10 messages, if you provide more than 10 messages only the first
    10 are saved. Each message can contain a maximum of 80 characters, any messages
    with more than 80 characters will be automatically truncated. {update} is the preferred
    option for Swarm 2019.3. For more details, see the note below.
{pass}
```

Tests pass callback URL. From Swarm 2019.3, **{update}** is preferred. For more details, see the note below.

{fail}

Tests fail callback URL. From Swarm 2019.3, **{update}** is preferred. For more details, see the note below.

Note

- Swarm 2019.3 and later still supports **{pass}** and **{fail}**, however **{update}** is preferred because you can also include a message with the test status.
- **{update}**, **{pass}**, and **{fail}** are composed automatically by Swarm. They include Swarm's own per-review authentication tokens.

6. **Optional:** specify any parameters that your automated tests require that must be sent via HTTP POST in the **POST Body** field. The POST parameters can include the special arguments listed above.

Select the format of the POST parameters, either **URL Encoded** or **JSON Encoded**.

- **URL Encoded:** POST parameters are parsed into name=value pairs.
- **JSON Encoded:** parameters are passed raw in the POST body.

Configuring Jenkins for Swarm integration

Important

Your Jenkins host needs to be able to communicate with the Swarm host, and the Swarm host needs to be able to communicate with the Jenkins host. Ensure that the appropriate DNS/host configuration is in place, and that each server can reach the other via HTTP/HTTPS.

1. Install the p4-plugin for Jenkins:

For instructions on installing the p4-plugin, see in the [Installation](#) chapter of the [Helix Plugin for Jenkins Guide](#).

2. Configure a Jenkins project:

Note

Helix Plugin for Jenkins 1.10.11 and later: Swarm must send the parameters for the build to Jenkins as a POST request. To do this, enter the parameters in the **Post Body** and select **URL Encoded**.

- a. Specify the job name so that it matches the project identifier used in the trigger URL, as defined [below](#).

For example, the computed value of `{projectName}_{branchName}`.

Or, edit the trigger URL to use the Jenkins job name you specify.

- b. Make the build parameterized to accept these parameters (note that these are named to match up with the script that is called):

```
{test}
    The name of the test
{testRunId}
    The test run id
{status}
    Whether the changelist to be tested is shelved or submitted
{change}
    Changelist # to run tests against
{review}
    The review's identifier
{version}
    The version of the review
{branchName}
    The branch name(s) impacted by the review, comma-separated
{update}
    The URL to POST to update the test run. You can include any or all of the following
    when calling the update url to update the test run: status, messages, and a url in the
    body that links to the CI system for that run. They should be formatted in JSON in
    the body of the POST request. Status: valid status values are running, pass,
    and fail. Messages: you can pass a maximum 10 messages, if you provide more
    than 10 messages only the first 10 are saved. Each message can contain a
    maximum of 80 characters, any messages with more than 80 characters will be
    automatically truncated. {update} is the recommended way of reporting test
    status.
```

Important

If your test system cannot POST to Swarm, you cannot use **update** and you must use **pass** and **fail** instead.

Note**When using {update}:**

- If your build script has access to any messages related to the test execution, pass the messages to Swarm using the **{update}** URL. Swarm uses the provided message(s) to add to the test results.
- If your build script has access to the results of test execution, include a POST parameter called **url** when calling the update URL. Swarm uses the provided url to link reviews to the test results. Valid test status values are **running**, **pass**, and **fail**.
- The **{update}** callback url accepts a JSON body where you can specify any or all of the following: messages, url, and status.

```
{
  "messages" : ["My First Message", "My Second
Message"],
  "url" : "http://jenkins_host:8080/link_to_run",
  "status": "pass"
}
```

{pass}

The URL to wget if the build succeeds. From Swarm 2019.3, **{update}** is preferred. For more details, see the note below.

{fail}

The URL to wget if the build fails. From Swarm 2019.3, **{update}** is preferred. For more details, see the note below.

Note

When using {pass} and {fail}: if your build script has access to the results of test execution, include a GET or POST parameter called **url** when calling the pass or fail URLs. Swarm uses the provided url to link reviews to the test results.

For example:

```
?url=https://jenkins.example.com/view/job/<
jobname>/<jobid>/
```

Tip

Swarm 2019.3 and later still supports **{pass}** and **{fail}**, however **{update}** is preferred because you can also include a message with the test status.

- c. Select **Perforce Software** for the **Source Code Management** section.

Important

You might see **Perforce** in the **Source Code Management** section. This represents an earlier community-provided Perforce plugin that does not include support for Swarm.

- d. Set up credentials and workspace behavior as needed.

For instructions on configuring credentials and workspaces, see the [Credentials](#) and [Workspaces](#) sections of the [Helix Plugin for Jenkins Guide](#).

Important

- If your Helix Core Server is configured for Helix Authentication Service, the service user credentials used for automated testing must not use Helix Authentication Service.
- The client workspace configured in Jenkins must have a view that includes the paths defined for that branch in Swarm.

3. Configure your Swarm project to run automated tests with a URL and parameters like this:

Tip

In this example, the build system is passed the changelist number of the latest version of the review in `{change}`. If the test is rerun at a later date, the build system will rerun the test against the same `{change}` number, and this will probably not be the latest version of the codebase. If you want the review to always be tested against the latest versions of submitted code, replace `change={change}` with `change=now`. Your CI system must support `change=now` to use this feature.

Automated Tests ☒ Enable

`http://jenkins_host:8080/job/{projectName}_{branchName}/review/build`

A URL that will trigger automated tests to run when reviews are created or updated.

Some special [arguments](#) are supported. [See help for more details](#).

Optional data to POST to the above URL.

POST Body

`status={status}&review={review}&change={change}&update={update}`

URL Encoded ▼

Important

For Jenkins, the job name needs to match the job identifier in the URL. In the example above, this is the computed value of `{projectName}_{branchName}`.

If you prefer a different naming scheme in Jenkins, replace `{projectName}_{branchName}` in the URL above with the project name actually defined in Jenkins.

Important

If security is enabled in Jenkins, the trigger URL needs to include credentials. Follow these steps:

- Create a Jenkins user that will trigger Swarm builds. For example swarm.
- Log into Jenkins as the new user.
- Click on the user's username in the Jenkins toolbar.
- Scroll down to **API Token**.
- Click **Show API Token**.
- Incorporate the value of the **API Token** into the Swarm trigger URL.

For example, if the username is swarm and the API Token value is 832a5db7e5500c1288324c1441460610, the Swarm trigger URL and parameters should be:

Automated Tests ☒ Enable

```
https://swarm:832a5db7e5500c1288324c1441460610@jenkins_host:8080/job/{projectName}_{branchName}/review/build
```

A URL that will trigger automated tests to run when reviews are created or updated.

Some special [arguments](#) are supported. [See help for more details.](#)

Optional data to POST to the above URL.

POST Body

```
status={status}&review={review}&change={change}&pass={pass}&fail={fail}
```

URL Encoded ▼

Avatars

Swarm uses *avatars*, images that represent users and groups responsible for events in activity streams, projects, reviews, etc.

Avatars are retrieved from an avatar provider; the default provider is [gravatar.com](https://www.gravatar.com). Swarm sends an identifier to the avatar provider (for **gravatar.com**, an MD5 hash of the user's or group's email address), and the provider returns the configured image (if one exists). If no avatar is defined with the provider or the requests fails for any reason, Swarm selects an avatar from its internal collection.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

Configure the avatar lookups in the avatars configuration block in the [SWARM_ROOT/data/config.php](#) file. Here is an example:

```
<?php
    // this block should be a peer of 'p4'
    'avatars' => array(
        'http_url'  => 'http://www.gravatar.com/avatar/{hash}?s={size}&d={default}',
        'https_url' => 'https://secure.gravatar.com/avatar/{hash}?s={size}&d={default}',
    ),
```

- **http_url**: specifies the **http** URL Swarm uses to lookup avatars. If the URL is not defined, Swarm uses the default **gravatar.com** URL.
- **https_url**: specifies the **https** URL Swarm uses to lookup avatars. If the URL is not defined, Swarm uses the default **gravatar.com** URL.

The **avatars** array URL that Swarm uses depends on the settings of the Swarm Apache server protocol and the ["external_url" on page 611](#) protocol, the securest of these protocols is used.

Several replacement values are available for inclusion in the URLs:

- **{user}**
The current Swarm userid, Perforce groupid, or empty string
- **{email}**
The current Swarm user's or group's email address, or empty string

- **{hash}**

The MD5 hash of the Swarm user's or group's email address, or 00000000000000000000000000000000 if no email address is configured

- **{default}**

The value blank for a transparent GIF (allowing users or groups without avatars to fallback to Swarm's internal avatars) or the value **mm** for a *mystery man* used in circumstances where no user or group identifier is known

- **{size}**

the size Swarm would like in pixels for both the width and height, without units, e.g. 64

The URL you specify must include one of **{user}**, **{email}**, or **{hash}** to properly select a user-specific or group-specific avatar. The URL should include **{size}** to assist Swarm's presentation. **{default}** is not necessary, but helps provide a consistent avatar experience.

Tip

You can adjust the appearance of avatars by using custom CSS, see ["Adjust the appearance of avatars" on page 791](#).

Note

By default, gravatar.com serves only G-rated avatar images. If your Swarm users and groups wish to use PG-, R-, or X-rated images, you need to configure the avatar lookup URLs with the appropriate rating flag. For example, to allow avatars with G or PG ratings, the configuration would look like:

```
<?php
    // this block should be a peer of 'p4'
    'avatars' => array(
        'http_url' => 'http://www.gravatar.com/avatar/
{hash}?r=pg&s={size}&d={default}',
        'https_url' => 'https://secure.gravatar.com/avatar/
{hash}?r=pg&s={size}&d={default}',
    ),
```

For more information on gravatar.com image requests, see: <https://en.gravatar.com/site/implement/images>

Disable avatar lookups

If you wish to disable avatar lookups altogether and simply use Swarm's internal avatars, set each URL to ' ' (empty string). For example:

```
<?php
    // this block should be a peer of 'p4'
```

```
'avatars' => array(
    'http_url' => '',
    'https_url' => '',
),
```

Backups

Swarm stores all of the information it requires to operate within the Helix Core Server. This includes project definitions, code reviews, comments, followers, and more. Code reviews are largely built on top of Helix Core Server's shelving feature, and most other records are stored in custom counters called *keys*.

Therefore, the standard recommendations for backing up your Helix Core Server also apply when backing up your Swarm data.

For more information about backup and recovery, see [Backup and recovery in *Helix Core Server Administrator Guide*](#).

In addition, your Swarm configuration and any modifications you might make to the provided modules, templates, CSS, JavaScript, etc. also need to be backed up. The `SWARM_ROOT/data` directory contains the configuration, as well as temporary working files and browser session storage.

Changelist files limit

Changelists or reviews with many files present a challenge to Swarm; it can take a notable amount of time for the file listing to reach Swarm from the Helix Core Server, and a notable amount of time to provide the file listing HTML to a user's browser. When the file listing is many thousands of files, Swarm may run out of memory. Increasing the amount of memory for Swarm can help, but the Swarm interface may work slowly or not at all depending on the amount of memory available to the user's browser.

With the 2015.1 release, Swarm limits the number of files presented, even if a changelist or review involves more files than the limit.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

You can adjust the limit to suit your needs, by specifying a value for `max_changelist_files` in the `SWARM_ROOT/data/config.php` file:

```
<?php
    'p4' => array(
        'max_changelist_files' => 1000,
    ),
```

The default value of `1000` should suffice for most Swarm installations. Consider the effects of changing this value:

- Problems in the Swarm UI still exist with arbitrarily large values; there is likely no advantage to using a value over `10000`.
Performance will be better with Helix Core Server 2014.1 or later, as the file listing will be limited by Helix Core Server.
- Smaller values can potentially interfere with reviews and reading changelists; the file limit may cause interesting files in a changelist or review to no longer be displayed in Swarm.

Client integration

Note Helix Swarm 2022.3 or later only works with Helix Core Visual Client (P4V) 2021.3 or later.

P4V and P4VS can integrate with Swarm. To indicate how these applications should connect with Swarm, Swarm sets the `P4.Swarm.URL` property set in Helix Core Server. P4V and P4VS read this property, and if set, they connect to the specified URL to make Swarm API calls. If the property is unset, Swarm integration features are disabled.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

When `P4.Swarm.URL` is set, P4V provides the following integration features:

- **Request a review**: requests a review for pending or committed changelists.
- **Update a review**: updates a review from the current state of a pending changelist. This works for changelists that are already associated with a review, or for unassociated changelists.
- **Open review in Swarm**: opens the review associated with the selected changelist in your system's default web browser.
- **Review Id and State columns**: adds **Review Id** and **Review State** columns to both the **Pending** and **Submitted** tabs.

By default, the first Swarm worker auto-detects the URL it is running under and sets `P4.Swarm.URL` accordingly. Swarm checks and updates this property every 10 minutes.

Customized Swarm installations

For customized Swarm installations, the auto-detected URL might not use the correct protocol, hostname, or port. In these scenarios, you can disable the URL auto-detection by editing the `SWARM_ROOT/data/config.php` file and setting the `auto_register_url` item to `false` in the `p4` configuration block. For example:

```
<?php
    'p4' => array(
        'auto_register_url' => false,
    ),
```

If you disable this feature, you should:

- "Manually set P4.Swarm.URL in the Helix Core Server" below
- "Set the external_url in Swarm" below

Manually set P4.Swarm.URL in the Helix Core Server

If you disable `auto_register_url`, you should manually set the `P4.Swarm.URL` property in Helix Core Server to the URL for your Swarm installation:

```
p4 property -a -n P4.Swarm.URL -v https://myswarm.url:port/
```

Replace `https://myswarm.url:port/` with the URL for your Swarm installation.

To find out what the current `P4.Swarm.URL` value is for all users, run:

```
p4 property -Al -n P4.Swarm.URL
```

Note

P4V uses an integration timeout, specified in the `P4.Swarm.Timeout` property, to limit delays in the P4V user interface. The default timeout is 10 seconds.

To change the integration timeout, run:

```
p4 property -a -n P4.Swarm.Timeout -v 10
```

Replace the `10` with the desired timeout in seconds. Increasing the timeout could cause notable delays in the P4V user interface, and decreasing the timeout could cause sporadic integration failures if Swarm's API responses take longer than the specified timeout.

Set the external_url in Swarm

If you disable `auto_register_url`, you should usually set the `external_url` for Swarm. For information about setting the `external_url` for Swarm, see "external_url" on page 611.

Comment attachments

Swarm supports attaching arbitrary files to comments in code reviews and jobs.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button"](#) on page 732.

Comment attachment storage

To store files attached to comments, Swarm looks for a depot named `// .swarm`. As Swarm does not create this depot, you need to create it, or specify another depot that the Swarm *admin* user can write to.

The depot used to store Swarm attachments must be:

- a **classic depot** (**stream**, **unload**, **archive**, and **spec** depots are not supported)
- a **local** depot (**remote** depots are not supported)
- a depot that the Swarm *admin* user can write to but other users cannot access. For more information about permissions, see [Authorizing access](#) in the *Helix Core Server Administrator Guide*.

Tip

It is best practice to use a depot directory that is not used to store any other files, this stops other users having access to the attachments stored in the depot. We recommend that you use `// .swarm`

Create a directory for the comment attachments

We recommend that you use `// .swarm` because it is best practice to use a depot directory that is not used to store any other files. This prevents other users from having direct access to the comment attachments.

1. To create a `// .swarm` depot, run the following as a user with *admin*-level privileges:

```
$ p4 depot .swarm
```

2. Ensure that the Swarm *admin* user can write to the `// .swarm` depot and make sure that other users cannot access it.

For information on creating depots, see [Working with depots](#) in the *Helix Core Server Administrator Guide*.

Specify the location for the comment attachments

By default Swarm looks for a depot named `// .swarm`, to specify a different depot path for comment attachments, use the `depot_storage` configuration block in the `SWARM_ROOT/data/config.php` file.

Replace `depot_name` with the depot where comment attachments are stored. The Swarm *admin* needs to be able to write to this depot but make sure that other users cannot access it:

```
<?php
// this block should be a peer of 'p4'
'depot_storage' => array(
    'base_path'  => '//depot_name',
),
```

Tip

The `base_path` can be more than just the depot name, for example `//depot/perforce/.swarm` would work.

Maximum comment attachment size

By default, the maximum file size of a single comment attachment is limited to:

- **Ubuntu, RHEL 7, and CentOS:** 8Mb
- **RHEL 8 and later:** 2Mb

Increasing the maximum attachment file size

You can increase the maximum attachment size if required.

Ubuntu, CentOS, RHEL 7

Edit the `php_value upload_max_filesize` and `php_value post_max_size` configurables in the `.htaccess` file.

RHEL 8 and later

Edit the `upload_max_filesize` and `post_max_size` configurables in the `/etc/php.ini` file.

Comment mentions

This section describes how to customize how Swarm treats user `@mentions` and group `@@mentions` in auto-complete drop-downs, comments, changelists, and review descriptions.

Note

- `@mentions` for users is only available for Swarm 2017.1 and later.
- `@@mentiiions` for groups is only available for Swarm 2017.3 and later.
- If your Helix Core Server is configured to be case sensitive this will also apply to the names specified in your exclude lists.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

By default:

- When a user uses an `@mention` or `@@mention` in a comment, Swarm will auto-complete for all Swarm users or groups respectively.
- When a user edits the reviewers on a review, Swarm will auto-complete for all Swarm users and groups.
- When a user edits the author of a review, Swarm will auto-complete for all Swarm users.

This behavior can be fine tuned with the **mentions** configuration block in the [SWARM_ROOT/data/config.php](#) file. For example:

```
<?php
    'mentions' => array(
        'mode' => 'global',
        'user_exclude_list' => array('super', 'swarm-admin'),
        'group_exclude_list' => array('testers', 'writers'),
    ),
```

- **mode** controls the scope of the auto-complete dropdown that is shown when the user starts typing an `@mention` or an `@@mention` in a review comment:
 - **global** displays all users/groups in Swarm, this is the default value
 - **projects**: only displays project members (users and groups) in the auto-complete list
 - **disabled** disables the auto-complete drop down so that it is not displayed
- **user_exclude_list** users in the exclude list:
 - are not suggested when you `@mention` them in a comment, edit reviewers, or edit the author in Swarm.

- are not added to the review if you manually add them as an *@mention* in a Swarm comment or a review description.
- are not added to the review if they are *@mentioned* in a changelist description.
- cannot be added to the review with the Swarm API.

Tip
When you add a user to the exclude list:

- the user is not removed from any reviews that they are already on.
- the user will stop receiving email notifications for those reviews.

■ **group_exclude_list** groups in the exclude list:

- are not suggested when you *@@mention* them in a comment or edit reviewers in Swarm.
- are not added to the review if you manually add them as an *@@mention* in a Swarm comment or review description.
- are not added to the review if they are *@@mentioned* in a changelist description.
- cannot be added to the review with the Swarm API.

Tip
When you add a group to the exclude list:

- the group is not removed from any reviews that it is already on.
- the group members will stop receiving email notifications for those reviews.
- if a member of an excluded group is also on the review as a user or as a member of another group, that user will continue to receive emails for the review.

Regular expressions in exclude lists

Regular expressions can be used in user and group exclude lists but they should be used with care to avoid accidentally excluding the wrong users or groups. Swarm exclude lists can use any valid PHP regular expression but not all of them make sense for the exclude lists.

Note

Swarm adds a `^` character before your regular expression and a `$` character after it. For example: `^your-reg-ex-pattern$` You must take this into account when constructing your regular expressions.

Tip

- You can test the results of your regular expressions before using them, see <https://regex101.com/>
- Case sensitivity for regular expressions in the exclude lists is determined by the case sensitivity setting of the Helix Core Server.

Regular expression examples

Useful regular expressions are shown below:

Exact match only:

'**super**' only users or groups named *super* are excluded.

Beginning with:

'**admin.***' users or groups beginning with *admin* are excluded.

For example: the following would be excluded:

administrator, admin-group, and admin

Ending with:

'**.*admin**' users or groups ending with *admin* are excluded.

For example: the following would be excluded:

test-admin, buildadmin, and admin

Ending with digits

'**.*[0-9]+**' users or groups ending with digits are excluded.

For example: the following would be excluded:

test03, admin-10, tester_123456

Containing:

'**.*admin.***' users or groups that contain *admin* are excluded.

For example: the following would be excluded:

test-admin, buildadmin, administrator, admin-group, and admin

Comments

This section provides information on how to delay comment notifications, and configure comment threading.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

Comment notification delay

By default, comment notifications are delayed to allow reviewers to add or edit comments as they progress through a review without sending a notification for each individual comment on the review. Comment notifications are rolled up into a single notification and sent either manually by the reviewer, or automatically after the notification delay time has been exceeded.

The delay countdown is reset each time the reviewer adds or edits a comment on the review, by default the notification delay time is set to 30 minutes.

- If you are commenting on more than one review, each of the reviews that you are commenting on has its own notification delay countdown that only applies to the comments that *you* make on *that* review.
- If another reviewer is making comments on the same review as you, *that* reviewer has their own notification delay timer for *that* review.
- If you manually send a delayed comment notification, the notification will only contain the comments that *you* made on *that* review.

Tip

The comment notification delay does not delay the posting of the comments, only the comment notification is delayed.

Note

Comment notifications are only delayed for comments on reviews. Comments on commits or jobs produce notifications immediately.

To change the comment notification delay time, update the `SWARM_ROOT/data/config.php` file to include the following configuration item within the comments block:

```
<?php
    'comments' => array(
        'notification_delay_time' => 1800, //Default to 30 minutes
```

```
1800 seconds
),
```

notification_delay_time: Specifies the comment notification delay time in seconds.

- Set to **0** to send the comment notification immediately the **Post** button is clicked.
- The default value if **notification_delay_time** is not specified is **1800** seconds (30 minutes).

Comment threading

By default, you can reply to comments, replies are displayed in a thread below the parent comment. Comment thread depth is set to 4 by default, this means you can have up to 4 levels of replies for a parent comment. The **Reply** link is not displayed for replies at or above the maximum thread depth set for Swarm. If the parent comment is archived, replies are archived with the parent, see "[Archiving comments](#)" on page 362. Comment threading depth is set by the **max_depth** configurable.

Tip

If the thread depth is reduced by a Swarm administrator, earlier replies at a deeper level will continue to be displayed but you cannot reply to them.

To configure comment threading, update the `SWARM_ROOT/data/config.php` file to include the following configuration item within the comments block:

```
<?php
    'comments' => array(
        'threading' => array(
            'max_depth' => 4, // default depth 4, to disable comment
threading set to 0
        ),
    ),
```

max_depth: Specifies maximum depth of a comment thread.

- Set to **0** to disable comment threading.
- The default value if **max_depth** is not specified is **4**.

Commit credit

When you use Swarm to commit a review, but you are not the review's author, Swarm gives credit to the review author by default. Activity stream entries and email notifications include both the committer and review author's details.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

If you prefer Swarm's original behavior, which was to give credit only to the committer, you can do so by editing the `$SWARM_ROOT/data/config.php` file and setting the `commit_credit_author` item to `false` in the `reviews` configuration block. For example:

```
<?php
// this block should be a peer of 'p4'
'reviews' => array(
    'commit_credit_author' => false,
),
```

Commit-edge deployment

Swarm can connect to a Helix Core Server configured to use the *commit-edge architecture*, which is a specific replication configuration that employs a *commit* server and one or more *edge* servers. This configuration distributes the compute, storage, and network requirements for improved performance and geographic distribution.

For more information on Helix Core Server's commit-edge architecture, see the [Commit-edge](#) chapter in the *Helix Core Server Administrator Guide*.

Tip

P4V users may have problems when P4V connects to an edge server and your commit and edge tickets are different. For more information on P4V authentication for commit-edge deployment, see ["P4V Authentication" on the next page](#).

Swarm commit-edge configuration

This section describes what configuration changes need to be made to Swarm and the Helix Core Servers to run Swarm with a commit-edge deployment. Before you begin making changes ensure that your Helix Core Server commit-edge deployment is working correctly, and that Swarm has been installed and configured.

Configure Swarm for a commit-edge deployment:

1. Configure Swarm to connect to the Helix Core Commit server, see ["Swarm configuration" on page 193](#). This is called the **Commit Server Swarm Instance**.

When Swarm is connected to the **Commit Server Swarm Instance**, the first worker detects this situation and sets a key in the Helix Core Server, `P4.Swarm.CommitURL`, to an auto-detected URL.

2. Configure the Swarm triggers:

- a. Configure the Swarm triggers on all of the Commit and Edge servers, see ["Installing Triggers" on page 204](#).

Important

- Save all of the Swarm trigger scripts in exactly the same place on each of the Commit and Edge servers. This is important because they all share the same trigger table.
- You must install the trigger dependencies on all of the Commit and Edge servers, see ["Trigger dependencies" on page 108](#).

- b. Configure `SWARM_HOST` in the `swarm-trigger.conf` file of all of the Commit and Edge servers to point to the **Commit Server Swarm Instance**, see ["Installing Triggers" on page 204](#).
3. Configure all of the Commit and Edge servers to point to the **Commit Server Swarm Instance** URL. This is done by setting the `P4.Swarm.URL` to the **Commit Server Swarm Instance** URL on each of the Commit and Edge servers, see [Swarm integration properties](#) in the *Helix Core Server Administrator Guide*.
4. Your Swarm configuration is now complete but you must check that it is working correctly before using it in production, see ["Validate your Swarm installation" on page 239](#).

P4V Authentication

When using P4V's Swarm integration in a commit-edge deployment, users may encounter authentication errors; such errors can result from incorrect configuration of login tickets in distributed environments. Essentially, the problem is that while P4V is connected to an edge server, Swarm is connected to the commit server, and the login tickets do not match.

If P4V users see the **error Host requires authentication**, the solution we recommend is to forward login requests to the commit server. This can be achieved by executing the following commands as a user with *operator* or *super* privileges in the Helix Core Server.

For example, the following `p4 configure set` commands run against the commit server set `auth.id` globally and enable login forwarding from the `edge1` and `replica1` servers:

```
p4 configure set auth.id=myAuthName
p4 configure set edge1#rpl.forward.login=1
p4 configure set replica1#rpl.forward.login=1
```

Replace `myAuthName` with the authentication identifier for your Helix Core Server.

For more information, see our Knowledge Base article [Single Ticket Login in Distributed Environments](#), and the `p4 serverid` command in the [Helix Core Command-Line \(P4\) Reference](#).

Commit timeout

When a code review contains many files, or large files, or both, committing the review within Swarm can take some time. The default configuration, within the `SWARM_ROOT/data/config.php` file:

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

```
<?php
    // this block should be a peer of 'p4'
    'reviews' => array(
        'commit_timeout' => 1800, // 30 minutes
    ),
```

The `commit_timeout` key is expressed in seconds. If a commit operation takes longer than this limit, it is terminated. It is likely that a terminated commit requires administrator intervention to complete the commit using another client.

Configuration overview

This section provides an overview of all the possible configuration blocks in the `SWARM_ROOT/data/config.php` file. Click on any link in the configuration overview below to see a detailed description of that configurable.

Warning

While the syntax of this example is correct, it includes configuration values that **cannot work**. Ensure that you adjust the configuration appropriately for your Swarm installation before using this example in testing or production.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

```
<?php
    return array(
```

```
'activity' => array(
    'ignored_users' => array(
        'p4dtguser',
        'system',
    ),
),
'archives' => array(
    'max_input_size'    => 512 * 1024 * 1024, // 512M (in
bytes)
    'archive_timeout'   => 1800,             // 30 minutes
    'compression_level' => 1,                // 0-9
    'cache_lifetime'    => 60 * 60 * 24,     // 1 day
),
'avatars' => array(
    'http_url'  => 'http://www.gravatar.com/avatar/{hash}?s=
{size}&d={default}',
    'https_url' => 'https://secure.gravatar.com/avatar/
{hash}?s={size}&d={default}',
),
'comments' => array(
    'notification_delay_time'    => 1800, //Default to 30
minutes 1800 seconds
    'threading'    => array(
        'max_depth' => 4, // default depth 4, to disable
comment threading set to 0
    ),
),
'depot_storage' => array(
    'base_path'  => '//depot_name',
),
'diffs' => array(
    'max_diffs'    => 1500,
),
'environment' => array(
    'mode'         => 'production',
    'hostname'     => 'myswarm.hostname',
```

```

        'external_url' => null,
        'base_url'      => null,
        'logout_url'    => null, // defaults to null
        'vendor'        => array(
            'emoji_path' => 'vendor/gemoji/images',
        ),
    ),
    'files' => array(
        'max_size'          => 1048576,
        'download_timeout' => 1800,
        'allow_edits' => true, // default is true
    ),
    'groups' => array(
        'super_only' => true, // default value is false
    ),
    'http client options' => array(
        'timeout'          => 10, // default value is 10 seconds
        'sslcapath'        => '', // path to the SSL certificate
directory
        'sslcert'          => '', // the path to a PEM-encoded SSL
certificate
        'sslpassphrase' => '', // the passphrase for the SSL
certificate file
        'hosts'           => array(), // optional, per-host
overrides. Host as key, array options as values
    ),
    'jira' => array(
        'host'             => '', // URL for your installed Jira
web interface (start with https:// or http://)
        'api_host'         => '', // URL for Jira API access,
'host' is used for Jira API access if 'api_host' is not set
        'user'             => '', // Jira Cloud: the username used
to connect to your Atlassian account
                                // Jira on-premises: the username
required for Jira API access

```

```

        'password'          => '', // Jira Cloud: a special API
token, obtained from https://id.atlassian.com/manage/api-tokens
                                // Jira on-premises: the password
required for Jira API access
        'job_field'        => '', // optional, if P4DTG is
replicating Jira issue IDs to a job field, list that field here
        'link_to_jobs'     => false, // set to true to enable
Perforce job links in Jira, P4DTG, and job_field required
        'delay_job_links' => 60, // delay in seconds, defaults to
60 seconds
        'relationship'    => '', // Jira subsection name links are
added to defaults to empty, links added to the "links to" subsection
    ),
    'linkify' => array(
        'word_length_limit' => 2048, // limit on the number of
characters which a text to be linkified can have
        'target' => '_self',          // opens the URL in the same
tab or a new tab, defaults to '_self'. To open the URL in a new tab,
set to '_blank'
        'markdown' => array(
            array(
                'id'      => 'jobs',
                'regex' => '',          // the regular expression
used to match the job keyword, default is empty
                'url'    => '',          // url that matching job
numbers are appended to, default is empty
            ),
        ),
    ),
    'log' => array(
        'priority'      => 3, // 7 for max, defaults to 3
        'reference_id' => false // defaults to false
    ),
    'mail' => array(
        // 'recipients' => array('user@my.domain'),
        'notify_self'  => false,

```

```

        'transport' => array(
            'host' => 'my.mx.host',
        ),
    ),
    'markdown' => array(
        'markdown' => 'safe', // default is 'safe'
        'disabled'|'safe'|'unsafe'
    ),
    'mentions' => array(
        'mode' => 'global',
        'user_exclude_list' => array('super', 'swarm-admin'),
        'group_exclude_list' => array('testers', 'writers'), //
defaults to empty
    ),
    'menu helpers' => array(
        'MyMenu01' => array( // A short recognizable name for the
menu item
            'id' => 'custom01', // A unique id
for the menu item. If not included in the array, parent array name is
used.
            'enabled' => true, // When set to
true, the menu item is visible. Defaults to true if not included in
the array.
            'target' => '/module/MyMenuItem/', // The URL or
custom module route a menu click takes you to.
// If not
included in array, id is used. If id not included, parent array name
is used.
            'cssClass' => 'custom_menu', // The custom
CSS class name added to the menu item, appended to h2.menu- in Swarm
CSS
            'title' => 'MyMenuItem', // The text that
will be shown on the button.
// If not
included in array, id is used. If id not included, parent array name

```

```

is used.

        'class'      => '',                                // If not
included in array or empty, the menu item is added to the main menu.
                                                // To add the
menu item to the project menu for all of the projects, set to
'\Projects\Menu\Helper\ProjectContextMenuHelper'
        'priority'   => 155,                                // The position
the menu item is displayed at in the menu.
                                                // If not
included in the array, the menu item is placed at the bottom of the
menu.
        'roles'      => null,                                //
null|'authenticated'|'admin'|'super'
                                                // If not
included in the array, null is the default.
                                                // Specifies the
minimum level of Perforce user that can see the menu item.
                                                //
'authenticated' = any authorized user, null = unauthenticated users
    ),
),
'notifications' => array(
    'honor_p4_reviews'      => false,
    'opt_in_review_path'    => '///depot/swarm',
    'disable_change_emails' => false,
),
'p4' => array(
    'port'      => 'my-helix-core-server:1666',
    'user'      => 'admin_userid',
    'password'   => 'admin user ticket or password',
    'sso'       => 'disabled', //
'disabled'|'optional'|'enabled'
                                                // default value is 'disabled'
'proxy_mode' => true, // defaults to true
'slow_command_logging' => array(
    3,

```

```

        10 => array('print', 'shelve', 'submit', 'sync',
'unshelve'),
    ),
    'max_changelist_files' => 1000,
    'auto_register_url'    => true,
),
'projects' => array(
    'mainlines' => array(
        'stable', 'release', // 'main', 'mainline', 'master',
and 'trunk' are hardcoded, there is no need to add them to the array
    ),
    'add_admin_only'          => false,
    'add_groups_only'         => array(),
    'edit_name_admin_only'    => false,
    'edit_branches_admin_only' => false,
    'readme_mode'             => 'enabled',
    'fetch'                   => array('maximum' => 0), //
defaults to 0 (disabled)
    'allow_view_settings'     => false, // defaults to false
),
'queue' => array(
    'workers'                  => 3,      // defaults to 3
    'worker_lifetime'         => 595,    // defaults to 10 minutes
(less 5 seconds)
    'worker_task_timeout'     => 1800,    // defaults to 30 minutes
    'worker_memory_limit'     => '1G',    // defaults to 1 gigabyte
),
'redis' => array(
    'options' => array(
        'password' => null, // Defaults to null
        'namespace' => 'Swarm',
        'server' => array(
            'host' => 'localhost', // Defaults to 'localhost'
or enter your Redis server hostname
            'port' => '7379', // Defaults to '7379' or enter

```

```

your Redis server port
    ),
    ),
    'items_batch_size' => 100000,
    'check_integrity' => '03:00', // Defaults to '03:00' Use
one of the following options:
                                //'HH:ii' (24 hour format
with leading zeros), the time the integrity check starts each day
                                // positive integer, the
time between integrity checks in seconds. '0' = integrity check
disabled
    'population_lock_timeout' => 300, // Timeout for initial
cache population. Defaults to 300 seconds.
    ),
    'reviews' => array(
        'patterns' => array(
            'octothorpe' => array( // #review or #review-1234
with surrounding whitespace/eol
                'regex' => '/(?P<pre>(?:\s|^)\(?\)\#
(?P<keyword>review|append|replace) (?:- (?P<id>[0-9]+)) ?(?P<post>
[.,!?:;])* (?:\s|$))/i',
                'spec' => '%pre##keyword%-%id%post%',
                'insert' => "%description%\n\n#review-%id%",
                'strip' => '/^\s*\#(review|append|replace) (-[0-
9]+)?(\s+|$)|(\s+|^)\#(review|append|replace) (-[0-9]+)?\s*$/i',
            ),
            'leading-square' => array( // [review] or [review-
1234] at start
                'regex' => '/^(?P<pre>\s*)\
[ (?P<keyword>review|append|replace) (?:- (?P<id>[0-9]+)) ?\]
(?P<post>\s*)/i',
                'spec' => '%pre%[%keyword%-%id%]post%',
            ),
            'trailing-square' => array( // [review] or
[review-1234] at end
                'regex' => '/(?P<pre>\s*)\

```



```
[ (?P<keyword>review|append|replace) (?::-(?P<id>[0-9]+)) ?\]
(?P<post>\s*) ?$/i',
    'spec'    => '%pre%[%keyword%-%id%]%post%',
),
),
'filters' => array(
    'filter-max' => 15,
    'result_sorting' => true,
    'date_field' => 'updated', // 'created' displays and sorts
by created date, 'updated' displays and sorts by last updated
),
'cleanup' => array(
    'mode'      => 'user', // auto - follow default,
user - present checkbox(with default)
    'default'    => false, // clean up pending
changelists on commit
    'reopenFiles' => false, // re-open any opened files
into the default changelist
),
'statistics' => array(
    'complexity' => array(
        'calculation' => 'default', // 'default|disabled'
        'high' => 300,
        'low' => 30
    ),
),
'allow_author_change'      => true,
'allow_author_obliterate'  => false,
'commit_credit_author'     => true,
'commit_timeout'           => 1800, // 30 minutes
'disable_approve_when_tasks_open' => false,
'disable_commit'           => true,
'disable_self_approve'     => false,
'end_states'               => array('archived',
'rejected', 'approved:commit'),
```

```

        'expand_all_file_limit'           => 10,
        'expand_group_reviewers'         => false,
        'ignored_users'                   => array(),
        'max_secondary_navigation_items'  => 6, // defaults to 6
        'moderator_approval'              => 'any', // 'any|each'
        'more_context_lines'              => 10, // defaults to 10
lines
        'process_shelf_delete_when'       => array(),
        'sync_descriptions'                => true,
        'unapprove_modified'               => true,
    ),
    'search' => array(
        'maxlocktime'           => 5000, // 5 seconds, in milliseconds
        'p4_search_host'        => '',   // optional URL to Helix Search
Tool
    ),
    'security' => array(
        'disable_system_info'     => false,
        'email_restricted_changes' => false,
        'emulate_ip_protections'  => false, // defaults to false
        'https_port'               => null,
        'https_strict'              => false,
        'https_strict_redirect'    => true,
        'require_login'            => true,
        'prevent_login'            => array(
            'service_user1',
            'service_user2',
        ),
    ),
    'session' => array(
        'cookie_lifetime'          => 0, // lifetime in seconds,
default value is 0=expire when browser closed
        'remembered_cookie_lifetime' => 60*60*24*30, // lifetime
in seconds, default value is 30 days
        'user_login_status_cache'  => 10, // Set in seconds,

```

```

default value is 10 seconds.

                                                                    // Set to 0 to disable
the cache and make Swarm

                                                                    // check the user
login status for every call to Helix Server.
    'gc_maxlifetime'          => 60*60*24*30, // lifetime
in seconds, default value is 30 days
    'gc_divisor'              => 100, // 100 user requests
),
    'short links' => array(
        'hostname'           => 'myho.st',
        'external_url'       => 'https://myho.st:port/sub-folder',
    ),
    'slack' => array(
        'token'              => 'TOKEN',
        'project_channels'   => array(
                                'myproject' => array
('myproject-channel',), //The Swarm project name must be in lower case
letters.

                                //For project 'myproject' the slack notification

                                //will go into the Slack channel 'myproject-
channel'.
                                ),
        'summary_file_names'   => false, //Attaches the
file to the original notification message sent to a Slack channel.
        'bypass_restricted_changelist' => false, //Allows Swarm
to post notification messages to a Slack channel when a change is
committed

                                                                    //or a review is
created for a restricted changelist, default value is false.
        'summary_file_limit'   => 10, //Limits the number
of files shown in the original notification message sent to a Slack
channel, default value is 10.

```

```

        'user' => array(
            'enabled' => true, //Forces the
Swarm app to use the custom username, overrides the Swarm app details.
            'name' => 'Helix Swarm', //This
is the username shown in the Slack channel when a notification is
posted.
            'icon' => 'URL', //This is the
avatar icon shown in the Slack channel

                                //when a
notification is posted, overrides the avatar set in the Swarm app.
            'force_user_header' => false, //The Slack
notification shows the username and avatar only for the first post by
a user, default value is false.
        ),
    ),
    'tag_processor' => array(
        'tags' => array(
            'wip' => '/(^|\s)+#wip($|\s)+/i'
        ),
    ),
    'test_definitions' => array(
        'project_and_branch_separator' => ':',
    ),
    'translator' => array(
        'detect_locale' => true,
        'locale' => array("en_GB", "en_US"),
        'translation_file_patterns' => array(),
        'non_utf8_encodings' => array('sjis', 'euc-jp',
'windows-1252'),
        'utf8_convert' => true,
    ),
    'upgrade' => array(
        'status_refresh_interval' => 10,    //Refresh page every 10
seconds
        'batch_size' => 1000,              //Fetch 1000 reviews to lower memory

```

```

usage
    ),
    'users' => array(
        'dashboard_refresh_interval' => 300000, //Default 300000
milliseconds
        'display_fullname'            => true,
        'settings' => array(
            'review_preferences' => array(
                'show_comments_in_files'            => true,
                'view_diffs_side_by_side'            => true,
                'show_space_and_new_line_characters' => false,
                'ignore_whitespace'                  => false,
            ),
            'time_preferences' => array(
                'display' => 'Timeago', // Default to 'Timeago' but
can be set to 'Timestamp'
            ),
        ),
    ),
    'workflow' => array(
        'enabled' => true, // Switches the workflow feature on.
Default is true
    ),
    'xhprof' => array(
        'slow_time'            => 3,
        'ignored_routes' => array('download', 'imagick',
'libreoffice', 'worker'),
    ),
    'saml' => array(...), // This configurable will be removed
from Swarm in a later release.
);

```

Note

The Swarm configuration file does not include PHP's standard closing tag (`?>`). This is intentional as it prevents unintentional whitespace from being introduced into Swarm's output, which would interfere with Swarm's ability to control HTTP headers. Debugging problems that result from unintentional whitespace can be challenging, since the resulting behavior and error messages often appear to be misleading.

Diff configuration

Swarm's [Diffs](#) feature can be customized via the following config item.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button"](#) on page 732.

max_diffs

If the number of diffs is greater than the configured maximum, then the list of diffs is truncated by default. An option will be given to show all the diffs for that file.

```
'diffs' => array(
    'max_diffs' => 1500,
),
```

max_size_chunk

This is the maximum size allowed for a diff block in a file. If the size of a diff block is greater than 500 KB then the file is truncated by default. An option is provided to view all the diffs.

```
'diffs' => array(
    'max_size_chunk' => 1024 * 512, //0.5 MB is the maximum size
    of the diff block and the value must be specified in bytes
),
```

Email configuration

Swarm's email delivery is controlled by the mail configuration block in the [SWARM_ROOT/data/config.php](#) file.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button"](#) on page 732.

Example:

```
<?php
    // this block should be a peer of 'p4'
    'mail' => array(
        // 'sender' => 'swarm@my.domain',    // defaults to
'notifications@hostname'
        'transport' => array(
            'name' => 'localhost',            // name of SMTP host
            'host' => '127.0.0.1',            // host/IP of SMTP host
            'port' => 587,                    // SMTP host listening
port
            'connection_class' => 'plain', // 'smtp', 'plain',
'login', 'crammd5'
            'connection_config' => array(    // include when auth
required to send
                'username' => 'user',        // user on SMTP host
                'password' => 'pass',        // password for user on
SMTP host
                'ssl' => 'tls',              // empty, 'tls', or 'ssl'
            ),
            // override email deliveries and store all messages in
this path
            // for more information about available options when using
the path parameter, see "Transport" on page 605 section
            // 'path' => '/var/spool/swarm',
        ),

        // override regular recipients; send email only to these
addresses
        // 'recipients' => array(
        //     'user1@my.domain',
```

```
//      'user2@my.domain',  
// ),  
  
// send notifications of comments to comment authors?  
'notify_self' => false,  
  
// blind carbon-copy recipients  
// 'use_bcc' => true,  
  
// suppress reply-to header  
// 'use_replyto' => false,  
  
// change the email subject prefix, the default prefix is '  
[Swarm]'  
// 'subject_prefix' => '[Swarm]',  
  
// Control email thread indexing  
// 'index-conversations' => true, // ['true'|'false'] default  
is true, email thread indexing is turned on  
),
```

Note

Without any mail configuration in the `SWARM_ROOT/data/config.php` file, Swarm attempts to send email according to PHP's configuration, found in the `php.ini` file. By default, the configuration in `php.ini` relies on SendMail being installed.

Important

Email delivery for events related to restricted changes is disabled by default. See ["Restricted Changes" on page 703](#) for details on how to enable restricted change notifications.

Sender

The sender item within the `mail` block specifies the sender email address that should be used for all notification email messages. The default value is:

```
notifications@hostname
```

hostname is the name of the host running Swarm, or when specified with the ["Environment" on page 609](#) configuration.

Transport

The `transport` block within the `mail` block defines which mail server Swarm should use to send email notifications. Most of the items in this block can be omitted, or included as needed.

Important

When setting the path parameter, ensure that you do not use any other transport configuration options. See the Laminas Framework's [File Transport Options](#) and [SMTP Transport Options](#) for more information about the supported options.

Note

Swarm supports the Laminas component versions in the `LICENSE.txt` file, features and functions in the Laminas documentation that were introduced in later versions of Laminas will not work with Swarm. The `LICENSE.txt` file is in the `readme` folder of your Swarm installation.

Swarm uses the custom path item to direct all email messages to a directory instead of attempting delivery via SMTP. For details, see ["Save all messages to disk" on page 607](#).

Recipients

The `recipients` item within the `mail` block allows you to specify a list of recipients that should receive email notifications, overriding the normal recipients. This is useful if you need to debug mail deliveries.

```
<?php
// this block should be a peer of 'p4'
'mail' => array(
    'recipients' => array(
        'user1@my.domain',
        'user2@my.domain',
    ),
),
```

Any number of recipients can be defined. If the array is empty, email notifications are delivered to the original recipients.

notify_self

The `notify_self` item within the `mail` block specifies whether comment authors should receive an email for their comments. The default value is `false`. When set to `true`, comment authors receive an email notification for their own comments.

When a group mailing list is enabled it overrides the group member's preference set for the `notify_self` configurable in the `SWARM_ROOT/data/config.php` file. So even when the `notify_self` configurable is set to false the group member will receive an email notification.

There is a caveat that if a user is a group member and the same user is added individually to the review and has set the `notify_self` configurable to true, the user receives two email notifications.

Use BCC

The `use_bcc` item within the `mail` block allows you to address recipients using the Bcc email field. Setting the value to `true` causes Swarm to use the `Bcc:` field in notifications instead of the `To:` field, concealing the email addresses of all recipients.

```
<?php
    // this block should be a peer of 'p4'
    'mail' => array(
        'use_bcc' => true,
    ),
```

Use Reply-To

The `use_replyto` item within the `mail` block allows you to suppress populating the Reply-To email field. Setting the value to `false` causes Swarm to omit the `Reply-To:` field in notifications; by default, it is populated with the author's name and email address. When this field is `true`, users receiving an email notification can simply reply to the email and their response will be addressed to the author.

```
<?php
    // this block should be a peer of 'p4'
    'mail' => array(
        'use_replyto' => false,
    ),
```

Email subject prefix

The `subject_prefix` item within the `mail` block sets the prefix for the subject line of emails sent by Swarm. By default this is set to `[Swarm]`.

```
<?php
    // this block should be a peer of 'p4'
    'mail' => array(
        'subject_prefix' => '[Swarm]',
    ),
```

Save all messages to disk

For testing purposes, you may want to send all email to disk without attempting to send it to recipients. Use the following configuration block to accomplish this:

```
<?php
    // this block should be a peer of 'p4'
    'mail' => array(
        'transport' => array('path' => MAIL_PATH),
    ),
```

<MAIL_PATH> should be replaced with the absolute path where email messages should be written, the mail transport will create a new file for each email message and write it to the path directory. This path must already exist and be writable by the web server user.

Note

Use of the path item causes Swarm to ignore **all** other configuration within the transport block. This is why path is commented out in the main example.

Email headers

Swarm sends out notification emails that contain custom headers which email programs can use for automatic filtering. Emails also contain headers to ensure they are correctly threaded in email clients which support doing so.

All Swarm emails contain the following headers, which can be used to identify which Swarm server they came from:

```
X-Swarm-Host: swarm.perforce.com
X-Swarm-Version: SWARM/2019.1/1798019 (2019/05/09)
```

The exact values may differ according to the version of Swarm you are running, and its configuration.

If a notification is applicable to one or more projects, then each project will be listed in the **X-Swarm-Project** header, which contains a list of one or more project names. Reviews may span multiple projects, so in this case a single email is sent out with each project listed.

```
X-Swarm-Project: gemini, apollo
X-Swarm-Host: swarm.perforce.com
X-Swarm-Version: SWARM/2019.1/1798019 (2019/05/09)
```

If one or more of the applicable projects is private, then two or more emails may be sent. In order for the existence of the private project to be hidden from non-members, any email sent to them will not contain references to the private project. Members of each private project will receive an email tailored to them which contains references to that private project. The email to the non-private projects will not contain references to any of the private projects in the **X-Swarm-Project** header.

For example, if a review spans three projects, called **Gemini**, **Apollo**, and **Ultra**, where **Ultra** is a private project, then members of projects **Gemini** and **Apollo** will receive an email with the following headers:

```
X-Swarm-Project: gemini, apollo
X-Swarm-Host: swarm.perforce.com
X-Swarm-Version: SWARM/2019.1/1798019 (2019/05/09)
```

Members of the **Ultra** project will receive an email with the following header:

```
X-Swarm-Project: ultra
X-Swarm-Host: swarm.perforce.com
X-Swarm-Version: SWARM/2019.1/1798019 (2019/05/09)
```

This can result in users receiving two notification emails (if they are members of **Ultra** and one of the other two projects), but privacy for the private project is preserved.

Email thread indexing

By default, Swarm indexes emails so that email conversations are threaded correctly in your email clients. If you find that your email clients are not displaying Swarm email threads correctly, disable thread indexing. To disable indexing, use the following configuration block and set to **false**:

```
<?php
    // this block should be a peer of 'p4'
    'mail' => array(
        'index-conversations' => false, // ['true'|'false'] default is
    'true'
    ),
```

Emoji

By default, Swarm uses a font to provide Emoji images. Swarm can make use of Emoji images from the [Gemoji project](#). Gemoji provides support for more Emojis and works on more browsers and platforms than the font Swarm normally uses.

Note

By default, Gemoji is only supplied with a small number of custom emojis. For instructions on adding more emojis to Gemoji, see the documentation for the [Gemoji project](#).

Due to licensing issues, Gemoji cannot be distributed with Swarm. You can use the Gemoji images after following these steps:

1. Download the latest release (currently 3.0.0) of the [Gemoji project](#) from their [releases page](#).
2. Unpack the release archive into `SWARM_ROOT/public/vendor`.
After unpacking, you should see a new folder: `SWARM_ROOT/public/vendor/gemoji-3.0.0`.
`3.0.0` represents the version of Gemoji, which may differ if you downloaded a different or newer release.

Tip

By default, Gemoji is stored in the `SWARM_ROOT/public/vendor/gemoji` directory.

If the `"emoji_path"` on [page 613](#) configurable has been changed: unpack the Gemoji release archive into the directory specified in the `emoji_path` configurable. Modify the commands in the following steps to match that directory.

3. Rename the new folder, default location shown in the example below:

```
$ cd SWARM_ROOT/public/vendor
$ mv gemoji-3.0.0 gemoji
```

Replace `3.0.0` in the above command if you have downloaded a different or newer release of Gemoji.

4. Ensure that the new images are readable.

```
$ cd SWARM_ROOT/public/vendor
$ chmod -R +r gemoji
```

Swarm detects and uses Gemoji images automatically.

Environment

This section describes the *environment* configuration items available for Swarm:

- **mode**: whether Swarm operates in *production* or *development* mode.
- **hostname**: specifies the canonical hostname Swarm should use, such as in links to Swarm in email notifications.
- **external_url**: specifies the canonical URL Swarm should use, such as in links to Swarm in email notifications. Swarm can often auto-detect the correct URL, but use of `external_url` might be necessary in complex web hosting environments.
- **base_url**: specifies the folder name Swarm is installed within when Swarm is not installed in the web server's document root.
- **logout_url**: specifies the URL that users are redirected to when they logout of Swarm.
- **emoji_path**: specifies the location the Gemoji images are stored in.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button"](#) on page 732.

Add the following configuration block to the `SWARM_ROOT/data/config.php` file:

```
<?php
    // this block should be a peer of 'p4'
    'environment' => array(
        'mode'          => 'development',      // defaults to
'production'
        'hostname'      => 'myswarm.hostname', // defaults to requested
hostname
        'external_url' => null,                  // defaults to null
        'base_url'     => null,                  // defaults to null
        'logout_url'   => null,                  // defaults to null
        'vendor'       => array(
            'emoji_path' => 'vendor/gemoji/images',
        ),
    ),
```

mode

By default, Swarm operates in *production* mode. When **mode** is set to **development**, Swarm displays greater error detail in the browser. Also, Swarm switches from including aggregated and minified JavaScript and CSS to requesting each JavaScript and CSS resource for all active modules. The default value is production. Any value other than **development** is assumed to mean production.

development mode makes it easier to discover problems and to identify their source, but also incurs additional browser overhead due to many more JavaScript and CSS requests for larger files. We recommend that you do not use development mode in production environments, unless directed to do so by Perforce technical support.

hostname

The hostname item allows you to specify Swarm's hostname. This could be useful if you have multiple virtual hosts deployed for a single Swarm instance; Swarm uses the hostname you configure when generating its web pages and email notifications.

Note

The value specified for the `hostname` item should be just the hostname. It should not include a scheme (e.g. "`http://`"), nor should it include a port (e.g. "`:80`").

external_url

The `external_url` item allows you to specify Swarm's canonical URL. This is useful if your Swarm instance is proxied behind another web service, such as a load balancer, caching proxy, etc., because Swarm's auto-detection of the current hostname or port could otherwise result in incorrect self-referencing URLs.

When specified, Swarm uses the `external_url` item as the prefix for any URLs it creates that link to itself in its web pages and email notifications.

Note

- Any path components included in `external_url` are ignored. If you specify `https://myswarm.url:8080/a/b/c`, Swarm only uses `https://myswarm.url:8080/` when composing URLs.
- If you set the `external_url` for Swarm, disable the `auto_register_url` feature in Swarm. For information about disabling the `auto_register_url` feature for Swarm, see ["Client integration" on page 578](#).

Important**Multiple Swarm instances connected to a single Helix Core Server instance:**

If you specify `base_url` along with `external_url`, ensure that all Swarm instances specify the same `base_url`. Varying `base_url` amongst cooperating Swarm instances is not supported.

base_url

The `base_url` item allows you to specify Swarm's folder within the web server's document root. This is useful if you cannot configure Swarm to operate within its own virtual host, such as when you have an existing web service and Swarm must exist alongside other applications or content. For information on configuring Swarm to run within a sub-folder of an existing website, see ["Run Swarm in a sub-folder of an existing web site" on page 230](#).

Tip

The `swarm-trigger.conf` file must be updated to include the `base_url` in the `SWARM_HOST` path, see ["Swarm trigger configuration" on page 232](#)

By default, `base_url` is null, which is equivalent to specifying `/`. If you specify a folder, include the leading `/`. For example, `/swarm`.

Warning**Multiple Helix Core Server instances on a single Swarm instance:**

Issue: Swarm will lose connection to all of the Helix Core Servers if you edit the `base_url` configurable value in the `environment` block of `<swarm_root>/data/config.php`. This will stop your system working.

Fix: Remove the `base_url` configurable from the `environment` block of `<swarm_root>/data/config.php`.

Important**Multiple Swarm instances connected to a single Helix Core Server instance:**

If you specify `external_url` along with `base_url`, ensure that all Swarm instances specify the same `base_url`. Varying `base_url` amongst cooperating Swarm instances is not supported.

logout_url

The `logout_url` configurable is used to automatically redirect users to a custom URL after they have logged out of Swarm.

When the `logout_url` item is set and you log out of Swarm:

- Using **Log out** from Swarm: you are logged out by Swarm and then redirected to the custom URL.
- Using **Log out** from a single Swarm instance on the global dashboard: you are logged out of the Helix Core Server instance by Swarm. You are only redirected to the custom URL if you are not logged in to any other Helix Core Server instances on the global dashboard.
- Using **Log out from all instances** on the global dashboard: you are logged out of all of the Helix Core Server instances by Swarm and then redirected to the custom URL.

To redirect users to a custom URL after they have logged out of Swarm, edit the `SWARM_ROOT/data/config.php` file and set the `logout_url` configurable to the URL you want to redirect users to:

```
<?php
    // this block should be a peer of 'p4'
    'environment' => array(
        'logout_url'    => <http[s]://your/redirect/logout/url>, //
defaults to null
    ),
```

The default value is `null`, users are not redirected when they log out of Swarm.

emoji_path

The `emoji_path` configurable defines the location of the Gemoji images in `SWARM_ROOT/public/`.

To change the location of the Gemoji images, edit the `SWARM_ROOT/data/config.php` file and set the `emoji_path` configurable to the location the Gemoji images are stored in:

```
<?php
// this block should be a peer of 'p4'
'environment' => array(
    'vendor'    => array(
        'emoji_path' => 'vendor/gemoji/images',
    ),
),
```

The default location is `vendor/gemoji/images`, this means that by default images are stored in `SWARM_ROOT/public/vendor/gemoji/images`.

Further information

For instructions on downloading Gemoji and configuring Swarm to use the images, see [Emoji](#).

Excluding Users from Activity Streams

Larger Helix Core Server installations often have one or more *service* users that perform automated tasks, such as build systems, continuous integration test servers, or integrations with 3rd-party databases via P4DTG.

As Swarm reports the activity of users, and these service users can generate significant volumes of activity entries, Swarm provides a mechanism to ignore activity from specified users.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the **"Reload Configuration button"** on page 732.

Update the `SWARM_ROOT/data/config.php` file to include the following configuration block:

```
<?php
// this block should be a peer of 'p4'
'activity' => array(
    'ignored_users' => array(
        'p4dtguser',
```

```
        'system',
    ),
),
```

After **`SWARM_ROOT/data/config.php`** is updated, Swarm no longer records activity for any of the listed userids. Any previously recorded activity is included in activity streams.

Extensions Management

Note

You must be logged in as a user with *super* user privileges to configure the Helix Core Server Extensions.

This section details how to manage the Helix Core Server Extensions. You can use the **`swarm-extctl.sh`** script file located in **`/opt/perforce/swarm/sbin`** directory to perform different actions on the Helix Core Server Extensions.

You can use the following commands to manage the Helix Core Server Extensions:

- "List" below
- "Delete " on the facing page
- "Install" on page 616
- "Config" on page 617
- "Tokens" on page 620
- "Disable" on page 620
- "Enable" on page 620
- "Save" on page 621
- "Load" on page 621
- "Ping" on page 622

List

Description

Lists all the installed Extensions and their configurations.

Commands run

Runs the **`list`** command.

```
swarm-extctl list
```

The **swarm-extctl.sh** script responds with:

```
... config swarm
... extension Perforce::helix-swarm
... uuid 4532BC59-7BC8-478F-ADF6-0A563C42563D
... revision 1
... owner super
... type change-commit
... arg //...
... debugStatus none

... config swarm
... extension Perforce::helix-swarm
... uuid 4532BC59-7BC8-478F-ADF6-0A563C42563D
... revision 1
... owner super
... type change-content
... arg //...
... debugStatus none

... config swarm
... extension Perforce::helix-swarm
... uuid 4532BC59-7BC8-478F-ADF6-0A563C42563D
... revision 1
... owner super
... type change-submit
... arg //...
... debugStatus none
```

Delete

Description

Deletes the Swarm Helix Core Server extension and their configurations.

Commands run

Runs the **delete** command.

swarm-extctl delete

The **swarm-extctl.sh** script responds with:

```
Checking to see if you are a super user... super
Deleting the configuration
Are you really sure (y/n)?
y
Extension 'Perforce::helix-swarm and its configurations' successfully
deleted.
```

Install

Description

Installs the Swarm Helix Core Server extension. The **helix-swarm.p4-extension** file located in "[swarm_root](#)" on [page 724](#)/**p4-bin/extensions** directory is used to install Extensions.

For more information about installing Swarm Helix Core Server extension, see "[Installing the Swarm Helix Core Server extension \(recommended\)](#)" on [page 198](#).

After installation, the saved configurations are automatically applied to the extension. For more information about the saved configurations, see "[Save](#)" on [page 621](#) command.

Commands run

Runs the **install** command.

swarm-extctl install

The **swarm-extctl.sh** script responds with:

```
Checking to see if you are a super user... super
Installing the configuration
Are you really sure (y/n)?
y
Extension 'Perforce::helix-swarm' version '2023.1/2414875' installed
successfully.
Perform the following steps to turn on the Extension:

# Create a global configuration if one doesn't already exist.
p4 extension --configure Perforce::helix-swarm

# Create an instance configuration to enable the Extension.
```

```
p4 extension --configure Perforce::helix-swarm --name swarm
```

For more information, visit:

<https://www.perforce.com/manuals/v23.1/extensions/Content/Extensions/Home-extensions.html>

```
Reading configuration from /opt/perforce/etc/swarm-extension.conf
```

```
Extension config helix-swarm saved.
```

```
Reading configuration from /opt/perforce/etc/swarm-extension.conf.inst
```

```
Extension config swarm saved.
```

Config

Description

Lists the global configuration and instance configuration for the installed Extensions.

The global configuration is retrieved from `/opt/perforce/etc/swarm-extension.conf` file and the instance configuration is retrieved from `/opt/perforce/etc/swarm-extension.conf.inst` file. If no configuration files are available, the extension specification file opens in your text editor. You can edit the specification file and save it in your text editor.

Commands run

Runs the `config` command.

```
swarm-extctl config
```

The `swarm-extctl.sh` script responds with:

The global configuration and instance configuration opened in your text editor.

Here is a sample global configuration specification for Swarm Helix Core Server extension:

```
# A Perforce Extension Specification.
#
# ExtName:           The name of the Extension being configured.
# ExtDescription:    The description of the Extension being
configured.
```

```
# ExtVersion:      The version of the Extension being configured.
# ExtUUID:         The UUID/key of the Extension being configured.
# ExtRev:          The revision of the Extension being configured.
# ExtMaxScriptTime: Maximum seconds the Extension may be run.
# ExtMaxScriptMem: Maximum megabytes the Extension may use.
# ExtAllowedGroups: Groups whose members may configure the Extension.
# ExtEnabled:      Enable/Disable this Extension.
# ExtP4USER:       Perforce user account for the Extension to use.
# Name:            The name of this Extension config.
# Owner:           The user who created this Extension config.
# Update:          Update time for the Extension config spec.
# Description:     The description of this Extension config.
# ExtConfig:       Extension-supplied configuration fields.
#
# See 'p4 help extension' for detailed information.

ExtName:          helix-swarm

ExtDescription:
    Helix Swarm Extension

ExtVersion:       2023.1/2414875

ExtUUID:          4532BC59-7BC8-478F-ADF6-0A563C42563D

ExtRev: 1

ExtMaxScriptTime:      unset

ExtMaxScriptMem:       unset

ExtAllowedGroups:

"/tmp/tmp.28273.140308554466304.89" 63 lines, 1847 characters
```

```
ExtEnabled:      true

ExtP4USER:       super

Name:    helix-swarm

Owner:    super

Update: 2023/03/07 15:35:18

Description:
    This is the Helix Swarm extension, which replaces the old
Perl
    trigger that was used previously. If this is installed and
    configured, then the Swarm triggers MUST be removed.
    This file configures the connection to the Swarm server.
Only one
    Swarm instance can manage a single Helix server.

ExtConfig:
    Swarm-URL:      http://localhost/

    Swarm-Token:     785D3DD6-304D-5FCA-AF47-751578849960

    Swarm-Secure:    true

    Debug:    3
```

E163: There is only one file to edit

After you edit the specification file and save it in your text editor.

```
Extension config helix-swarm saved.
Extension config swarm saved.
```

Tokens

Description

Lists the trigger tokens available for a Helix Core Server including the multiple Helix Core Server (P4D) instances. The tokens are saved in `/opt/perforce/swarm/data/queue/tokens` directory. The Swarm installation folder is `/opt/perforce/swarm/`.

Commands run

Runs the `tokens` command.

```
swarm-extctl tokens
```

The `swarm-extctl.sh` script responds with:

```
Checking to see if you are a super user... super
Available tokens:
785D3DD6-304D-5FCA-AF47-751578849960
```

Disable

Description

Disables the Swarm Helix Core Server extension configuration.

Commands run

Runs the `disable` command.

```
swarm-extctl disable
```

The `swarm-extctl.sh` script responds with:

```
Checking to see if you are a super user... super
Disabled
```

Enable

Description

Enables the Swarm Helix Core Server extension configuration.

Commands run

Runs the **enable** command.

```
swarm-extctl enable
```

The **swarm-extctl.sh** script responds with:

```
Checking to see if you are a super user... super
Enabled
```

Save

Description

Saves the global configuration to `/opt/perforce/etc/swarm-extension.conf` file and the instance configuration to `/opt/perforce/etc/swarm-extension.conf.inst` file.

Commands run

Runs the **save** command.

```
swarm-extctl save
```

The **swarm-extctl.sh** script responds with:

```
Checking to see if you are a super user... super
Extension config helix-swarm saved.
Extension config swarm saved.
```

Load

Description

Loads the global configuration from `/opt/perforce/etc/swarm-extension.conf` file and the instance configuration from `/opt/perforce/etc/swarm-extension.conf.inst` file.

Commands run

Runs the **Swarm load** command.

```
swarm-extctl load
```

The **swarm-extctl.sh** script responds with:

```
Checking to see if you are a super user... super
Extension config helix-swarm saved.
Extension config swarm saved.
```

Ping

Description

Lists the version and status of the Swarm Helix Core Server extension.

Commands run

Runs the **ping** command.

```
swarm-extctl ping
```

The **swarm-extctl.sh** script responds with:

```
Checking to see if you are a super user... super
OK
Swarm Version: SWARM/2023.1/2414875 (2023/03/07)
Extension Version: 2023.1/2414875
```

Files configuration

Swarm can be customized using the following config items:

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button"](#) on page 732.

max_size

Swarm limits the size of files that are displayed in a review or standard file view. This defaults to 1MB, but can be configured to be larger or smaller. When the files are in the add or delete state, files larger than this will be truncated. The lowest **max_size** value defaults to 50 bytes. If you set the **max_size** value lower than 50 bytes then Swarm will ignore it.

When a file is in the edit state, depending on the **max_size** value, the following two flags are set in the response:

- **fetchContent**: This flag determines if the size of a file is larger than the **max_size** value and accordingly enables or disables the show more context buttons in the file diff panel. For more information about the show more context buttons, see ["File diff panel" on page 459](#).

Files larger than the **max_size** value sets the **fetchContent** flag to false. When the **fetchContent** flag is set to false, all the show more context buttons in the file diff panel are disabled. When the **fetchContent** flag is set to true, all the show more context buttons in the file diff panel are enabled.

- **isCut**: This flag determines if a file diff size is larger than the **max_size** value and must be truncated.

File diff sizes larger than the **max_size** value sets the **isCut** flag value to the **max_size** value which truncates the diff. When the file diff size is smaller than the **max_size** value, the **isCut** flag is set to false.

If the **max_size** value is set to zero, then the file size is unlimited.

```
<?php
// this block should be a peer of 'p4'
'files' => array(
    'max_size' => 1 * 1024 * 1024, // 1MB (in bytes)
),
```

download_timeout

When downloading files, there is a timeout which defaults to 30 minutes. This can be changed by setting the **download_timeout** configurable to a value in seconds. Alternatively, setting it to zero removes the timeout.

```
<?php
// this block should be a peer of 'p4'
'files' => array(
    'download_timeout' => 1800, // seconds (30 minutes)
),
```

allow_edits

Important

File edit is a Technology Preview feature.

Features offered in Technology Preview are experimental and not guaranteed to always work as expected. If you have feedback and functionality suggestions, email techpreview@perforce.com.

By default, you can edit a file from the files page and shelve or commit it.

When set to false, files cannot be edited from the Swarm files page.

```
<?php
// this block should be a peer of 'p4'
'files' => array(
    'allow_edits' => false, // default is true
),
```

Groups configuration

By default, Swarm allows all users to see the **Groups** menu item and use it to view the list of user groups. Users can view the members of a group, as well as their activities.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button"](#) on page 732.

You configure the visibility of this menu item with the **groups** configuration block in the `SWARM_ROOT/data/config.php` file, as in the following example:

```
<?php
// this block should be a peer of 'p4'
'groups' => array(
    'super_only' => true, // default value is false
),
```

Setting the **super_only** option to **true** hides the groups tab from non-super users.

Tip

If you need more fine grained control over the visibility of the **Groups** menu item, create a **groups** block in the `menu_helpers` configuration block. This enables you to restrict the visibility to only authenticated users, only users with admin privileges, or only users with super privileges. It is important to note that visibility of the **Groups** menu is controlled by a combination of the **super_only** setting and the **role** setting for groups in the `menu_helper` configuration block. Swarm uses the most restrictive of these settings to control the visibility of the **Groups** menu item.

Ignored users for reviews

Automated test environments may inadvertently participate in code reviews if they copy user-generated change descriptions. For example, if an automated system copied a change description containing `#review` and subsequently shelved or committed files, a new review would be started. Similarly copying a description with `#review-123` could inadvertently update an existing review. As test environments may involve thousands or millions of tests, such interactions can potentially generate far too many Swarm notifications.

To mitigate this problem, Swarm can be configured to ignore specified users for the purposes of starting or updating a review. Edit the `SWARM_ROOT/data/config.php` file, and provide the list of users to ignore in the `ignored_users` item in the reviews configuration block.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

Edit the `SWARM_ROOT/data/config.php` file, and provide the list of users to ignore in the `ignored_users` item in the reviews configuration block. For example:

```
<?php
// this block should be a peer of 'p4'
'reviews' => array(
    'ignored_users' => array('build_user1', 'build_user2'),
),
```

Job keywords

By default, if you add `jobnnnnnn` or `@jobnnnnnn` to a job description, changelist description, review description, or comment, Swarm automatically creates a link to the Swarm job page if a match is found. The `jobnnnnnn` or `@jobnnnnnn` keywords are hardcoded in Swarm and always present.

If your job numbers are not formatted as `jobnnnnnn`, you can configure Swarm to recognize other keywords in addition to `job` and `@job`. For example, if your job numbers are formatted as `defectnnnnnn`, Swarm can be configured to recognize them and link to the correct job page.

When you create a new job keyword, this is in addition to the default job keywords. Multiple job keywords can be specified. The first successful match is used, starting with the default job keywords. Once Swarm finds a match, none of the other keywords are evaluated.

The job keyword is configured with a regular expression so that almost any keyword can be used. When you add a new job keyword to Swarm, take care to choose syntax and terminology that is unlikely to occur naturally in a description to avoid unexpected links being created.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the "Reload Configuration button" on page 732.

Adding a new job keyword

Tip

- You can test the results of your regular expression before using it, see <https://regex101.com/>
- Regular expressions are case insensitive.

Job parameters:

- **'word_length_limit'**: The limit on the number of characters which a text to be linkified can have.
- **'target'**: Opens the URL in the same tab or a new tab, defaults to **'_self'**. To open the URL in a new tab, set to **'_blank'**.
- **'regex'**: The regular expression used to match the job keyword.
- **'url'**: The URL to append the regex match to. This URL plus the job number matched by the regex is used to display the job page. This URL can reference a [Swarm page](#) or an [external defect tracking system](#).

Optional: Specify the regex capture group to append to the URL. For example, adding **`/ {2}`** after the **url** parameter tells Swarm to append the second regex capture group result to the **url**. For more information on capture groups, see <https://www.regular-expressions.info/brackets.html>.

Swarm page link example

Tip

In this example, only the second regex capture group is appended to the **url** parameter. This means that the **@** character is only used to help identify the job and is not appended to the URL.

For example, to add "defect" as a jobs keyword so that both **@defectnnnnn** and **defectnnnnn** mentions are linked to a job page in Swarm, add the following block to the [SWARM_ROOT/data/config.php](#) file:

```
<?php
// this block should be a peer of 'p4'
'linkify' => array(
    'word_length_limit' => 2048,
```

```

        'target' => '_blank',
        'markdown' => array(
            array(
                'id'      => 'jobs',
                'regex'   => '(@?)(defect[0-9]+)', // the regular
expression used to match the job keyword, default is empty
                'url'     => '{baseurl}/{2}',      // url that matching
job numbers are appended to, default is empty
            ),
        ),
    ),
),

```

External page link example

Tip

In this example, the `/ {0}` character string at the end of the `url` is important because it helps the external link to get built dynamically.

For example, to add "issue" as a jobs keyword so that mentions of **issuennnnn** are linked to jobs in your defect tracking system, add the following block to the [SWARM_ROOT/data/config.php](#) file:

```

<?php
    // this block should be a peer of 'p4'
    'linkify' => array(
        'word_length_limit' => 2048,
        'target' => '_blank',
        'markdown' => array(
            array(
                'id'      => 'jobs',
                'regex'   => '(issue[0-9]+)',          // the
regular expression used to match the job keyword, default is empty
                'url'     => 'example.defect-tracker-url/{0}', // url
that matching job numbers are appended to, default is empty
            ),
        ),
    ),
),

```

Jira issue link example

Tip

In this example, the `{0}` character string at the end of the `url` is important because it helps the external link to get built dynamically.

For example, to add your Jira issue format so that mentions of `xxx-nnnnn` are linked to Jira issues in your Jira system, add the following block to the `SWARM_ROOT/data/config.php` file:

```
<?php
// this block should be a peer of 'p4'
'linkify' => array(
    'word_length_limit' => 2048,
    'target' => '_blank',
    'markdown' => array(
        array(
            'id' => 'jira',
            'regex' => '[A-Z1-9]{2,}-\d+', // the
regular expression used to match the job keyword, default is empty
            'url' => 'example.jira-system-url/browse/{0}', //
url that matching numbers are appended to, default is empty
        ),
    ),
),
```

License

Helix Swarm has always been free to use with an unlicensed Helix Core Server. With the 2014.4 release, Swarm is free to use with any Helix Core Server. You no longer need to purchase a Swarm-specific license, and any existing Swarm-specific licenses are no longer evaluated.

An unlicensed Helix Core Server provides unlimited use for up to 5 users and 20 workspaces. When the user or workspace limit is crossed, Helix Core Server imposes a maximum of 1,000 files. Swarm honors the restrictions of Helix Core Server.

Localization & UTF8 character display

Important

To ensure that all characters, including Unicode characters, are displayed and handled correctly by Swarm, configure your Helix Core Server in Unicode mode. For information on configuring your Helix Core Server in Unicode mode, see [Setting up and managing Unicode installations](#) in the *Helix Core Server Administrator Guide*.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the "Reload Configuration button" on page 732.

Swarm supports localization and the display of some non-UTF8 characters.

Localization

Important

If you are using a web browser that uses the **en-US** locale, Helix Swarm will automatically redirect to **en** to get locale files.

Swarm is fully localized; with an appropriate language pack installation, Swarm can support users in multiple languages.

A language pack consists of **gettext-style default.po** and **default.mo** files, placed in a folder named for the locale they represent, within the language folder in the Swarm root directory. In addition, language packs contain two JavaScript files to provide translation strings for the in-browser UI, **locale.js** and **locale.jsgz**, which both appear in the [SWARM_ROOT/public/build/language](#) folder.

The following example illustrates the directory layout of a language pack:

```
SWARM_ROOT/
  language/
    locale/
      default.mo
      default.po
  public/
    build/
      language/
        locale.js
        locale.jsgz
```

You can configure certain localization behaviors with the **translator** configuration block in the [SWARM_ROOT/data/config.php](#) file. Here is an example:

```
<?php
// this block should be a peer of 'p4'
'translator' => array(
    'detect_locale'           => true,
    'locale'                  => "",
    'translation_file_patterns' => array(),
),
```

The **detect_locale** key determines whether Swarm attempts to detect the browser's locale. The default value is **true**. Set the value to **false** to disable browser locale detection.

The **locale** key is a string specifying the default locale for Swarm. Alternately, an array of 2 strings can be used to specify the default locale, and a fallback locale. For example:

```
<?php
// this block should be a peer of 'p4'
'translator' => array(
    'locale'                  => array("en_GB", "en_US"),
),
```

The **translation_file_patterns** key allows you to customize the Laminas translation infrastructure, which you might do if you are developing your own language pack. For details, see [Laminas\l18n](#).

Note

Swarm supports the Laminas component versions in the **LICENSE.txt** file, features and functions in the Laminas documentation that were introduced in later versions of Laminas will not work with Swarm. The **LICENSE.txt** file is in the **readme** folder of your Swarm installation.

Non-UTF8 character display

Swarm supports the display of non-UTF8 characters in file content when all of the following are true:

- The Helix Core Server is Unicode enabled.
- The character encoding is supported by Swarm.
- The character encoding is listed in the **translator** array.

Multiple character encodings can be specified in the array, if Swarm fails to find a matching encoding in the array it will fall back to **windows-1252**. Unsupported and invalid character encodings in the array are ignored.

Tip

For information on character encodings supported by Swarm, see <http://php.net/manual/en/mbstring.supported-encodings.php>

Configure character encodings by adding the following block to the *SWARM_ROOT/data/config.php* file. For example:

```
<?php
    // this block should be a peer of 'p4'
    'translator' => array(
        'non_utf8_encodings' => array('sjis', 'euc-jp', 'windows-
1252'),
    ),
```

Important

windows-1252 is the default character encoding, it must always be the last entry in the array.

UTF8 conversion

If Swarm is connected to a non-unicode enabled Helix Core Server, Swarm converts special characters to UTF8 by default. For ANSI and ASCII files, this can result in some special characters being displayed in the diff view incorrectly or not at all. To stop Swarm converting special characters to UTF8 and to display special characters as unknown (), disable UTF8 conversion by setting **utf8_convert** to **false**.

Disable UTF8 character conversion by adding the following block to the *SWARM_ROOT/data/config.php* file. For example:

```
<?php
    // this block should be a peer of 'p4'
    'translator' => array(
        'utf8_convert' => false,
    ),
```

By default, the value of this configurable is **true**.

Logging

Helix Swarm is a web application, so there are several types of logging involved during the course of Swarm's normal operations.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the "Reload Configuration button" on page 732.

Web server logging

All accesses to Swarm may be logged by the web server hosting Swarm. As web server log configuration is web server specific, refer to your web server's documentation. Since we recommend the use of Apache, more information regarding log configuration in Apache can be found here: <http://httpd.apache.org/docs/2.4/logs.html>

Helix Core Server logs

Swarm communicates with the associated Helix Core Server for every page request. Review the Swarm-generated requests on your Helix Core Server by enabling logging.

For more information, see [Configuring logging](#) and [Auditing user file access](#) in the *Helix Core Server Administrator Guide*.

Swarm logs

Depending on the log configuration you provide to Swarm, Swarm can log its own operations. Swarm's logs are much more helpful if you encounter problems.

Swarm stores its log data in the `$SWARM_ROOT/data/log` file. Each request in the Swarm logs contains the following:

- URI (Uniform Resource Identifier)
- Process ID (PID) from the system
- IP address of the system that created the request

The worker event requests contain additional details such as:

- Worker ID
- Task information

The volume of entries recorded in the log depends on the configuration stored in the `$SWARM_ROOT/data/config.php` file. Here is an example:

```
<?php
// this block should be a peer of 'p4'
'log' => array(
```

```
'priority' => 3, // 7 for max, defaults to 3
),
```

The `logpriority` specifies the importance of the messages to be logged. When `priority` is set to `0` (the lowest value), only the most important messages are logged. When `priority` is set to `7` (the highest value), all messages, including the least important messages, are logged. The default priority is `3`.

The different priority levels are:

Priority	Description
0	Emergency: a message about a system instability
1	Alert: a message about a required immediate action
2	Critical: a message about a critical condition
3	Error: a message about an error
4	Warning: a message about a warning condition
5	Notice: a message about a normal but significant condition
6	Info: an informational message
7	Debug: a debugging message

Note

Setting `priority` to a value lower than `0` does not result in reduced logging.

Logrotate

From Swarm 2021.1, Swarm package installations and upgrades install logrotate to manage your Swarm log rotation. The logrotate configuration supports Swarm installations with single and multiple Helix Core Server instances.

Logrotate is on by default and is configured to:

- Rotate your Swarm logs daily.
- Keep Swarm logs for 14 days.

The `helix-swarm` logrotate configuration file can be modified or disabled and is located in `/etc/logrotate.d`. Any changes you make will be preserved when you upgrade Swarm in the future.

Reference ID

When `reference_id` is set to `true`, every log message is appended with `referenceId:<id>` where `<id>` is a hash value based on the web request id that generated the log message. The `<id>` is appended to all of the log messages related to that web request for the life of the web request allowing you to follow the request in the log. This can help to diagnose which request was being processed when the log message was created. It is particularly useful when you have lots of requests updating the log because it helps you to see which request each log message relates to.

Example: when a Swarm web request starts a worker the `<id>` that is generated is attached to all of the log entries related to that worker for the lifetime of the worker. This enables you to see all of the events the worker processes in its lifetime. For this example the web request `reference Id` is `"066ae44103ced2ab05c28578a36b6854"`:

```
2019-02-27T09:07:21+00:00 INFO (6): Worker 1 startup.
{"referenceId":"066ae44103ced2ab05c28578a36b6854"}
2019-02-27T09:07:21+00:00 INFO (6): Worker 1 event: task.commit(3276)
{"referenceId":"066ae44103ced2ab05c28578a36b6854"}
2019-02-27T09:07:21+00:00 DEBUG (7): P4
(00000000557948260000000002e2ec483) start command: help
{"referenceId":"066ae44103ced2ab05c28578a36b6854"}
.
.
.
2019-02-27T09:17:16+00:00 INFO (6): Worker 1 shutdown.
{"referenceId":"066ae44103ced2ab05c28578a36b6854"}
```

To append the `referenceId:<id>` to Swarm log messages, edit the [SWARM_ROOT/data/config.php](#) file to include the `reference_id` configurable and set it to `true`:

```
<?php
    // this block should be a peer of 'p4'
    'log' => array(
        'reference_id' => true, // defaults to false
    ),
```

The default value is `false`.

Trigger Token Errors

If the trigger tokens are missing or invalid, the web server error log contains a suitable error:

```
queue/add attempted with invalid/missing token: token used
```

A token is *invalid* when it is not formed from the characters A through Z (in upper or lowercase), 0 through 9, or a dash (-).

A token is *missing* when a file using the token as its name does not exist in the [SWARM_ROOT/data/queue/tokens](#) directory.

Performance logging

Swarm logs warnings whenever commands issued to the Helix Core Server take longer than expected to complete. These warnings can be used to diagnose performance problems in the Helix Core Server.

Note

By default, warnings are not captured in the Swarm log. To capture warnings, the [log priority](#) must be set to 4 or higher.

The default configuration is:

```
<?php
// this block should be placed within the 'p4' block
'slow_command_logging' => array(
    3, // commands without a specific rule have a 3-second limit
    10 => array('print', 'shelve', 'submit', 'sync', 'unshelve'),
),
```

In this configuration block, the numeric key specifies the time threshold in seconds, and the value (if present) is an array of Helix Core Server commands that should use the threshold. Should a command be associated with multiple thresholds, the largest is used for that command.

In addition, Swarm automatically detects when the PHP extension **xhprof** is installed and collects profiling data for requests that take longer than expected. The profiling data could be helpful in diagnosing performance issues within Swarm itself, particularly when analyzed in combination with the slow command logging described above. When collected, profiling data is stored in the [SWARM_ROOT/data/xhprof](#) folder.

The default configuration is:

```
<?php
// this block should be a peer of 'p4'
'xhprof' => array(
    'slow_time' => 3, // the threshold in seconds
    'ignored_routes' => array('download', 'imagick',
'libreoffice', 'worker'),
),
```

`slow_time` specifies the threshold, in seconds, that should be used to determine when a Swarm request is slow. `ignored_routes` is an array that specifies a list of Laminas Framework `route` names that should not be profiled. For example, Swarm's `Files` module specifies that the download route should be ignored from profiling as downloads could take significantly longer than the default threshold.

Note

Depending on the performance of the server hosting Swarm, and particularly the performance of the associated Helix Core Server, you may want to monitor the `SWARM_ROOT/data/xhprof` folder for disk usage. Each request that exceeds the time threshold causes 200-600KB of data to be written.

Mainline branch identification

When viewing a project's files, the initial view is the list of the project's branches. The branches appear in alphabetical order, but the branch identified as the main branch appears first and in bold. By default, a project overview page is displayed when there is a README Markdown file in the project mainline, this can be disabled by your Swarm administrator, see ["Project readme" on page 672](#).

Swarm uses a list of names to identify which of a project branches should be considered as the main branch and which Markdown file is used for the project overview page. The `mainlines` array supports regular expressions, see ["Regular expressions" on the facing page](#). The default names are *main*, *mainline*, *master*, and *trunk*. For the steps Swarm goes through to find the mainline branch and the Markdown file for a project, see ["Project mainline and README check" on page 638](#) below.

Tip

The default *main*, *mainline*, *master*, and *trunk* names are hardcoded in Swarm, there is no need to add them to your array. Swarm will always try to find exact matches to them from your branch names.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

You can adjust the configuration to match your local branch naming convention . Here is an example of the configuration block:

```
<?php
// this block should be a peer of 'p4'
'projects' => array(
    'mainlines' => array(
        'stable', 'release', // 'main', 'mainline', 'master', and
```



```
'trunk' are hardcoded, there is no need to add them to the array
    ),
),
```

Regular expressions

The **mainlines** array supports any valid PHP regular expression, this can be useful if you have a large number of similar branch names across different projects that you are using as your main branches. Regular expressions must be used with care to avoid accidentally including a branch that is not a mainline branch.

Note

Swarm adds a `^` character before your regular expression and a `$` character after it. For example: `^your-reg-ex-pattern$` You must take this into account when constructing your regular expressions.

Tip

- You can test the results of your regular expressions before using them, see <https://regex101.com/>
- Regular expressions are case insensitive.

Regular expression examples

Useful regular expressions are shown below:

Exact match only:

'**mymainbranch**' branches named **mymainbranch** are treated as mainline branches.

Only match default **mainlines** array branch names:

' ' (an empty **mainlines** array) only branches that exactly match the default names of '**main**', '**mainline**', '**master**', or '**trunk**' are treated as mainline branches.

Beginning with:

'**stable.***' branch names beginning with **stable** are treated as mainline branches.

For example: the following are treated as mainline branches.

stablecode-01, *stable2019_2*, and *stable_branchxyz*

Ending with:

'**.*mastprj**' branch names ending with **mastprj** are treated as mainline branches.

For example: the following are treated as mainline branches.

productxy-mastprj, *assets_mastprj*, and *project07mastprj*

Ending with digits

'`.*[0-9]+`' branch names ending with digits are treated as mainline branches.

For example: the following are treated as mainline branches.

branch03, *product-10*, *stable22*, *main_123456*

Containing:

'`.*main.*`' branch names that contain **main** are treated as mainline branches.

For example: the following are treated as mainline branches.

swarm-main-branch, *branchmain_22*, *main05*, and *assetmain-brnch*

Project mainline and README check

Tip

Swarm checks the branches of a project in alphabetical order.

Swarm uses the following process to find the mainline branch and README Markdown file for a project:

1. Swarm searches the project branches in alphabetical order for a branch that matches a name in the **mainlines** array. The search includes the default names.
2. When Swarm finds a matching branch, that branch is used as the mainline when displaying the branches on the project page.
 - If the branch contains a README Markdown file, that file is used for the project overview page.
 - If the branch does not contain a README Markdown file, see step 3.
3. Swarm checks the remaining branches for a branch that matches a name in the **mainlines** array.
 - If Swarm finds a README file, that file is used for the project overview page and swarm stops checking the branches.
 - If a README file is not found, this step is repeated until a README file is found in a branch that matches a name in the **mainlines** array or all of the project branches have been checked.
 - If no README Markdown page is found, the project overview page is not displayed.

Markdown

The **markdown** configurable defines what can be rendered in Markdown by Swarm for project overview pages and files stored in the Helix Core Server. By default, **markdown** is set to **safe**, Markdown text is displayed, but Markdown support is limited to prevent execution of raw HTML and JavaScript content.

Tip

- Valid Markdown file extensions are: **md**, **markdown**, **mdown**, **mkdn**, **mkd**, **mdwn**, **mdtxt**, **mdtext**.
- By default, project overview pages are displayed when there is a **README** Markdown file in the project mainline. This can be disabled by your Swarm administrator, see ["Project readme" on page 672](#).

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

Add or update the following configuration block to the `SWARM_ROOT/data/config.php` file, at the same level as the `p4` entry:

```
<?php
    // this block should be a peer of 'p4'
    'markdown' => array(
        'markdown' => 'safe', // default is 'safe'
    ),
```

- safe**: Markdown content is displayed, but Markdown support is limited to prevent execution of raw HTML and JavaScript content. This is the default.
- unsafe**: Markdown support is unrestricted, allowing full HTML and JavaScript to be used. This is insecure as any person with access to Swarm can add script to the Markdown which would execute as the currently logged in user.
- disabled**: Markdown text is not rendered and is only displayed as plain text. This is the most secure setting.

Note

Markdown content is displayed in comments and review descriptions, but Markdown support is limited to prevent execution of raw HTML and JavaScript content. This is the equivalent of **safe** mode and cannot be changed.

Menu helpers

The **menu_helpers** configuration block is used to control the visibility of Swarm menu items and to create new menu items for URLs and Swarm modules.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

Add a menu item to a menu

Add Swarm menu items in the **menu_helpers** configuration block in the [SWARM_ROOT/data/config.php](#) file, as in the following example:

```
<?php
    // this block should be a peer of 'p4'
    'menu_helpers' => array(
        'MyMenu01' => array( // A short recognizable name for the menu
item
            'id'          => 'custom01',           // A unique id for
the menu item. If not included in the array, parent array name is
used.
            'enabled'    => true,                  // When set to true,
the menu item is visible. Defaults to true if not included in the
array.
            'target'     => '<https://MyWebite>', // The URL or custom
module route a menu click takes you to.
                                                    // If not included
in array, id is used. If id not included, parent array name is used.
            'cssClass'   => 'custom_menu',         // The custom CSS
class name added to the menu item.
            'title'      => 'MyMenuItem',         // The text that
will be shown on the button.
                                                    // If not included
in array, id is used. If id not included, parent array name is used.
            'class'      => '',                   // If not included
in array or empty, the menu item is added to the main menu.
```

```

// To add the menu
item to the project menu for all of the projects, set to
'\Projects\Menu\Helper\ProjectContextMenuHelper'
    'priority' => 155, // The position the
menu item is displayed at in the menu.

// If not included
in the array, the menu item is placed at the bottom of the menu.
    'roles' => null, //
null|'authenticated'|'admin'|'super'

// If not included
in the array, null is the default

// Specifies the
minimum level of Perforce user that can see the menu item.
// 'authenticated' =
any authorized user, null = unauthenticated users
    ),
),

```

Configurables

Tip

You do not need to enter values for all of the configurables in a menu item array, configurables not included in the array will be set to their default values. For an example of this, see ["Create a menu item using only the custom module name" on page 644](#).

- **id** a unique id used by Swarm for the menu item, this can be anything. If **id** is not included in the array, Swarm uses the name of the parent array.
- **enabled**
 - **true** the menu item is visible subject to restrictions imposed by the **roles** value. This is the default value if **enabled** is not included in the array.
 - **false** the menu item is not in use
- **target** the URL or [Swarm module](#) route used when the menu item is clicked. If **target** is not included in the array, Swarm uses the **id** for the route. If **id** is not included in the array, Swarm uses the parent array name for the route.
- **cssClass** specifies the CSS class for the menu item, used to replace the default icon with a custom icon for the menu item, see ["Add a custom icon to a menu item" on the next page](#).

- **title** specifies the menu item name displayed in the menu. If **title** is not included in the array, Swarm uses the **id** for the menu title. If **id** is not included in the array, Swarm uses the parent array name for the menu title.
- **class** specifies whether the menu item is displayed in the main menu or the project menu.
 - `' '` or not included in the array, adds the menu item to the main menu.
 - `\Projects\Menu\Helper\ProjectContextMenuHelper` adds the menu item to the project menu of all of your projects.

Tip

You can restrict the menu item to a specific project by extending the **ProjectAwareMenuHelper** module, see ["Add a menu item to the project menu of a specific project" on page 644](#).

- **priority** specifies the menu item position in the menu, lower numbers are displayed higher up the menu structure. If priority is not included in the array, the menu item is placed at the bottom of the Swarm menu. If you use the same priority value as an existing menu item (not recommended), the menu items with the same priority are displayed in alphabetical order.

Tip

Enter `<yourSwarmURL>/api/v10/menus` in your browser address bar to return a list of the priority values for the Swarm menu items.

- **roles** restrict the display of the menu item based on the permissions of the logged in user:
 - **null** menu item visible to any user even if they are not authenticated. if roles is empty or not included in the array, null is used as the default.
 - **authenticated** menu item is visible to any authenticated user
 - **admin** menu item is visible to any user with Perforce *admin* permissions or higher
 - **super** menu item is visible to any user with Perforce *super user* permissions

Add a custom icon to a menu item

Add some custom CSS to Swarm to replace the default icon with a custom icon for the menu item. For example, the CSS below replaces the default icon with the **double-speech-bubbles.svg** icon for the **custom01** menu item by modifying the **svg.svgIcon.custom01MenuItem** and **.menuItem-custom01** classes:

```
.swarmMenu .menuItem.menuItem-tests .svgIcon{
    display:none;
}.swarmMenu .menuItem.menuItem-tests a::before{
    content: '';
    background-image: url('/swarm/img/icons/line/double-speech-
bubbles.svg');
```

```

background-size: 16px;
background-repeat: no-repeat;
padding-left: 20px;
    padding-right: 4px;
    margin-left: 20px;
}

```

Save these lines in a file, perhaps `menu_custom01.css` in a folder called `custom_menus`, in the `SWARM_ROOT/public/custom` folder. Swarm automatically includes the CSS file, immediately using your custom menus for subsequent page views. If you are using your own custom images, we recommend you store them in the same folder as your custom css.

Make Groups menu item visible to admin-user and above only

Note

The visibility of the Swarm **Groups** menu is controlled by a combination of the `role` setting here and the `super_only` setting in the `groups` configuration block, see ["Groups configuration" on page 624](#). Swarm uses the most restrictive of these two settings to control the visibility of the **Groups** menu item.

If you need fine grained control over the visibility of the **Groups** menu item, create a `groups` block in the `menu_helpers` configuration block and set the `roles` as required. In this example `roles` is set to `'admin'` so that only users with `admin-user` access and higher will be able to see the **Groups** menu item:

```

<?php
    // this block should be a peer of 'p4'
    'menu_helpers' => array(
        'groups' => array(
            'roles' => 'admin',
        ),
    ),
),

```

Tip

Enter `<yourSwarmURL>/api/v10/menus` in your browser address bar to return a list of `roles` for the Swarm menu items.

Create a menu item using only the custom module name

If you have created a Swarm custom module you can create a menu item for it that uses the custom module route and you do not need to add any other configurables for it in the **menu_helpers** configuration block.

For example, if you have a custom module with a route of '**codeAnalytics**', the configuration block only needs to be:

```
<?php
    // this block should be a peer of 'p4'
    'menu_helpers' => array(
        'codeAnalytics' // links the menu item to a module called
        'codeAnalytics'
    );
```

All of the other configurable settings for the **codeAnalysis** menu item will be set to their default values because the **codeAnalysis** array is empty:

- **id**: parent array name used (**codeAnalysis**)
- **enabled = true**, menu item is enabled
- **target**: parent array name used (**codeAnalysis**), this must match the custom module name
- **cssClass**: no class is set and Swarm uses the default chevron > icon for the menu item
- **class**: no class is set, the menu item is displayed in the main menu
- **priority**: no priority is set, the menu item is displayed at the bottom of the Swarm menu
- **roles**: **null**, visible to all users even if they are not authenticated

Add a menu item to the project menu of a specific project

To add a menu item in the project menu of a specific project, extend the **ProjectAwareMenuHelper** and set the menu item **class** to use the extended helper.

The example in this section only displays the **Jenkins Build** menu item in the **Jam** project menu.

Extend the ProjectAwareMenuHelper

Extend the **ProjectAwareMenuHelper** so that the **Jenkins Build** menu item is only displayed in the **Jam** project menu. This helper can be called anything and saved anywhere but it should have a descriptive name and be in a logical place, in this example it is **/modules/Project/src/Menu/Helper/JamOnlyMenuItem.php**:

```
<?php
/**
```



```

* Perforce Swarm
* @copyright 2020 Perforce Software. All rights reserved.
* @license Please see LICENSE.txt in top-level folder of this
distribution.
*@version <release>/<patch>
*/

namespace Projects\Menu\Helper;

/**
 * Only show the menu item using this class for the specified
Project.
 * In this example it is for the project with the id of Jam.
 * Class JamOnlyMenuItem
 * @package Projects\Menu\Helper
 */
class JamOnlyMenuItem extends ProjectAwareMenuHelper
{
    /**
     * Modifies an item's target if the context is for a project, the
project supports the item and
     * the item already has a target. Otherwise, nullify the item
     * @return array|null
     */
    public function getMenu()
    {
        if ($this->project && $this->project->getId() === 'jam') {
            $item = parent::buildMenu();
            // Allow project characteristics to determine menu item
availability
            return !empty($this->project) && $this->isDisabled() === false ?
$item : null;
        }
        return null;
    }
}

```

```
}
```

Define your menu item

The following block defines the **Jenkins Build** menu item you are adding to the **Jam** project menu. The important part is that `class` is set to `\Projects\Menu\Helper\JamOnlyMenuItem` so that it is only displayed in the **Jam** project.

```
<?php
    // this block should be a peer of 'p4'
    'menu_helpers' => array(
        'jenkins_jam' => array( // A short recognizable name for the
menu item
            'id'          => 'jenkins_jam',
            'enabled'     => true,
            'target'      => 'http://my-jenkins.instance.jam.com',
            'cssClass'    => 'custom_menu',
            'title'       => 'Jenkins Build',
            'class'       => '\Projects\Menu\Helper\JamOnlyMenuItem',
            'priority'    => 160,
            'roles'       => ['authenticated'],
        ),
    ),
```

Delete Swarm config cache to enable the menu item

Your changes will not be used by Swarm until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

Multiple Helix Core Server instances

A single Swarm server can be connected one or more Helix Core Server (P4D) instances. This section describes how to configure Swarm to connect to multiple Helix Core Servers.

Warning


Issue: Swarm will lose connection to all of the Helix Core Servers if you edit the `base_url` configurable value in the `environment` block of `<swarm_root>/data/config.php`. This will stop your system working.

Fix: Remove the `base_url` configurable from the `environment` block of `<swarm_root>/data/config.php`.

Note

Known limitations, only applies if you want to use "Global Dashboard" on page 310:

- **Issue:** If Helix Authentication Service is enabled for Swarm and the **Try to login to all available servers with these credentials** checkbox or the **All available servers** option is selected in a login dialog, Swarm will not try to log in to any of the other Helix Core Server instances that are configured for Helix Authentication Service.

Workaround: Log in to them individually using the instance **Log in** button  in the sidebar or by including the server instance name in the URL, for example:

`https://swarm.company.com/serverA`.

- Swarm must be installed in its own virtual host.

Complete the following steps to connect Swarm to multiple Helix Core Server instances:

- "Set up the Swarm configuration file for the Helix Core Servers" below
- "Set the Swarm trigger token and Swarm host variable for each Helix Core Server" on page 657
- "Configure a cron job for each Helix Core Server instance" on page 658

Set up the Swarm configuration file for the Helix Core Servers

The Swarm Helix Core Server connection details are configured using the `p4` item in the `<swarm_root>/data/config.php` configuration file.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the "Reload Configuration button" on page 732.

When Swarm is configured for a single Helix Core Server instance the `p4` item looks like this:

```
<?php
    return array(
```

```

        'p4' => array(
            'port'      => 'my-helix-core-server:1666',
            'user'      => 'admin_userid',
            'password' => 'admin user ticket or password',
            'sso'       => 'disabled', //
            ['disabled'|'optional'|'enabled'] default value is 'disabled'
        ),
    );

```

Where:

- **port** is the P4PORT for your Helix Core Server.
- **user** is a *userid* that has *admin* privileges for the Helix Core Server.
- **password** is the ticket or password for the Helix Core Server user.
- **sso** if the Helix Core Server is configured for the Helix Authentication Service, set to one of the following:
 - **enabled** all users must use Helix Authentication Service to log in to Swarm.
 - **optional** Helix Authentication Service is available for users to log in to Swarm but is not enforced.
 - **disabled** Helix Authentication Service is not available to Swarm. This is the default value.

Important

From Swarm 2021.1, the **sso_enabled** configurable is deprecated but remains supported. It is replaced with the more flexible **sso** configurable. If the **sso_enabled** configurable and **sso** configurable are both present in the **p4** configuration block, Swarm uses the **sso** configurable value.

Multiple Helix Core Server instances can be added to the **p4** item in **config.php** but each Helix Core Server instance must have a label to identify the server instance. This enables Swarm to connect to the Helix Core Servers in the **p4** item.

Tip

The Helix Core Server instances do not have to all be at the same version but they must all be versions that Swarm supports. For a list of Helix Core Server versions supported by Swarm, see [Helix Core Server requirements](#).

The first Helix Core Server in the **p4** item is the primary Helix Core Server, this instance renders the global dashboard. If Swarm fails to connect to the primary Helix Core Server, Swarm cannot display any of the Helix Core Server instances in the global dashboard. You can still connect directly to the other P4D instances by including the Helix Core Server label in the URL for each instance, for example: **https://swarm.company.com/serverA**.

Configure the Swarm configuration file for multiple Helix Core Server instances:

Enter a label for each Helix Core Server instance and enter the server connection details under the server label.

For example:

```
<?php
    return array(
        'p4' => array(
            'serverA' => array(
                'port'      => 'p4d-A:1666',
                'user'      => 'admin_userid_serverA',
                'password' => 'admin user ticket or password for
serverA',
                'sso'      => 'disabled', //
                ['disabled'|'optional'|'enabled'] default value is 'disabled'
            ),
            'serverB' => array(
                'port'      => 'p4d-B:1666',
                'user'      => 'admin_userid_serverB',
                'password' => 'admin user ticket or password for
serverB',
                'sso'      => 'disabled', //
                ['disabled'|'optional'|'enabled'] default value is 'disabled'
            ),
        );
```

Where:

- **serverA** and **serverB** are labels that identify the individual Helix Core Servers.
 - The labels must be URL-friendly labels.
 - The labels must not contain any spaces or any characters that require URL encoding, and
 - The total length of the server label and **Redis namespace** combined is limited to ≤ 127 characters. The default **namespace** is **Swarm**.
- The **port**, **user**, **password**, and **sso** items are specific to the Helix Core Server instance they are nested under.

Configure a specific Helix Core Server

The root element in the `<swarm_root>/data/config.php` configuration file contains default values for all Helix Core Server instances. You can override the default values by using a server specific configuration file. For example, to override the global value of `jira` configuration for a specific server, you can provide a different value for `jira` configuration in the server specific configuration file.

To configure a specific Helix Core Server instance, add a `config.php` file in the `<swarm_root>/data/servers/<serverid>` directory. For the new Helix Core Server configuration to take effect, the configuration cache must be deleted and reloaded. For more information on how to delete and reload cache, see ["Cache Info" on page 731](#).

Important

- You must include a global `config.php` file to specify the multiple Helix Core Server instances.
- You must ensure that a `p4` item is not included in the Helix Core Server specific configuration. If the `p4` item is included, no error message is displayed and the `p4` item is ignored by the Helix Core Server.

Note

Helix Core Server specific configuration overrides the global configuration.

Configure event notification for each Helix Core Server

Important

You must use either Helix Core Server Extensions (recommended) or Helix Core Server Triggers to notify Swarm about events on the Helix Core Servers. You cannot use a mixture of the two methods.

Tip

Triggers are still supported, but we recommend you use Helix Core Server Extensions. Helix Core Server Extensions are easier to install and maintain than Triggers.

Swarm needs to know about a number of Helix Core Server events to operate correctly, this can be done by using Helix Core Server Extensions (recommended) or Helix Core Server Triggers. Swarm installs include the Swarm Helix Core Server extension file and trigger scripts required for Swarm to get the events it needs from your Helix Core Server.

You must install and configure the Swarm Helix Core Server extension or Triggers to complete the Swarm configuration.

Do one of the following so that Swarm is notified about events on the Helix Core Servers it is connected to:

- Install and configure the Swarm Helix Core Server extension for each Helix Core Server instance, see ["Install and configure the Swarm Helix Core Server extension \(recommended\)" below](#)
- Set the Swarm trigger token and host variable for each Helix Core Server instance, see ["Set the Swarm trigger token and Swarm host variable for each Helix Core Server" on page 657](#)

Install and configure the Swarm Helix Core Server extension (recommended)

Important

- If you are using the Swarm Helix Core Server extension, Swarm Helix Core Server Triggers must not be installed.
- The following commands must be run on each Helix Core Server instance Swarm is connected to.
- You must be a Helix Core Server *super* user to install and configure Helix Core Server Extensions.

Prerequisites

To install the Swarm Helix Core Server extension you need:

A compatible version of Helix Core Server for Helix Core Server Extensions:

- **Linux:** Helix Core Server 2021.2 and later. If you are using an earlier version of Helix Core Server, you must use triggers.
- **Windows:** Helix Core Server 2021.2 and later. If you are using an earlier version of Helix Core Server, you must use triggers.

You will also need:

- Extensions installed and configured on your Helix Core Server.
- A user with *super* permissions to install and configure the Swarm Helix Core Server extension.

Swarm Helix Core Server extension installation

The Swarm Helix Core Server extension file is included for all of the Swarm installation types. The **helix-swarm.p4-extension** file is located in ["swarm_root" on page 724/p4-bin/extensions](#)

On each Helix Core Server you are connecting Swarm to, run the following commands as a Helix Core Server user with *super* permissions:

1. Install the Swarm Helix Core Server extension on the Helix Core Server with:

```
p4 extension --yes --install helix-swarm.p4-extension
```

Example response:

```
Extension 'Perforce::helix-swarm' version '2022.1.20221215'
installed successfully.
Perform the following steps to turn on the Extension:

# Create a global configuration if one doesn't already exist.
p4 extension --configure Perforce::helix-swarm

# Create an instance configuration to enable the Extension.
p4 extension --configure Perforce::helix-swarm --name
Perforce::helix-swarm-instanceName

For more information, visit:
https://www.perforce.com/manuals/extensions/Content/Extensions/Ho
me-extensions.html
```

2. Create a global configuration if one doesn't already exist:

```
p4 extension --configure Perforce::helix-swarm
```

The spec file opens in your text editor, for example:

```
ExtName:          helix-swarm
ExtDescription:
    Helix Swarm Extension

ExtVersion:       2022.1.DEV.20220120

ExtUUID:          4532BC59-7BC8-478F-ADF6-0A563C42563D

ExtRev: 1

ExtMaxScriptTime:      unset

ExtMaxScriptMem:       unset

ExtAllowedGroups:
```



```

ExtEnabled:      true

ExtP4USER:       sampleExtensionsUser

Name:    helix-swarm

Owner:    super

Update: 2022/01/21 10:43:46

Description:
    The description of your config.

ExtConfig:
    Debug:
        2

    Swarm-Secure:
        true

    Swarm-Token:
        ... SWARM-TOKEN

    Swarm-URL:
        http://localhost/

```

3. Edit the spec file to match your system:

- **ExtP4USER:** Set to an existing user with *super* permissions.
- **ExtConfig:** Check and edit the values in this block:
 - **Debug:**

Debug levels **0** to **3** control the amount of debug information sent to the Helix Core Server Extensions log.

Important

A debug level of **10** and higher sends all debug information to the client. This is useful for debugging, but should not be run in a production environment.

- **SSL-CA-File:** (optional) Set the absolute path of the [Certificate Authority \(CA\) Privacy-Enhanced Mail \(PEM\)](#) file to validate the SwarmHelix Core Server's certificate. Ensure that this file is accessible to the Helix Core Server process owner.

If a suitable certificate is not found in the local SSL certificate store maintained by the operating system or if there are verification issues from using a self-signed certificate then the default CA file is used.

By default, the local certificate stored in the default CA file is used if the absolute path for the [Certificate Authority \(CA\)](#) is not specified.

- **Swarm-Secure:**
 - `true` will only accept secure SSL certificates.
 - `false` will accept insecure SSL certificates.
- **Swarm-Token:** Set to the Swarm trigger token value.

Important

To set multiple Helix Core Server token settings for Extensions, repeat the following steps for each Helix Core Server. The process to obtain an extension token of your Swarm instance is similar to obtaining a trigger token of your Swarm instance.

To obtain the trigger token of your Swarm instance:

- a. Log in to Swarm as a *super* user.
- b. Click your *userid*, found at the right of the main toolbar.
- c. Select **About Swarm**.
- d. The **About Swarm** dialog is displayed and Swarm generates an API token if it doesn't already exist.
- e. Copy the trigger token value displayed at the bottom of the dialog and paste the token into the spec file.

- **Swarm-URL:** Set to the URL of your Swarm instance.

4. When you have finished editing the spec file, save it in your text editor.
5. Create an instance configuration to enable the extension:

```
p4 extension --configure Perforce::helix-swarm --name swarm
```

The spec file opens in your text editor, for example:

```
ExtName:          helix-swarm

ExtDescription:
    Helix Swarm Extension

ExtVersion:       2022.1.DEV.20220120

ExtUUID:          4532BC59-7BC8-478F-ADF6-0A563C42563D

ExtRev: 1

ExtMaxScriptTime:      unset

ExtMaxScriptMem:       unset

ExtEnabled:           true

ExtDebug:             none

Name:    swarm

Owner:   super

Update: 2022/01/21 10:58:41

Description:
    The description of your config.

ExtConfig:
    depot-path:
        //...
    enableStrict:
        true
```

```
enableWorkflow:
    true
httpTimeout:
    30
ignoreErrors:
    false
```

We recommend that these settings are left at their default values, but the configuration needs to be run to set the default values.

- **enableStrict**: Enable strict checking of reviews for workflow. Verifies that the file content in a commit matches the file content of its associated approved review. If one or more files in a commit do not match the content of the file in its associated review, the commit is rejected.
 - **enableWorkflow**: Enable workflow.
 - **httpTimeout**: Timeout (in seconds) when communicating with the Swarm server.
 - **ignoreErrors**: If the server returns an error (timeout, not there), then default to allowing a request. It means workflow rules can be skipped, but if the Swarm server is down it will not block submits.
6. Save the spec file in your text editor.
 7. Confirm your Swarm Helix Core Server extension is installed and configured by running the following command on the Helix Core Server:
p4 extension --run swarm ping
 - If the Swarm Helix Core Server extension is working, the Helix Core Server responds with **OK**
 - If the Swarm Helix Core Server extension is not working, the Helix Core Server responds with an error message. For example, **BAD (Cannot reach https://swarm-example.com/)**For more information about Helix Core Server Extensions, see:
 - [Extension overview](#) in the *Helix Core Extensions Developer Guide*.
 - [p4 extension](#) in the *Helix Core Command-Line (P4) Reference*.
 8. Repeat these steps for each Helix Core Server.
 9. When you have configured all of the Helix Core Servers, configure a cron job to start Swarm workers for each Helix Core Server instance, see "[Configure a cron job for each Helix Core Server instance](#)" on page 658.

Set the Swarm trigger token and Swarm host variable for each Helix Core Server

Important

If you are using Swarm Helix Core Server Triggers, the Swarm Helix Core Server extension must not be installed.

Helix Core Server uses a Swarm trigger token to confirm that trigger requests from Swarm are valid. Each Helix Core Server instance must have a valid Swarm trigger token in its **swarm-trigger.conf** file. The **swarm-trigger.conf** file also contains the Swarm host URL for the Helix Core Server instance.

Set the Swarm trigger token and Swarm host variable for each Helix Core Server instance:

1. Navigate to the Swarm URL for the instance.
For example: **https://swarm.company.com/serverA**

2. Log in to Swarm as a *super* user.
3. Click your *userid*, in the main toolbar and select **About Swarm**.

The **About Swarm** dialog is displayed with a **Trigger Token**. Swarm generates the trigger token if it doesn't already exist.

4. Copy the trigger token value from the dialog.
5. Open the **swarm-trigger.conf** file for the Helix Core Server instance.

Tip

For more information about the location of the **swarm-trigger.conf** file, see ["Set up Swarm triggers with a Windows or Linux-hosted Helix Core Server" on page 205](#).

6. In the **swarm-trigger.conf** file:
 - a. Set the **SWARM_HOST** URL.
 - b. Set the **SWARM_TOKEN** by pasting the Swarm trigger token you copied in the steps above.

For example:

```
# SWARM_HOST (required)
# Hostname of your Swarm instance, with leading "http://" or
# "https://".
SWARM_HOST="https://swarm.company.com/serverA"

# SWARM_TOKEN (required)
# The token used when talking to Swarm to offer some security. To
# obtain the
# value, log in to Swarm as a super user and select 'About Swarm'
# to see the
# token value.
SWARM_TOKEN="TRIGGER-TOKEN-FOR-SERVERA"
```

7. Save the **swarm-trigger.conf** file.
8. Repeat these steps for each Helix Core Server.
9. When you have configured all of the Helix Core Servers, configure a cron job to start Swarm workers for each Helix Core Server instance, see ["Configure a cron job for each Helix Core Server instance" below](#).

Tip

Alternatively, you can simply touch a file in each **<Swarm root>/data/servers/<serverid>/tokens** folder. The file content is ignored, the filename itself is the token.

This allows you to specify a common token that is used by all of the Helix Core Server instances.

However, we recommend that you use a separate token for each Helix Core Server instance. This makes it easier to invalidate a token for a specific Helix Core Server by deleting the file in the **tokens** folder for that server if you need to.

Configure a cron job for each Helix Core Server instance

To automatically spawn workers for each of the servers connected to Swarm, you must manually configure a cron job for each of the servers.

Swarm package installations only

The **swarm-cron-hosts.conf** file specifies the connection type (HTTP or HTTPS), hostname, port number and, server label for Swarm cron jobs.

If you have installed Swarm via packages, you must specify all of the Helix Core Server URLs within **/opt/perforce/etc/swarm-cron-hosts.conf**.

1. Edit the **swarm-cron-hosts.conf** file so that it contains the actual Swarm hostname, ports, and server labels you have configured for Swarm.

The following format is used with one Helix Core Server on each line:

```
[http[s]://]<swarm-host>[:<port>] [/<base-url>]
```

Default if value not specified:

- [http[s]://] http
- <swarm-host> must be specified
- [:<port>] 80
- [/<base-url>] server label, must be specified

For example, for **serverA** and **serverB** configured earlier on a Swarm host of *https://swarm.company.com* with a default port value of 80. The entries in the **swarm-cron-hosts.conf** file would be:

```
https://swarm.company.com/serverA
https://swarm.company.com/serverB
```

2. Save the **swarm-cron-hosts.conf** file.
3. Swarm is now configured to connect to multiple Helix Core Server instances.

Tip

To check or modify Swarm worker configuration, see ["Worker configuration" on page 750](#).

Swarm Tarball installation only

If you have installed Swarm from a tarball or configured cron manually, you need create a file called **helix-swarm** and add all of the Helix Core Server instances connected to Swarm.

1. Create a file named **helix-swarm** in **/etc/cron.d** if it does not already exist.
2. Edit the **helix-swarm** file so that it has the following content:

```
#
# Cron job to start Swarm workers every minute
#
* * * * * nobody [ -x /opt/perforce/swarm/p4-bin/scripts/swarm-
cron.sh ] && /opt/perforce/swarm/p4-bin/scripts/swarm-cron.sh
```

3. Save the **helix-swarm** file.
4. Edit the **swarm-cron-hosts.conf** file so that it contains the actual Swarm hostname, ports, and server labels you have configured for Swarm.

The following format is used with one Helix Core Server on each line:

```
[http[s]://]<swarm-host>[:<port>][/<base-url>]
```

Default if value not specified:

- `[http[s]://]` http
- `<swarm-host>` must be specified
- `:<port>` 80
- `[/<base-url>]` server label, must be specified

For example, for **serverA** and **serverB** configured earlier on a Swarm host of `https://swarm.company.com` with a default port value of 80. The entries in the **swarm-cron-hosts.conf** file would be:

```
https://swarm.company.com/serverA
https://swarm.company.com/serverB
```

5. Save the **swarm-cron-hosts.conf** file.
6. Swarm is now configured to connect to multiple Helix Core Server instances.

Tip

To check or modify Swarm worker configuration, see ["Worker configuration" on page 750](#).

Further information

Users

- To visit the Global Dashboard, enter the basic Swarm URL without a Helix Core Server instance name, for example: `https://swarm.company.com`.
- To visit a specific Helix Core Server instance in Swarm without going via the global dashboard, include the server name in the URL, for example:
`https://swarm.company.com/serverA`.

Once you are viewing the Helix Core Server in Swarm, Swarm works as a standard single Helix Core ServerSwarm system.

Tip

- To browse jobs on **serverA**, navigate to:
`https://swarm.company.com/serverA/jobs`
- To browse reviews on **serverB**, navigate to:
`https://swarm.company.com/serverB/reviews`
- To view the dashboard for **serverB**, navigate to
`https://swarm.company.com/serverB/#actionable-reviews`

- If you don't include the server label in the URL you will be taken to the specified page for the primary Helix Core Server.

For example: navigating to `https://swarm.company.com/reviews` will redirect you to `https://swarm.company.com/serverA/reviews` because **serverA** is the [Primary Helix Core Server](#) in the **p4** configuration item.

Administrators

On the web server hosting Swarm, Swarm automatically creates a data folder for each Helix Core Server instance in `<Swarm_root>/data/servers`. This is because each Helix Core Server request is a field.

For example:

```
# ls -la /opt/perforce/swarm/data/servers/serverA
total 362296
drwx----- 6 apache apache    4096 Sep  8 17:15 .
drwx----- 6 apache apache    4096 Sep  6 15:41 ..
drwx----- 2 apache apache    4096 Sep 20 18:08 cache
drwx----- 3 apache apache    4096 Sep  7 13:25 clients
-rw-r--r-- 1 apache apache 3709543 Sep 26 14:19 log
-r----- 1 apache apache      84 Sep  6 15:41 p4trust
drwx----- 4 apache apache    4096 Sep 20 14:27 queue
drwx----- 2 apache apache    4096 Sep 20 11:18 sessions
```

Tip

It is important to understand that there will be a Swarm log file for each Helix Core Server instance.

Developers

To get a list of projects via the Swarm API for *serverA*, run bash:

- If you are using **wget**:

```
$ wget -u "apiuser:password"
https://swarm.company.com/serverA/api/v10/projects
```

- If you are using **curl**:

```
$ curl -u "apiuser:password"
https://swarm.company.com/serverA/api/v10/projects
```

Tip

If you don't include the server label in the URL you will be taken to the projects page for the primary Helix Core Server.

For example: navigating to `https://swarm.company.com/api/v10/projects` will redirect you to `https://swarm.company.com/api/v10/serverA/projects` because **serverA** is the [Primary Helix Core Server](#) in the **p4** configuration item.

Notifications

This section describes the configurables that can be set for notifications.

Note

For more information on how to delay comment notifications, and configure comment threading, see ["Comments" on page 585](#).

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

Swarm can be configured to provide generic notifications of committed changes in Helix Core Server, taking the role of a review daemon.

Notifications configuration is expressed with a **notifications** block in the [SWARM_ROOT/data/config.php](#) file, similar to the following example (defaults shown):

```
<?php
// this block should be a peer of 'p4'
'notifications' => array(
    'honor_p4_reviews'      => false, // defaults to false
    'opt_in_review_path'    => '',    // required if honor_p4_
reviews is true; defaults to ''
    'disable_change_emails' => false, // optional; defaults to
false
),
```

If **honor_p4_reviews** is set to **true**, then **opt_in_review_path** must be set to a path somewhere in the depot. This path does not need to point to an actual file that exists, but it must be accessible by all users who want to make use of this functionality. For example:

```
'notifications' => array(
    'honor_p4_reviews'      => true,
```

```
'opt_in_review_path' => '//depot/swarmReviews',
),
```

If these two values are set, then users can make use of the Perforce review functionality by subscribing to the **opt_in_review_path** in their user spec. Any user subscribed to that file, will receive notifications for all the other paths they are subscribed to.

We recommend that **opt_in_review_path** point to a file that does not exist. Ideally, it points to a file that no user is likely to want to create. It must however be in a valid depot.

For example, if a user has the following review paths set in their user spec:

```
$ p4 user -o asmith
User:      asmith

Email:     asmith@example.com

FullName:   Alice Smith

Reviews:
    //depot/swarmReviews
    //depot/main/acme/...
    //depot/main/orion/...
    //depot/dev/asmith/...
    //streams/main
```

The **//depot/swarmReviews** means that this user is subscribed to the path set in **opt_in_review_path**, and therefore will receive notifications.

The subscription lines ending in **/...**, define which paths in the depot the user is interested in. For example, this user will receive a notification for a change made to **//depot/main/acme/foo.txt**, but not a change made to **//depot/dev/acme/foo.txt**.

To receive notifications for stream spec changes in a path, leave the **/...** off of the end of the subscription line. For example, the **//streams/main** subscription line means the user will receive a notification for a change made to a stream spec in **//streams/main**.

To see which users are subscribed to receive notifications you can run **p4 reviews <path>** against the **opt_in_review_path** value:

```
$ p4 reviews //depot/swarmReviews
asmith <asmith@example.com> (Alice Smith)
bbrown <bbrown@example.com> (Bob Brown)
erogers <erogers@example.com> (Eve Rogers)
```

To see which users are subscribed to files in a particular changelist, you can run `p4 reviews -c <changelist>`. Swarm will notify the users who subscribe to both the review of this changelist and the review path, `opt_in_review_path`.

- **honor_p4_reviews**: When set to `true`, Swarm sends notification emails for every committed change to all users where the change matches one of their **Reviews**: paths.
- **opt_in_review_path**: Optional item, required only if **honor_p4_reviews** is set. This item specifies a special depot path, which typically would not exist in the Helix Core Server machine. When a path is specified, users must include this path (or a path that contains it) in the **Reviews**: field of their user spec to cause Swarm to send the user a notification for every committed change that matches one of their **Reviews**: paths.

For example, if the `opt_in_review_path` is set to `//depot/swarmReviews`, users can opt-in to Swarm review notifications by adding that path, or a path such as `//depot/...`, to the **Reviews**: field in their user spec.

- **disable_change_emails**: Optional item. When set to `true`, notifications for committed changes, based on the **Reviews**: field and the users and projects you follow, are disabled. Notifications for reviews and comments will still be sent.

Note

If your Helix Core Server machine already has a review daemon in operation, users receive two notifications for **Reviews**: paths. You may want to deprecate the review daemon in favor of Swarm's change notifications.

Note

"Groups" on page 373 have per-group notification settings. See "Add a group" on page 491 for details.

Global settings

There are many situations that can result in email notifications being sent out to users and groups. Whilst it is possible for a user and group owner to configure their own settings, it is also possible for the system owner to configure the defaults for all users and groups by modifying the settings in the `config.php`.

Note

- **If a group owner is not specified for the group**: only users with *super* privileges can configure group notification settings.
- **If one or more group owners are specified for the group**: only group owners and users with *super* privileges can configure group notification settings.

Each notification consists of an **Event** and a **Role**. The **Event** is what happened (for example, a new review was created, a file was submitted) and the **Role** is the role of the user or group who could receive a notification. A user or group can belong to multiple roles, in which case if any of them are set to send a notification, then the user or group will receive a notification.

For example, when a review is voted on (**review_vote**), there are a number of roles of users, and groups that could be notified:

- **Users:** the user that voted on the review (**is_self**), the author of the review (**is_author**), a user who is a moderator of the project branch the review is in (**is_moderator**), and a user who is a reviewer of the review (**is_reviewer**).
- **Groups:** members of a moderator group for the project branch the review is in (**is_moderator**), and members of a reviewer group for the review (**is_reviewer**).

Configuration options

By default, all notifications are enabled for all roles. The system-wide defaults can be changed by adding the following options into the notifications block of the [SWARM_ROOT/data/config.php](#). These options are in addition to those described above.

```
<?php
// this block should be a peer of 'p4'
'notifications' => array(
    'review_new' => array(
        'is_author' => 'Enabled',
        'is_member' => 'Enabled',
    ),
    'review_files' => array(
        'is_self'      => 'Enabled',
        'is_author'    => 'Enabled',
        'is_reviewer'  => 'Enabled',
        'is_moderator' => 'Enabled',
    ),
    'review_vote' => array(
        'is_self'      => 'Enabled',
        'is_author'    => 'Enabled',
        'is_reviewer'  => 'Enabled',
        'is_moderator' => 'Enabled',
    ),
    'review_required_vote' => array(
        'is_self'      => 'Enabled',
        'is_author'    => 'Enabled',
```

```
        'is_reviewer' => 'Enabled',
        'is_moderator' => 'Enabled',
    ),
    'review_optional_vote' => array(
        'is_self' => 'Enabled',
        'is_author' => 'Enabled',
        'is_reviewer' => 'Enabled',
        'is_moderator' => 'Enabled',
    ),
    'review_state' => array(
        'is_self' => 'Disabled',
        'is_author' => 'Enabled',
        'is_reviewer' => 'Enabled',
        'is_moderator' => 'Enabled',
    ),
    'review_tests' => array(
        'is_author' => 'Enabled',
        'is_reviewer' => 'Enabled',
        'is_moderator' => 'Enabled',
    ),
    'review_changelist_commit' => array(
        'is_author' => 'Enabled',
        'is_reviewer' => 'Enabled',
        'is_member' => 'Enabled',
        'is_moderator' => 'Enabled',
    ),
    'review_comment_new' => array(
        'is_author' => 'Enabled',
        'is_reviewer' => 'Enabled',
    ),
    'review_comment_update' => array(
        'is_author' => 'Enabled',
        'is_reviewer' => 'Enabled',
    ),
    'review_comment_liked' => array(
        'is_commenter' => 'Enabled',
```

```

    ),
    'review_opened_issue' => array(
        'is_self'          => 'Enabled',
        'is_author'        => 'Enabled',
        'is_reviewer'      => 'Enabled',
        'is_moderator'     => 'Enabled',
    ),
    'review_join_leave' => array (
        'is_self'          => 'Enabled',
        'is_author'        => 'Enabled',
        'is_reviewer'      => 'Enabled',
        'is_moderator'     => 'Enabled',
    ),
)
)

```

Each setting can have one of four possible values to either enable or disable notifications of that type. If multiple settings apply to a given event, then a user will receive a notification if *any* of them are enabled.

- **Enabled:** Notifications for this event and role are enabled, and a user will receive emails by default.
- **Disabled:** Notifications for this event and role are disabled, and a user will not receive emails by default.
- **ForcedEnabled:** Same as for Enabled, but this is forced enabled for all users. An individual user is not able to disable this notification type on their settings page.
- **ForcedDisabled:** As for Disabled, but this is forced disabled for all users. An individual user will not be able to enable this notification type on their settings page.

Unless one of the forced options is used, system wide options can be overridden by individual users and group owners, who can configure which notifications they receive.

Notification Roles

The various roles are as follows:

- **is_self:** This role is the user who changed the state of the review.
- **is_author:** This role is the user who was the author of the review.
- **is_reviewer:** This role includes all users and groups who are listed as being a reviewer on the review.
- **is_member:** This role includes all users who are a member of the project in which the event happened.

- **is_moderator:** This role includes all users and groups who are listed as being a moderator of the project branch in which the event happened.
- **is_follower:** This role includes all users who are followers of the project in which the event happened.
- **is_commenter:** This role is the user who was the author of a comment.

Notification events

An event is the action that causes the notification:

- **review_new:** A new review has been created.
- **review_files:** Files have been added to a review.
- **review_vote:** A user has voted on a review.
- **review_state:** The status of a review has been changed.
- **review_tests:** Automated tests have failed for a review. The first time automated tests pass for a review after a test failure for that review.
- **review_changelist_commit:** Files on a review have been committed.
- **review_comment_new:** A comment has been added to a review.
- **review_comment_update:** A comment on a review has been updated.
- **review_comment_liked:** A user has liked a comment.
- **review_join_leave:** A user or group has joined or left a review.

Obliterate Review

Note

- By default, you must be a user with *admin* or *super* user rights to obliterate a review.
- **Optional:** Swarm can be configured to allow users to obliterate reviews that they have authored. This can be configured by your Swarm administrator. See "[Allow author obliterate review](#)" on page 688.

Obliterate is used to permanently delete reviews that have been created by mistake. For instance, if a review is associated with the wrong changelist, or a review contains sensitive information that should not be openly available.

When you obliterate a review


- **All reviews:**
 - Metadata associated with the review is permanently deleted.
 - Events associated with the review are permanently deleted from the activity feed.
 - The Swarm 404 page is displayed if a user navigates to the review from a Swarm notification email link, or by using the URL.
- **Review with a shelved changelist:**
 - The Swarm review changelist, and any associated shelved files are permanently deleted.
 - Files shelved in the user's associated changelist are left intact.
- **Review with a committed changelist:**
 - The Swarm review changelist is permanently deleted.
 - The committed changelist is left intact.

Obliterate a review

Important

Obliterate must be used with care, the review and all of its associated metadata are permanently deleted. An obliterated review cannot be reinstated, not even by Perforce Support.

To obliterate a review:

1. Navigate to the review.
2. Click the **Review actions**  button and select **Obliterate Review**.
3. Click **Yes** on the confirmation dialog to complete the obliterate action.
4. The review is obliterated.

P4TRUST

When Swarm is configured to connect to a Helix Core Server (P4D) using an SSL connection, Swarm automatically executes the `p4 trust` command, which accepts the SSL fingerprint and creates a `p4trust` file containing a list of trusted servers and their fingerprints.

The location the `p4trust` file is saved to depends on whether Swarm is connected to a single Helix Core Server or to multiple Helix Core Servers.

- **Single Helix Core Server:** saved as `SWARM_ROOT/data/p4trust`
- **Multiple Helix Core Servers:** saved as a separate file for each server. For example, for `serverA`, `serverB`, and `serverC` they are saved as:
 - `SWARM_ROOT/data/serverA/p4trust`
 - `SWARM_ROOT/data/serverB/p4trust`
 - `SWARM_ROOT/data/serverC/p4trust`

If a certificate changes

If a certificate for a Helix Core Server is changed for any reason then Swarm connections to that server will fail after that server is restarted.

The solution is to delete the `p4trust` file for that Helix Core Server from the location described above. Swarm will automatically run `p4 trust` on the next request if the `p4trust` file is not found.

Projects

This section describes the configurables that can be set for projects.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

Restrict project name and branch definition editing to administrators

By default, once a project has been created, any member of the project can edit or delete the project's settings. Projects can also be set to **Only Owners and Administrators can edit the project**, only owners and administrators can make changes to the project.

Instead of allowing any changes, or preventing all changes, you may want to prevent project members from making specific changes, such as to the project's name (and associated identifier), or adjusting the branch definition(s). This is useful when build infrastructure or other tooling treats these details as operational configuration, but you still want members to be able to adjust other aspects of the project configuration.

To do so, edit the `SWARM_ROOT/data/config.php` file, and set the following two items, similar to the following example:

```
'projects' => array(  
    'edit_name_admin_only'    => true,
```

```
'edit_branches_admin_only' => true,
),
```

- **edit_name_admin_only**: when set to true, only users with *admin* privileges in the Helix Core Server can modify a project's name.
- **edit_branches_admin_only**: when set to true, only users with *admin* privileges in the Helix Core Server can modify a project's branch definition(s).

Both items default to **false**.

Limit adding projects to administrators

By default, any authenticated user can add new projects. Swarm can restrict project creation to users with *admin*-level privileges or higher. Once restricted, Swarm prevents non-administrators from adding projects, and does not display the + icon to add a project to non-administrators.

Add or update the following configuration block to the `SWARM_ROOT/data/config.php` file, at the same level as the `p4` entry:

```
<?php
// this block should be a peer of 'p4'
'projects' => array(
    'add_admin_only' => true,
),
```

Important

If **add_admin_only** is enabled and **add_groups_only** has one or more groups configured, project creation is only available to users with administrator privileges and who are members of the specified groups.

Limit adding projects to members of specific groups

Swarm can restrict project creation to members of specific groups. The groups and membership need to be defined in the Helix Core Server.

Add or update the following configuration block to the `SWARM_ROOT/data/config.php` file, at the same level as the `p4` entry:

```
<?php
// this block should be a peer of 'p4'
'projects' => array(
    'add_groups_only' => array('wizards', 'slayers', 'phbs'),
),
```

Important

If `add_admin_only` is also enabled, project creation is only available to users with administrator privileges **and** who are members of the specified groups.

Project readme

By default, projects can have a **README** markdown file associated with them that is automatically displayed on the project overview page. This file is read from the root of the project's mainline if it is available, and displayed on the project page. See ["Mainline branch identification" on page 636](#) for details on configuring the project mainline.

Tip

- Valid Markdown file extensions are: `md`, `markdown`, `mdown`, `mkdn`, `mkd`, `mdwn`, `mdtxt`, `mdtext`.

If more than one **README** file is found, Swarm displays the first one it finds based on the order above.

- By default, Markdown support is limited to prevent execution of raw HTML and JavaScript content. The level of Markdown support can be configured by your Swarm administrator to **disabled**, **safe**, and **unsafe**, see ["Markdown" on page 639](#).

Add or update the following configuration block to the `SWARM_ROOT/data/config.php` file, at the same level as the `p4` entry:

```
<?php
// this block should be a peer of 'p4'
'projects' => array(
    'readme_mode' => 'enabled',
),
```

- enabled**: the use of **README** markdown files is enabled for project overview pages, text content is displayed for project overview pages with a README markdown file. This is the default.
- disabled**: the use of **README** markdown files is disabled for project overview pages, project overview pages are not displayed.

Note

In Swarm 2018.3 and earlier, `readme_mode` could be set to **disabled**, **restricted**, or **unrestricted**. If your `readme_mode` was originally set to **restricted** or **unrestricted** and you have upgraded to Swarm 2019.1 or later, Swarm treats `readme_mode` as if it is set to **enabled**. Swarm uses the `markdown` configurable to determine the level of Markdown support (default is **safe**), see ["Markdown" on page 639](#).

Projects tab initial fetch

By default, all of the projects are fetched for the **Projects** tab with a single call. If your Swarm system has a large number of projects, it can take some time to populate the **Projects** tab. The **fetch** configurable is used to tell Swarm to fetch and display the first **x** projects and then load the rest from cache in the background. This speeds up the initial display of projects in the **Projects** tab.

Add or update the following configuration block to the `SWARM_ROOT/data/config.php` file, at the same level as the `p4` entry:

```
<?php
    // this block should be a peer of 'p4'
    'projects' => array(
        'fetch' => array('maximum' => 50), // defaults to 0 (disabled)
    ),
```

The default value is **0**, all projects are fetched in a single call.

Allow project members to view project settings

By default, if the **Only Owners and Administrators can edit the project** checkbox is selected for a project, only the project owners and administrators can view the project **Settings** tab.

When **allow_view_settings** is set to **true** and the **Only Owners and Administrators can edit the project** checkbox is selected for a project, project members that are not owners or administrators can view a read-only version of the project **Settings** tab. The project **Automated Tests** and **Automated Deployment** details are hidden from the project members unless they are an owner or an administrator. This enables project members to check the project settings but not change them.

Add or update the following configuration block to the `SWARM_ROOT/data/config.php` file, at the same level as the `p4` entry:

```
<?php
    // this block should be a peer of 'p4'
    'projects' => array(
        'allow_view_settings' => true, // defaults to false
    ),
```

The default value is **false**, project members cannot see the project settings page if the project is set to **Only Owners and Administrators can edit the project**.

Redis server

Swarm requires Redis to manage its caches and by default Swarm uses its own Redis server on the Swarm machine. Swarm caches data from the Helix Core Server to improve the performance of common searches in Swarm and to reduce the load on the Helix Core Server.

Redis configuration:

- "Swarm connection to Redis " below: `SWARM_ROOT/data/config.php`.
- "Redis server configuration file" on page 676: `/opt/perforce/etc/redis-server.conf`.

Tip

- If required, you can use your own Redis server instead of the Swarm Redis server. For instructions on how to configure Swarm to use your Redis server, see ["Use your own Redis server" on page 235](#).
- If users and groups are added, edited, or deleted in the Helix Core Server while the Swarm server is shutdown, the Swarm user and group Redis caches will be out of sync. After the Swarm server has restarted, run a curl request to verify that the Redis caches are in sync with the Helix Core Server. For the Redis cache verification steps, see ["Manually verify the Redis caches" on page 677](#).
- When Swarm starts it verifies the Redis cache, during this time you cannot log in to Swarm. The time taken to verify the Redis cache depends on the number of users, groups, and projects Swarm has. Start-up time can be improved by persisting the memory cache. You can persist the memory cache by disabling background saves and enabling append saves in the `redis-server.conf` file, see ["Redis server configuration file" on page 676](#).

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

Swarm connection to Redis

The `redis` block of the `SWARM_ROOT/data/config.php` file contains the `password`, `namespace`, `host` and `port` details of the Swarm Redis server:

```
<?php
// this block should be a peer of 'p4'
'redis' => array(
    'options' => array(
        'password' => null, // Defaults to null
        'namespace' => 'Swarm',
        'server' => array(
            'host' => 'localhost', // Defaults to 'localhost' or
enter your Redis server hostname
```

```

        'port' => '7379', // Defaults to '7379' or enter your
Redis server port
    ),
),
    'items_batch_size' => 100000,
    'check_integrity' => '03:00', // Defaults to '03:00' Use one
of the following two formats:
                                // 1) The time of day that the
integrity check starts each day. Set in 24 hour format with leading
zeros and a : separator
                                // 2) The number of seconds
between each integrity check. Set as a positive integer. Specify '0'
to disable the integrity check.
    'population_lock_timeout' => 300, // Timeout for initial cache
population. Defaults to 300 seconds.
),

```

Configurables:

- **password:** Redis server password. Defaults to **null** and should be left at default if using the Swarm Redis server.
- **namespace:** the prefix used for key values in the Redis cache. Defaults to **Swarm** and should be left at default if using the Swarm Redis server.

Note

If you have multiple Swarm instances running against a single Redis server, each Swarm server must use a different Redis **namespace**. This enables the cache data for the individual Swarm instances to be identified. The **namespace** is limited to ≤ 128 characters.

If one or more of your Swarm instances is connected to multiple Helix Core Servers, the Redis **namespace** includes the **server label** and the character limit is reduced to ≤ 127 characters, see ["Multiple Helix Core Server instances" on page 646](#).

- **host:** Redis server hostname. Defaults to **localhost** and should be left at default if using the Swarm Redis server.
- **port:** Redis server port number. Defaults to **7379**. Swarm uses port **7379** as its default to avoid clashing with other Redis servers that might be on your network. It should be left at default if using the Swarm Redis server. The default port for a non-Swarm Redis server is **6379**.
- **items_batch_size:** Maximum number of key/value pairs allowed in an **mset** call to Redis. Sets exceeding this will be batched according to this maximum for efficiency. Defaults to **100000**.

Note

The default value of **100000** was chosen to strike a balance between efficiency and project data complexity. This value should not normally need to be changed, contact support before making a change to this value.

- **check_integrity**: In some circumstances, such as when changes are made in the Helix Core Server when Swarm is down or if errors occur during updates, the Redis cache can get out of sync with the Helix Core Server. Swarm can run a regular integrity check to make sure that the Redis caches and Helix Core Server are in sync. If an integrity check finds an out of sync cache file, Swarm automatically updates the data in that cache.

The **check_integrity** configurable specifies when the Redis cache integrity check is run. Set as a specific time of day (24 hour format with leading zeros) or a number of seconds (positive integer) between checks. Disable the integrity check with **'0'**. Defaults to **'03:00'**.

- **population_lock_timeout**: specifies the timeout, in seconds, for initial cache population. If you have a large Swarm system, increase this time if the initial cache population times out. Defaults to **300** seconds.

Redis server configuration file

Tip

To fine-tune your Redis server settings, make your changes in the Laminas Redis adapter. For information about the Laminas Redis adapter, see the [Laminas Redis Adapter documentation](#).

The Redis server configuration for Swarm is defined in `/opt/perforce/etc/redis-server.conf` and defaults to:

```
bind 127.0.0.1
port 7379
supervised auto
save ""
dir /opt/perforce/swarm/redis
```

Default values:**Tip**

- The default settings are shown below, the `redis-server.conf` file contains more detailed information about the Redis configuration for Swarm.
- On Swarm systems with a large number of users, groups, and projects, start-up time can be improved by persisting the memory cache. You can persist the memory cache by disabling background saves and enabling append saves, see the `redis-server.conf` file comments for detailed information.

- `bind 127.0.0.1` - Redis server IP address (loopback interface)
- `port 7379` - Redis server port number
- `supervised auto` - detects the use of `upstart` or `systemd` automatically to signal that the process is ready to use the supervisors
- `save ""` - background saves disabled, recommended.
- `dir /opt/perforce/swarm/redis` - the directory the Redis cache database is stored in.

Manually verify the Redis caches

Important

- Only *admin* and *super* users can verify the Redis cache.
 - On Swarm systems with a large number of users, groups, and projects verification of the caches can take some time and can impact performance.
- In normal operation, the Redis cache and Helix Core Server data should stay in sync and there is no need to verify the caches.
- If you have a large number of users, groups, and projects, we recommend that you verify the individual user and group caches rather than using **Verify All**.

In normal operation the Redis cache and Helix Core Server data should stay in sync and there is no need to verify the caches. However, in some circumstances, such as when changes are made in the Helix Core Server when Swarm is down or if errors occur during updates, the Redis cache can get out of sync with the Helix Core Server.

If you suspect that you have a cache issue, you can verify the cache manually from the System Information page.

For example, to verify the User cache:

1. Navigate to the **User id** dropdown menu.
2. Select **System Information**.
3. Click the **Cache Info** tab.
4. Click the **Verify User** button.

Swarm responds with a message to say that it has successfully verified the user cache.

If verification finds the cache is out of sync with the Helix Core Server, Swarm automatically updates the data in that cache.

For more information, see ["Verify buttons" on page 732](#).

Check progress of the verification request

To check the progress of the verification request:

1. From the **Cache Info** tab.
2. Click the **Refresh Status** button.

Review cleanup

This section describes the configurables that can be set for review cleanup.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

When a review is created in Swarm, it creates its own version of the changelist, leaving the user's own changelist untouched so as not to interfere with the user's ongoing work. Each time new work is submitted to the review, Swarm creates a new changelist so that there is a versioned history of the review.

When the review is finally committed it is Swarm's own changelist that is committed. Swarm's changelists are generally hidden from the users, but the user's changelist will remain open. By default it is necessary for the user to tidy up and remove this changelist themselves after the review has been completed.

There is an option to do this automatically when the review is committed. Configuration for this is expressed with a reviews block in the `SWARM_ROOT/data/config.php` file, similar to the following example:

```
<?php
    'reviews' => array(
        'cleanup' => array(
            'mode' => 'user', // auto - follow default, user -
present checkbox(with default)
            'default' => false, // clean up pending changelists
on commit
            'reopenFiles' => false, // re-open any opened files into
the default changelist
        ),
    ),
```

By default, this option is enabled but defaults to no clean up (so a user can select the option if they want to when they commit a review).

If the Helix Core Server user that Swarm is configured to use is a super user, then the user can clean up all user changelists associated with a review. If this is not the case, then the user who commits a review can only clean up changelists that are in their name.

By default, Swarm only cleans up changelists that are owned by the committing user. In the case where a user is committing a review that has been contributed to by other users, their changelists will not be cleaned up.

If you want to configure Swarm to clean up all changelists contributing to a review, regardless of whether they are owned by the committing user, you can do this by granting the Helix Core Server user that Swarm is configured as 'super' permissions/privileges (rather than just the 'admin' permissions/privileges that Swarm requires for its other operations).

Note

There is an API option that allows full cleanup to be executed with *super user* permissions using an external script. This removes the need for the Swarm Helix Core Server user to have super privileges. See [Pending review cleanup example](#) for details.

■ mode

If the mode is **user** then a checkbox is displayed when a review is committed, and the user has the option as to whether to clean up changelists or not.

If the mode is **auto**, then no checkbox is displayed, and all committed reviews will either always be tidied up, or never will be, depending on the value of **default**.

■ default

If set to **true**, changelists will always be cleaned up.

If set to **false**, changelists will never be cleaned up.

■ reopenFiles

If a changelist has files checked out (not shelved), then it cannot be deleted. Setting **reopenFiles** to **true** will mean that when a changelist is cleaned up, any opened files will first be moved to the default changelist, allowing the changelist to be removed.

If set to **false**, then a changelist with checked out files will not be cleaned up.

If users normally revert their files after shelving them, then this option may not be needed. If set to **true** then it may result in files appearing in the user's default changelist unexpectedly.

The review cleanup feature is designed to help users keep their workspaces clean, and prevent the proliferation of unwanted changelists after reviews have been approved and committed.

The following caveats should be kept in mind when using this feature:

- The changelists created by Swarm itself will not be touched by this process. There will always be a record of the review history that is kept.
- If the user who commits a review is normally different to the user that did the work, then unless Swarm runs as a super user then many changelists will not be cleaned up.

- If the committer created some of the changelists, then those changelists will be removed. However, changelists created by other users will not be removed.
- If the user has shelved files into changelists without reverting them, then they will still remain in the user's local workspace, and will need to be cleaned up manually.

Review keyword

This section describes how to configure keywords, and how to use keywords in a changelist description to create a review, and add a changelist to a review.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

Review keyword configuration

The keyword can be configured with a regular expression so that most any keyword syntax can be used. If you choose to customize the review keyword, take care to choose syntax and terminology that is unlikely to occur in a changelist description, to avoid unexpected Swarm activity.

To configure the review keyword, add the following block to the `SWARM_ROOT/data/config.php` file:

```
<?php
// this block should be a peer of 'p4'
'reviews' => array(
    'patterns' => array(
        'octothorpe' => array( // #review or #review-1234 with
surrounding whitespace/eol
            'regex' => '/(?P<pre>(?:\s|^)\(?\)\#
(?P<keyword>review|append|replace) (?::-(?P<id>[0-9]+))?(?P<post>
[.,!?:;])* (?:\s|$))/i',
            'spec' => '%pre%##keyword%-%id%%post%',
            'insert' => "%description%\n\n#review-%id%",
            'strip' => '/^\s*\#(review|append|replace) (-[0-9]+)?
(\s+|$) | (\s+|^)\#(review|append|replace) (-[0-9]+)?\s*$/i',
        ),
        'leading-square' => array( // [review] or [review-
```

```

1234] at start
    'regex'    => '/^(?P<pre>\s*)\
[ (?P<keyword>review|append|replace) (?::-(?P<id>[0-9]+)) ?\]
(?P<post>\s*)/i',
    'spec'     => '%pre%[%keyword%-%id%]%post%',
    ),
    'trailing-square' => array(      // [review] or [review-
1234] at end
    'regex'    => '/(?P<pre>\s*)\
[ (?P<keyword>review|append|replace) (?::-(?P<id>[0-9]+)) ?\]
(?P<post>\s*)?$/i',
    'spec'     => '%pre%[%keyword%-%id%]%post%',
    ),
    ),
    ),
    ),

```

Multiple patterns can be specified; the first successful match is used and none of the other patterns are evaluated.

The keyword types are grouped under their identifiers. In each group, the **regex** item specifies the regular expression to be used to identify the review keyword in the changelist description. The **spec** item is used when the review keyword needs to be updated.

Note
The use of named capture groups in the **regex**, for example **(?<pre>\s*)**. The values captured during regex matching are used to replace any identically named placeholder values in the spec item that are surrounded by percent % characters. In the example configuration above, the pre and post capture groups and placeholders maintain any whitespace surrounding the review keyword.

The use of named capture groups in the **regex**, for example **(?<pre>\s*)**. The values captured during regex matching are used to replace any identically named placeholder values in the spec item that are surrounded by percent % characters. In the example configuration above, the pre and post capture groups and placeholders maintain any whitespace surrounding the review keyword.

For **octothorpe** (or "hashtag") review keywords, these can appear anywhere in the changelist description. The **strip** item is used to ensure that the keyword is removed from the review description if it appears at the start or end of the changelist description. The **insert** item is currently not used; it is included here to prevent future upgrade issues. The intended use case is when a review is started and the changelist does not already contain a review keyword, the insert item would be used to add the review keyword to the changelist description.

For more information on named capture groups in PHP, see: <http://www.regular-expressions.info/named.html>

Create a review

By default, including the keyword `#review` within a changelist description (separated from other text with whitespace, or on a separate line) tells Swarm that a review should begin when the changelist is shelved or committed.

Once a review has begun, Swarm adds the review identifier to the `#review` keyword, for example `#review-1234`. This tells Swarm which review should be updated whenever the original changelist is re-shelved or committed.

Note

- If you create a pre-commit review and a user appends a changelist to your review, the default add mode of your review is changed from replace to append, see "[Add a changelist to a review](#)" below for details.
- Swarm acts on the first valid keyword it finds in the changelist description, Swarm then ignores any further valid keywords it finds in the description.
- Swarm can also accept `[review]` at the start or end of the changelist description, but this form of review keyword is now deprecated and is likely to be removed in a future version of Swarm.

Add a changelist to a review

Once a review has been started you can add a changelist to the review. It can be useful to add changelists to an existing review. For example, if follow up changes are made to files in a review or if you need to group a number of changelists under a single review.

Note

The changelist must not be part of another review, if it is Swarm will reject it.

Note

Swarm acts on the first valid keyword it finds in the changelist description, Swarm then ignores any further valid keywords it finds in the description.

By default, the `#append-`, `#replace-`, and `#review-` keywords along with the *review identifier* (separated from other text with whitespace, or on a separate line) can be used in a changelist description to add changelists to an existing review. The options available depend on whether the review is pre-commit or post-commit:

■ Pre-commit reviews:

- `#append-` + *review identifier*. When you add a changelist to a review with the append option, the files in the changelist are appended to the existing files in the review. This changes the default add mode of the review to append.

- **#replace-** + *review identifier*. When you add a changelist to a review with the replace option, all of the files in the review are replaced with the files in the changelist you are adding to the review. This changes the default add mode of the review to replace.

Note

If you replace a pre-commit review with a committed changelist, the new version of the review will be a post-commit review.

- **#review-** + *review identifier*. When you add a changelist to a review with this option, the review's default add method is used. The default add method can be either replace or append and it is set by the most recent add method used on the review:
 - When a review is first created the default add mode for the review is replace.
 - If a changelist is added to the review using append, the default add mode for the review is set to append.
 - If a changelist is added to the review using replace, the default add mode for the review is set to replace.
 - If you do not know what the current default add mode of the review is, delete the **#review-** keyword and specify the add mode you want by using either **#append-** + *review identifier*, or **#replace-** + *review identifier*.

■ Post-commit reviews:

- **#review-** or **#replace-** + *review identifier*. When you add a changelist to a review with the replace changelist option, all of the files in the review are replaced with the files in the changelist you are adding to the review.

Note

If you replace a post-commit review with a pending changelist, the new version of the review will be a pre-commit review.

Tip

When the content of a review is changed, Swarm checks to see which branches are in the new version of the review:

■ If a new branch was added to the review:

- Default reviewers on the new branch are added to the review.
- Moderators from the added branch become moderators for the review alongside the existing moderators.
- **Only if workflow is enabled:** if the new branch is associated with a workflow, the [workflow is merged](#) with the existing workflow. The most restrictive workflow is used for the review.

- **If a branch is no longer part of the review:**

- Reviewers for the review are not changed.
- Moderators from the removed branch no longer moderate the review.
- **Only if workflow is enabled:** if the branch was associated with a workflow, the branch workflow is removed from the review.

Example: Appending a changelist to review 1234:

Add `#append-1234` to the changelist description.

- Original review contains the following files: A^{#1}, B^{#2}, C^{#1}, D^{#1}, and E^{#4}
- Changelist contains the following files: A^{#2}, C^{#3}, E^{#4} (marked for delete), and F^{#1}
- Appending the changelist to the original review generates a new version of the review that contains the following files: A^{#2}, B^{#2}, C^{#3}, D^{#1}, E^{#4} (marked for delete), and F^{#1}

Important

Committing a review with Swarm (recommended): Swarm automatically commits the files in the approved version of the review.

Committing a review outside of Swarm:

Before you commit the review:

1. Unshelve the review into the pending changelist associated with the review.
2. Reshelve the files in the pending changelist.
3. Commit the pending changelist.
4. This process ensures that all of the files in the approved version of the review are committed.

Workflow feature [enabled \(default\):](#)

Swarm can be configured to automatically check that the files being committed match the files in the approved version of the review by using the **On commit with a review** workflow rule. For information about the **On commit with a review** rule, see ["Workflow rules" on page 526](#).

Workflow feature [disabled:](#)

Swarm can be configured to automatically check that the files being committed match the files in the approved version of the review by using the [strict](#) trigger option.

Example: Replacing the files in a review 1234 with the files in a changelist:

Add `#replace-1234` to the changelist description.

- Original review contains the following files: A^{#1}, B^{#2}, C^{#1}, D^{#1}, and E^{#4}
- Changelist contains the following files: A^{#2}, C^{#3}, E^{#4} (marked for delete), and F^{#1}

- Replacing the original review with the changelist generates a new version of the review that contains the following files: A#2, C#3, E#4 (marked for delete), and F#1

Tip

When you replace a review with a changelist, the base revisions of the files in the new version of the review are the base revisions of the files in the replacement changelist.

Keyword workflow

For workflow examples using keywords, see ["Review creation and modification outside of Swarm" on page 483](#).

Remove mentions from review descriptions

This example shows how to use regular expressions in the `patterns` block to remove @mentions in changelist descriptions so they don't get added to the review description.

Tip

This method can be adapted to remove anything you want from the review description.

Add the following block to the `SWARM_ROOT/data/config.php` file:

```
<?php
// this block should be a peer of 'p4'
'reviews' => array(
    'patterns' => array(
        'mentions' => array(
            'regex' => '/(?P<mention>\s@\S+|@\S+(\$|\s))/',
            'spec' => '%mention%'
        ),
    ),
),
```

Example

Using the above pattern on the following changelist description:

```
Updated copyright date to 2021.
#review
@bruno
@mei
```

If the changelist is shelved, it becomes:

```
Updated copyright date to 2021.
#review-12345
@bruno
@mei
```

The review description would be:

Updated copyright date to 2021.

Reviews filter

This section describes the configurables that can be set for the reviews filter.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

Swarm review sorting on the ["Reviews list" on page 428](#) page is customized by using the following configuration block in the `SWARM_ROOT/data/config.php` file:

```
<?php
// this block should be a peer of 'p4'
'reviews' => array(
    'filters' => array(
        'filter-max' => 15,
        'result_sorting' => true,
        'date_field' => 'updated', // 'created' displays and sorts by
created date, 'updated' displays and sorts by last updated
    ),
),
```

filter-max

When you visit the ["Reviews list" on page 428](#) page, reviews are displayed incrementally **x** reviews at a time until the white space on the page is filled with reviews or there are no more reviews to display. This process continues as you scroll down the page. The **filter-max** configurable sets this number.

Default value is **15**.

result_sorting

Controls whether the **Result order** button is displayed on the "Reviews list" on page 428 page.

- Set `result_sorting` to `true` to display the **Result order** button on the "Reviews list" on page 428 page, this is the default setting.
- Set `result_sorting` to `false` to remove the **Result order** button from the "Reviews list" on page 428 page, reviews are sorted based on the `date_field` setting.

date_field

If `result_sorting` is set to `true` (default), the `date_field` sets the initial setting for the **Result order** button on the "Reviews list" on page 428 page.

If `result_sorting` is set to `false`, the `date_field` sets the order reviews are displayed in on the "Reviews list" on page 428 page.

- **created**: reviews are sorted by when they were created. This is the default setting.
- **updated**: reviews are sorted by when they were last updated.

When a user starts a new session the **Result order** button is set to the initial sort order set in `date_field`. The review sort order can be changed by the user from the **Result order** button. This setting is remembered even if the user navigates away from the "Reviews list" on page 428 page. When the user starts a new session, the **Result order** button is reset back to the setting in `date_field`.

Reviews

This section describes the configurables that can be set for reviews.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the "Reload Configuration button" on page 732.

Allow author change

When enabled, the author of a review can be changed. This is useful in the case where the original author is no longer available, and someone else needs to take over ownership.

To allow the author of a review to be changed, update the `SWARM_ROOT/data/config.php` file to include the following configuration item within the reviews block:

```
<?php
    // this block should be a peer of 'p4'
```

```
'reviews' => array(
    'allow_author_change' => true,
),
```

The default value is **false**.

Allow author obliterate review

When enabled, users can obliterate reviews that they have authored. For details about how to obliterate a review, see ["Obliterate Review" on page 668](#).

Important

By default, only users with *admin* and *super* user rights can obliterate a review. Perforce recommends that the **allow_author_obliterate** configurable is kept at its default value of **false**. An obliterated review cannot be reinstated, not even by Perforce Support.

To allow users to obliterate reviews that they have authored, update the [SWARM_ROOT/data/config.php](#) file to include the following configuration item within the reviews block:

```
<?php
// this block should be a peer of 'p4'
'reviews' => array(
    'allow_author_obliterate' => true,
),
```

The default value is **false**.

Disable approve for reviews with open tasks

When enabled, Swarm will not allow you to **Approve**, or **Approve and Commit** a review that has open tasks. To approve, or approve and commit a review with open tasks, you must address the tasks first and then set them to **Task Addressed**, or **Not a Task**, see ["Set a task to Task addressed or Not a task" on page 352](#) for details.

Important

If the open tasks on the review are archived, the review can then be approved, or approved and committed.

To prevent reviews being approved, or approved and committed when the review has open tasks, update the [SWARM_ROOT/data/config.php](#) file to include the following configuration item within the reviews block:

```
<?php
// this block should be a peer of 'p4'
```

```
'reviews' => array(
    'disable_approve_when_tasks_open' => true,
),
```

The default value is `false`, in this case a warning is displayed for the open tasks but the action is allowed.

Disable self-approval of reviews by authors

The Swarm 2015.2 release provides the ability to disable review approval by authors, even if they are moderators or administrators. This is useful for development workflows where review by others is of paramount importance.

To disable review approval by authors, update the `SWARM_ROOT/data/config.php` file to include the following configuration item within the reviews block:

```
<?php
// this block should be a peer of 'p4'
'reviews' => array(
    'disable_self_approve' => true,
),
```

The default value is `false`.

Disable tests on approve and commit

Important

- From Swarm 2020.1, the `disable_tests_on_approve_commit` configurable has been removed from Swarm. This functionality has been moved to workflows and is set on a per workflow per test basis. A test can be set to run **On Update** or **On Submit** for a specific workflow or globally by setting the test on the [global workflow](#).
- If the `disable_tests_on_approve_commit` configurable is in your `"swarm_root"` on [page 724](#)/`data/config.php` file, it is ignored.

Protected end states

By default, the content of a review can be updated no matter what state the review is in. Reviews can be protected from automatic content change if they are in a specified state by using the `end_states` configurable.

Tip

When a review is in a protected end state, it can still be updated manually by a user from the Swarm UI.

Workflow:

- **If workflow is enabled** (default): the `end_states` configurable sets the protected review states. A combination of the [On update of a review in an end state rule](#) and [workflows](#) determines which projects and branches have protected reviews.
- **If workflow is disabled**:: when the content of a review is updated, the `end_states` array is checked to see if the review can be updated. If the review state matches a state in the `end_states` array the submit is rejected.

To set the protected review states, update the `SWARM_ROOT/data/config.php` file to include the following configuration item within the reviews block. In this example the **Archived**, **Rejected**, and **Approved : Committed** states are specified:

```
<?php
// this block should be a peer of 'p4'
'reviews' => array(
    'end_states' => array('archived','rejected',
'approved:commit'),
),
```

Valid review states for the array:

- **archived**: the review state is **Archived**.
- **rejected**: the review state is **Rejected**.
- **approved:commit**: the review has been approved and committed, the review state is **Approved : Commit**.
- **approved**: the review state is **Approved**.

By default the array is empty, review content can be updated no matter what state the review is in.

Expand all file limit

The review page has an **Expand All** button that opens all the files within that review. If the number of files is large, clicking the button might affect performance.

The `expand_all_file_limit` disables the button if the number of files in the review exceeds the given value. If the value is set to zero, the button is always enabled and can therefore open all the files.

To change the expand all file limit, update the `SWARM_ROOT/data/config.php` file to include the following configuration item within the reviews block:

```
<?php
// this block should be a peer of 'p4'
'reviews' => array(
    'expand_all_file_limit' => 10,
),
```

The default value if the option is not specified is **10**.

Expand group reviewers

By default, reviewer group members are not displayed in the *Individuals* area of the reviews page when they interact with a review (vote, comment, update, commit, archive, etc.). This avoids overloading the *Individuals* area with individual avatars if you have large reviewer groups.

Note

An exception to this behavior is when a member of a reviewer group is also an individual *required reviewer*, in this case their avatar will be displayed in the *Individuals* area.

When `expand_group_reviewers` is set to `true`, reviewer group members are added to the *Individuals* area of the review page when they interact with the review (vote, comment, update, commit, archive, etc.). If you have large reviewer groups, this might affect performance.

To expand group reviewers, update the `SWARM_ROOT/data/config.php` file to include the following configuration item within the reviews block:

```
<?php
// this block should be a peer of 'p4'
'reviews' => array(
    'expand_group_reviewers' => true,
),
```

The default value is **false**.

Maximum initial reviewers and reviewer groups in secondary navigation bar

By default, a maximum of 6 individual reviewers and 6 reviewer groups are initially displayed in the secondary navigation pane on the Review page. If there are more individual reviewers or reviewer groups on the review, the `+ n more` link is displayed as required so that you can display all of the reviewers and reviewer groups.

For example, with the default value of 6:

- **8 individual reviewers on the review:** 6 displayed in the secondary navigation bar and the remainder are available on the + 2 **more** link
- **9 reviewer groups on the review:** 6 reviewer groups are displayed in the secondary navigation bar and the remainder are available on the + 3 **more** link

To edit the maximum number of reviewers and reviewer groups initially displayed in the secondary navigation pane, update the `SWARM_ROOT/data/config.php` file to include the following configuration item within the reviews block. For example, to set the maximum number of reviewers and reviewer groups to 10:

```
<?php
    // this block should be a peer of 'p4'
    'reviews' => array(
        'max_secondary_navigation_items' => 10, // defaults to 6
        reviewers
    ),
```

The default value if the option is not specified is **6**.

Moderator behavior when a review spans multiple branches

By default, when a review spans multiple project branches that have different moderators, only one moderator from any one of the branches needs to approve the review. Set the configurable to **each** to require one moderator from each branch approve the review. If a moderator belongs to more than one of the branches spanned by the review, their approval will count for each of the branches they belong to.

To require one moderator from each branch to approve a review, edit the `SWARM_ROOT/data/config.php` file and set the `moderator_approval` configurable to **each** in the **reviews** section of the codeblock:

```
<?php
    // this block should be a peer of 'p4'
    'reviews' => array(
        'moderator_approval' => 'each', // 'any|each'
    ),
```

The default value is **any**.

More context lines

The **Review display** and **Changelist display** pages have **Show More Lines above** and **Show More Lines Below** buttons on the **file diff view** when a file is expanded. The buttons are used to display more lines of context for the file being viewed. By default 10 extra lines are displayed.

To change the number of lines displayed when the buttons are clicked, update the [SWARM_ROOT/data/config.php](#) file to include the following configuration item within the reviews block:

```
<?php
    // this block should be a peer of 'p4'
    'reviews' => array(
        'more_context_lines' =>15, // defaults to 10 lines
    ),
```

The default value if the option is not specified is **10**.

Process shelf file delete when

By default, when you delete files from a shelved changelist, the files are not removed from the associated review.

When you have a review state set in the `process_shelf_delete_when` array, Swarm will automatically carry out the following steps when a file is deleted from a shelf:

1. Check to see if the shelf is associated with a review.
2. The deleted file is only removed from the review if the review state matches the review state in the `process_shelf_delete_when` array.

Tip

- More than one review state can be specified in the `process_shelf_delete_when` array.
- When files are removed from a review, any votes on the review become stale. The vote counts are reset, and the vote indicators are muted.
- If `process_shelf_delete_when` is configured and the review is in a state that is specified in the array: When the review contains a file that is a common to multiple changelists that are associated with the review, if the owner of any of the associated pending changelists deletes the common file from their shelf, that file will be removed from the review, and a new version of the review will be created. This is true even if the file is not the most recent version in the review.

Important

Required for `process_shelf_delete_when`:

- A supported version of Helix Core Server (standard maintenance), see ["Helix Core Server requirements" on page 106](#).
- You must have the `swarm.shelvedel shelve-delete` trigger line in the Helix Core Server trigger table, see ["Installing Triggers" on page 204](#).
- You must satisfy the Swarm trigger dependencies, see ["Trigger dependencies" on page 108](#).

Important

Do not use the Swarm user that is configured in the [Swarm configuration file](#) when deleting shelves, or deleting files from shelves. The Swarm logic processes the *shelf-delete* trigger event, if the event is invoked by the Swarm user it is rejected. The delete operations will fail.

To remove files from a review when they have been deleted from an associated shelved changelist and the review is in a specified state, update the `SWARM_ROOT/data/config.php` file to include the following configuration item within the reviews block. In this example the **Needs Review** and **Needs Revision** states are specified:

```
<?php
    // this block should be a peer of 'p4'
    'reviews' => array(
        'process_shelf_delete_when' => array('needsReview',
        'needsRevision'),
    ),
```

Valid review states for the array:

- **needsReview**: the review state is **Needs Review**.
- **needsRevision**: the review state is **Needs Revision**.
- **rejected**: the review state is **Rejected**.
- **approved**: the review state is **Approved**.

Leave the array empty to disable this feature, see below. This is the default setting for the array.

```
<?php
    // this block should be a peer of 'p4'
    'reviews' => array(
        'process_shelf_delete_when' => array(),
    ),
```

Note

If a review is **Approved and Committed** or **Archived**, files will not be removed from the review when they are deleted in an associated shelf. This is true even if the **approved:commit** or **archived** states are in the **process_shelf_delete_when** array.

Synchronize review description

By enabling the synchronization of review descriptions, it becomes possible to update the description of a review by updating the description of a changelist associated with the review. Whenever an associated changelist is saved, the text of the review will be updated to match.

Note

The update is not triggered if:

- another changelist is attached to a review.
- files are updated in a changelist attached to a review
- a user selects **Update pending changelist** when editing a review description.

To enable synchronize review descriptions, update the `SWARM_ROOT/data/config.php` file to include the following configuration item within the reviews block:

```
<?php
    // this block should be a peer of 'p4'
    'reviews' => array(
        'sync_descriptions' => true,
    ),
```

The default value is `false`.

Review complexity

Review complexity is displayed for each review in the **Complexity** column on the "Reviews list" on [page 428](#) to help users to quickly see how complicated each review in the list is. By default, review complexity is based on the number of lines changed in the files in a review and simply displayed as high, medium, or low for each review. The values for high and low complexity are set using the `high` and `low` configurables.

Tip

- Review complexity is only calculated for a review when the review is updated and the file content has changed.
- Review complexity is only stored for the current version of a review.
- If you change the `high` or `low` complexity value, complexity is only recalculated for a review when the review is updated and the file content has changed.

To change the complexity values, update the `SWARM_ROOT/data/config.php` file with item within the reviews block:

```
<?php
    // this block should be a peer of 'p4'
    'reviews' => array(
        'statistics' => array(
            'complexity' => array(
                'calculation' => 'default', // 'default|disabled'
```

```
        'high' => 300,  
        'low' => 30  
    ),  
),  
,
```

- **calculation**

- **default** the complexity is based on the number of lines changed in the files in the review. For the purposes of complexity, a new file is counted as the number of lines in the file. Default value is **default**.
- **disabled** set to **disabled** to disable the complexity calculation.
- **high** set the number (integer) of line changes that are considered high for your organization. Default value is **300**.
- **low** set the number (integer) of line changes that are considered low for your organization. Default value is **30**.

By default, complexity is calculated as:

- **High:** number of lines changed in files in the review \geq **high**.
- **Medium:** number of lines changed in files in the review $<$ **high** and $>$ **low**.
- **Low:** number of lines changed in files in the review \leq **low**.

Maximum files limit

When creating or updating a review you can set a maximum file limit. This limits the maximum number of files that can be added to a review. If a request exceeds the maximum file number limit then a review is not created or updated.

By default, **max_files** configurable is set to 0. To edit the maximum file limit, update the **max_files** configurable under the reviews section in the `SWARM_ROOT/data/config.php` file.

Tip

Do not set **max_files** configurable to a low number as this will allow commits to bypass workflow and moderators that are blocking approval. It will not bypass review association. A commit must meet the workflow rule of being associated to a review.

Search

Swarm's search feature combines user, group, project, and file path searching. The standard Swarm searches can be extended to search file content and changelist description by using the optional Helix Core Search API.

You can download the Helix Core Search API installer from the Helix Search [download page](#) on the Perforce website. For more information, see the [Helix Core Search Developer Guide](#) and the Helix Search [release notes](#).

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the "Reload Configuration button" on page 732.

Configure Swarm search with the following configuration block in the [SWARM_ROOT/data/config.php](#) file:

```
<?php
    // this block should be a peer of 'p4'
    'search' => array(
        'maxlocktime'      => 5000, // 5 seconds, in milliseconds
        'p4_search_host'   => 'http://myhelixsearch.mydomain.com:4567',
    // optional URL to the Helix Search host
    ),
```

maxlocktime

The `maxlocktime` key specifies the maximum amount of time, in milliseconds, that any table within the Helix Core Server should be locked while performing `fstat` command searching. Increasing this value might allow better search results at the expense of potentially blocking other queries on the Helix Core Server. Decreasing this value impacts the Helix Core Server less, but may be insufficient for returning the desired search results.

p4_search_host

The `p4_search_host` configurable specifies the URL of your Helix Search server. When configured, Swarm issues API calls to Helix Search to take advantage of its file content and description indexing.

Comment attachment thumbnail and blur generation

Swarm generates the blur and thumbnail images for Swarm comment attachments and stores them in the [comment attachment depot](#).

To prevent Swarm clashing with Helix Search thumbnail and blur generation, set the Helix Core Server protections for the Swarm comment attachments depot directory so only the **swarm** user has access to it.

Tip

Swarm 2022.1 Patch 1 and later: By default, the Swarm package post-installation configuration script offers to create the comment attachment depot directory as `// .swarm` and set the protections so only the **swarm** user has access to it.

For example, if the comment attachment depot directory is `// .swarm`, set Helix Core Server protections for `// .swarm` to:

```
list user * * -//.swarm/...
admin user swarm * //.swarm/...
super user super * //.swarm/...
```

This prevents Helix Search creating thumbnails and blurs for files stored in the comment attachment depot directory.

More information

- For instructions on configuring the Swarm comment attachments depot directory, see ["Comment attachments" on page 580](#).
- For instructions on setting Helix Core Server protections, see the [Setting protections with p4 protect](#) section of the [Helix Core Server Administrator Guide](#).

Security

There are many strategies for securing a Swarm installation. This section provides guidance on security features Swarm controls, and recommendations for several areas for the system hosting Swarm.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

Require login

Important

Prior to Swarm's 2016.1 release, `require_login` defaulted to `false`. For 2016.1 and later releases, the default is `true`.

By default, Swarm prevents anonymous users from viewing any Helix Core Server resources; users must login to see commits, reviews, etc.

Swarm can be configured to allow anonymous users to access any readable resources (creating or editing resources by anonymous users is not permitted). Add the following configuration block to the `SWARM_ROOT/data/config.php` file, at the same level as the p4 entry:

```
<?php
    // this block should be a peer of 'p4'
    'security' => array(
        'require_login' => false, // defaults to true
    ),
```

There is one exception: the `/queue/worker` endpoint is available to any user.

Note

Service and operator users are not permitted to login. For more information on these user types, see [User types](#) in *Helix Core Server Administrator Guide*.

Prevent login

When your Helix Core Server has users that should not be able to log in to Swarm, for example *service* users involved with Helix Core Server replicas, the `prevent_login` configuration item can be used to prevent successful authentication.

Add or update the following configuration block to the `SWARM_ROOT/data/config.php` file, at the same level as the p4 entry:

```
<?php
    // this block should be a peer of 'p4'
    'security' => array(
        'prevent_login' => array(
            'service_user1',
            'service_user2',
        ),
    ),
```

`prevent_login` defaults to `array()`, which means no users in your Helix Core Server are prevented from logging into Swarm.

For more information, see [Service users](#) in *Helix Core Server Administrator Guide*.

Auto-create new users

Swarm supports the automatic-creation of new users in Perforce on login if the Helix Core Server is configured to allow it and the user exists in LDAP. This function is configured in the Helix Core Server and no Swarm configuration is required. For instructions on how to configure Helix Core Server to automatically create new users in Perforce on login, see `auth.ldap.userautocreate` and `auth.default.method` in the [Defining LDAP-related configurables](#) section of the *Helix Core Server Administrator Guide*.

Sessions

Swarm manages logged-in sessions using cookies in the browser, and PHP session storage on the server. Swarm uses reasonable defaults for the cookie and session lifetimes (measured in seconds); when the lifetime is exceeded users need to login again. To specify session lifetimes and garbage collection frequency, add the following configuration block to the `SWARM_ROOT/data/config.php` file, at the same level as the `p4` entry:

```
<?php
    // this block should be a peer of 'p4'
    'session' => array(
        'cookie_lifetime'           => 0, // lifetime in seconds,
default value is 0=expire when browser closed
        'remembered_cookie_lifetime' => 60*60*24*30, // lifetime in
seconds, default value is 30 days
        'user_login_status_cache'   => 10, // Set in seconds, default
value is 10 seconds.

                                                // Set to 0 to
disable the cache and make Swarm

                                                // check the user
login status for every call to Helix Server.
        'gc_maxlifetime'           => 60*60*24*30, // lifetime in
seconds, default value is 30 days
        'gc_divisor'               => 100, // 100 user requests
    ),
```

■ `cookie_lifetime`

Optional: Limits the lifetime of session cookies (in seconds).

Default is `0`, which causes the session cookie to expire when the user's browser is closed.

Maximum value is `2147483647` seconds.

Note

When the **Remember Me** checkbox on the [login dialog](#) is checked, the `remembered_cookie_lifetime` value will be used for `cookie_lifetime`.

- `remembered_cookie_lifetime`

Optional: Limits the lifetime of session cookies when the **Remember Me** checkbox on the [login dialog](#) is checked.

Default is `60*60*24*30` seconds (30 days). You can enter the value in seconds if you prefer. For example, `2592000` seconds (30 days) is also valid.

Maximum value is `2147483647` seconds.

- `user_login_status_cache`

Swarm 2022.1 and later: Swarm caches the Helix Core Server login status of users. When Swarm makes any single request, it tests the user's login status many times. The default 10 second cache value is intended to reduce the number of checks performed, reducing the load on the Helix Core Server.

Default is `10` seconds.

Set to `0` to disable the cache and make Swarm check the login status for every call to the Helix Core Server.

Note

If a user requests a Swarm operation on the Helix Core Server when logged out of the Helix Core Server, but their login cache status shows them as logged in, the operation will fail. If you get too many reports of failed operations, reduce your cache setting.

- `gc_maxlifetime`

Optional: If a session is inactive for the specified number of seconds, the user is logged out. User sessions are stored in `SWARM_ROOT/data/sessions/`.

Default is `60*60*24*30` seconds (30 days). You can enter the value in seconds if you prefer. For example, `2592000` seconds (30 days) is also valid.

Maximum value is `2147483647` seconds.

Note

By default, the user's Perforce ticket expires after 12 hours, which also causes them to be logged out.

- `gc_divisor`

Optional: Sets how often garbage collection is run based on the number of user requests that are made. The setting range is 1 to 100. Garbage collection deletes user session files that are older than the `gc_maxlifetime` setting from `SWARM_ROOT/data/sessions/`.

- `gc_divisor = 100`: Garbage collection runs after every 100th user request. 100 is the default setting.
- `gc_divisor = 1`: Garbage collection runs after every user request.

Important

If your Swarm system has a large number of users, setting `gc_divisor` to a low number can result in performance issues.

X-Frame-Options header

By default, Swarm emits a `X-Frame-Options` HTTP header set to `SAMEORIGIN`. This prevents embedding of the Swarm interface into other web pages, which avoids *click-jacking* attacks.

If your deployment of Swarm needs to be integrated into another web interface, you can adjust the `X-Frame-Options` header by adjusting the `x_frame_options` item within the security configuration block, found in the `SWARM_ROOT/data/config.php` file. For example:

```
<?php
// this block should be a peer of 'p4'
'security' => array(
    'x_frame_options' => value,
),
),
```

Where value can be one of:

- `'SAMEORIGIN'` - Swarm can only be displayed in a frame hosted on the same domain.
- `'DENY'` - Swarm cannot be displayed in a frame.
- `'ALLOW-FROM URI'` - Swarm can only be displayed in a frame hosted on the specified URI.
- `false` - The `X-Frame-Options` header is not emitted, so Swarm can be embedded without restriction.

For more information on the `X-Frame-Options` header, see [this Mozilla Developer Network article](#).

For more information on click-jacking attacks, see this [Wikipedia article](#).

Disable commit

Swarm provides the ability to commit reviews within the Swarm interface. You may want to disable this capability to prevent reviews from being committed by someone other than the review's author. When disabled, the **Approve and Commit** (and **Commit** if the review is already approved) option is removed from the list of [states](#) available to a code review.

To disable commits, set `disable_commit` to `true` within the reviews item in the `SWARM_ROOT/data/config.php` file. For example:

```
<?php
// this block should be a peer of 'p4'
'reviews' => array(
    'disable_commit' => true,
),
),
```

Restricted Changes

The Helix Core Server provides two changelist types: `public` (the default), and `restricted`. Swarm honors restricted changelists by preventing access to the changelist, and any associated comments or activity related to the changelist.

If a user has *list*-level privileges to at least one file in the changelist, Swarm allows the user to see the changelist and any of the files they have permission to see.

To prevent unintended disclosures, email notifications for restricted changes are disabled by default. To enable email notifications for restricted changes, set `email_restricted_changes` to `true` within the security item in the `SWARM_ROOT/data/config.php` file. For example:

```
<?php
// this block should be a peer of 'p4'
'security' => array(
    'email_restricted_changes' => true,
),
),
```

Note

When `email_restricted_changes` is set to `true`, email notifications for restricted changes are sent to all interested parties with no permissions screening. These notifications might disclose sensitive information.

Swarm can only report on changes that the configured *admin*-level user has access to. When using restricted changes, we advise that you grant the *Swarmadmin*-level user access to the restricted files and set `require_login = true` to avoid leaking information to unauthenticated users.

Limit adding projects to administrators

Important

For Swarm 2016.1, the configuration item `add_project_admin_only` was moved from the **security** block to the **projects** block, and the item was renamed to `add_admin_only`. The functionality of this configuration item remains unchanged.

If you do not update your `SWARM_ROOT/data/config.php` configuration file, the old configuration for restricting project creation to administrators continues to work.

If you add the new configuration item `add_admin_only` to the **projects** block, it takes precedence over any remaining `add_project_admin_only` setting in the **security** block.

Limit adding projects to members of specific groups

Important

For Swarm 2016.1, the configuration item `add_project_groups` was moved from the **security** block to the **projects** block, and the item was renamed to `add_groups_only`. The functionality of this configuration item remains unchanged.

If you do not update your `SWARM_ROOT/data/config.php` configuration file, the old configuration for restricting project creation to specific groups continues to work.

If you add the new configuration item `add_groups_only` to the **projects** block, it takes precedence over any remaining `add_project_groups` setting in the **security** block.

IP address-based protections emulation

A Helix Core Server can be configured via *protections* to restrict access to a depot in a variety of ways, including by IP address. As Swarm is a web application acting as a client to the Helix Core Server, often with *admin*-level privileges, Swarm needs to emulate IP address-based restrictions. It does so by checking the user's IP address and applying any necessary restrictions during operations such as browsing files, viewing file content, viewing and adding comments on files.

Swarm also emulates proxy-based protections, in addition to regular IP-based protections emulation. However, Swarm does not detect whether it is connecting to a Helix Proxy or not; it merely attempts to emulate protections table entries that use proxy syntax.

IP address-based protections emulation is disabled by default. To enable IP address-based protections emulation for your Swarm installation, include the following configuration item in the `SWARM_ROOT/data/config.php` file:

```
<?php
    // this block should be a peer of 'p4'
    'security' => array(
```

```
'emulate_ip_protections' => true,
),
```

Tip

If Swarm is connected to multiple Helix Core Server instances, you have the following options for **emulate_ip_protections**:

- To apply a global **emulate_ip_protections** setting to all Helix Core Server instances that Swarm is connected to, set **emulate_ip_protections** in the **security** block as shown above.
- To apply an individual **emulate_ip_protections** setting to a Helix Core Server instance, set **emulate_ip_protections** in the **serverid** block for that server. This value overrides the global setting in the **security** block for the Helix Core Server instance.

For more information about configuring Swarm for multiple Helix Core Server instances, see ["Set up the Swarm configuration file for the Helix Core Servers" on page 647](#).

Known limitations

- Notification e-mails for reviews or commits include the list of affected files. Swarm cannot reliably know the IP address used to retrieve that e-mail, and makes no attempt to filter the files and their depot paths nor any details included in the description. However, when a user follows a link from the notification e-mail to a restricted resource, that access is denied.
- Swarm filters comments from activity streams, but any comments created prior to upgrading to the 2013.3 release cannot be filtered and may leak sensitive information.
- Swarm displays a comment count in code the review list, code reviews, jobs, and activity streams, but the count does not account for any comments that may be hidden from the user due to association with files the user is restricted from viewing.
- Should Swarm users connect to Swarm via a proxy or VPN, the protections will generally use the IP address of the proxy/VPN.
- When the user's IP address and Swarm's IP address both have restrictions applied, the user experiences the most constraining of the two IP address-based restrictions; Swarm cannot bypass restrictions applied to itself.
- Swarm performs a variety of operations with *admin*-level privileges, on behalf of a user. Even if the Helix Core Server has IP-based, or userid-based protections, installed to prevent access to some or most of its versioned data, Swarm typically does have access to this data. Therefore, *Swarm cannot guarantee that no information leakage will occur*, particularly when custom modules are in use, or Swarm source has been customized.

For more information, see [Authorizing Access](#) in *Helix Core Server Administrator Guide*.

Disable system info

Swarm provides a **System Information** page, available to users with *admin* or *super* privileges, which displays information about Helix Core Server that Swarm is configured to use, as well as PHP information and the Swarm log file.

While this information can be invaluable when communicating with Perforce support engineers, you may wish to prevent disclosure of any system information. The **System Information** page can be disabled for all users by adding the following configuration block to the [SWARM_ROOT/data/config.php](#) file, at the same level as the p4 entry:

```
<?php
    // this block should be a peer of 'p4'
    'security' => array(
        'disable_system_info' => true, // defaults to false
    ),
```

Once disabled, the **System Information** link disappears from the [About Swarm](#) dialog, and 403 errors are generated for any attempts to browse to the **System Information** page.

HTTP client options

Swarm permits configuration of options that are passed through to the underlying Laminas Framework HTTP client. These options can be used to specify SSL certificate locations, request timeouts, and more, and can be specified globally or per host.

Here is an example configuration:

```
<?php
    // this block should be a peer of 'p4'
    'http_client_options' => array(
        'timeout'          => 10, // default value is 10 seconds

        // path to the SSL certificate directory
        'sslcapath'        => '',

        // the path to a PEM-encoded SSL certificate
        'sslcert'          => '',

        // the path to Certificate Authority file
        'sslcafile'        => '',

        // the passphrase for the SSL certificate file
        'sslpassphrase'    => '',
```

```
// optional, per-host overrides;
// host as key, array of options as value
'hosts'      => array(
    'jira.example.com' => array(
        'sslcafile'      => '/path/to/certs/jira.pem',
        'sslpassphrase' => 'keep my JIRA secure',
        'timeout'        => 15,
    ),
    'jenkins.example.com' => array(
        'sslcafile'      => '/path/to/certs/jenkins.pem',
        'sslpassphrase' => 'keep my jenkins very secure',
        'timeout'        => 15,
    ),
),
),
```

See the [Laminas Framework Socket Adapter documentation](#) for more information.

Note

Swarm supports the Laminas component versions in the **LICENSE.txt** file, features and functions in the Laminas documentation that were introduced in later versions of Laminas will not work with Swarm. The **LICENSE.txt** file is in the **readme** folder of your Swarm installation.

Warning

While it is possible to use a self-signed SSL certificate, adding the configuration to do so disables certificate validity checks, making connections to the configured host less secure. **We strongly recommend against using this configuration option.**

However, if you need to configure continuous integration, deployment, or JIRA connections and those connections must use a self-signed SSL certificate, set the **sslallowselfsigned** item to **true** for the specific host that needs it, as in the following example:

```
<?php
// this block should be a peer of 'p4'
'http_client_options' => array(
    'hosts'      => array(
        'jira.example.com' => array(
            'sslallowselfsigned' => true,
        ),
    ),
),
```

```
    ),
  ),
```

Strict HTTPS

To improve the security when users work with Swarm, particularly if they need to do so outside of your network, Swarm provides a mechanism that tries to force web browsers to use HTTPS. When enabled, Swarm's behavior changes in the following ways:

- HTTP requests to Swarm include a meta-refresh to the HTTPS version. If a load balancer handles encryption before requests reach Swarm, the meta-refresh should be disabled. See [below](#).
- A strict transport security header is included for all requests, which pins the browser to using HTTPS for your Swarm installation for 30 days.
- All qualified URLs that Swarm produces use HTTPS for the scheme.
- Cookies are flagged as HTTPS-only.

Here is an example of how to enable strict HTTPS:

```
<?php
    // this block should be a peer of 'p4'
    'security' => array(
        'https_strict'           => true,
        'https_strict_redirect' => true, // optional; set false to
avoid meta-refresh
        'https_port'            => null, // optional; specify if HTTPS
is
                                   // configured on a non-
standard port
    ),
```

When the `https_strict_redirect` item is set to `false`, Swarm does not add a meta-refresh for HTTP clients. This prevents an endless redirect when a load balancer in front of Swarm applies HTTPS to the client-to-load balancer connection, but not the load balancer-to-Swarm connection.

Apache security

There are several Apache configuration changes that can improve security for Swarm:

■ Disable identification

By default, each Apache response to a web request includes a list of tokens identifying Apache and its version, along with any installed modules and their versions. Also, Apache can add a signature line to each response it generates that includes similar information. By itself, this identification information is not a security risk, but it helps would-be attackers select attacks that could be successful.

To disable Apache identification, add the following two lines to your Apache configuration:

```
ServerSignature Off
ServerTokens ProductOnly
```

■ Disable TRACE requests

TRACE requests cause Apache to respond with all of the information it has received, which is useful in a debugging environment. **TRACE** can be tricked into divulging cookie information, which could compromise the credentials being used to login to Swarm.

To disable **TRACE** requests, add the following line to your Apache configuration:

```
TraceEnable off
```

■ Update SSL configuration

Swarm works correctly with an SSL-enabled Apache. Several attacks on common SSL configurations have been published recently. We recommend that you update your Apache configuration with the following lines:

```
<IfModule mod_ssl.c>
    SSLHonorCipherOrder On
    SSLCipherSuite ECDHE-RSA-AES128-SHA256:AES128-GCM-
    SHA256:RC4:HIGH:!MD5:!aNULL:!EDH
    SSLCompression Off
</IfModule>
```

PHP security

There are several PHP configuration changes that can improve security for Swarm:

■ Disable identification

By default, PHP provides information to Apache that identifies that it is participating in a web request, including its version.

To disable PHP identification, edit your system's `php.ini` file and change the line setting `expose_php` to:

```
expose_php = Off
```

■ Remove scripts containing `phpinfo()`

During module development or other debugging, you may need to call `phpinfo()`, which displays PHP's active configuration, compilation details, included modules and their configuration. Typically, you would add a script to Swarm's public directory containing:

```
<?php phpinfo() ?>
```

Any such scripts should be removed from a production instance of Swarm.

proxy_mode

By default, Swarm works in proxy mode, passing the end-user's browser IP address to the Helix Core Server. The Helix Core Server checks the IP address against its IP protection table to work out what permissions the end-user has.

Set the Swarm **proxy_mode** in the **p4** configuration block of the [SWARM_ROOT/data/config.php](#) file:

```
<?php
    'p4' => array(
        'proxy_mode' => true, // defaults to true
    ),
```

- **true**: Swarm passes the end-user's browser IP address to the Helix Core Server allowing the Helix Core Server to work out the what permissions the end-user has. This is the default setting.
- **false**: the Swarm server IP address is passed to the Helix Core Server. The Helix Core Server checks the Swarm IP address against its protections table to work out the permissions the end-user has. The result is that all of the Swarm users will have the same permissions.

Tip

If Swarm is connected to multiple Helix Core Server instances, you have the following options for **proxy_mode**:

- To apply a global **proxy_mode** setting to all Helix Core Server instances that Swarm is connected to, set **proxy_mode** in the **p4** block as shown above.
- To apply an individual **proxy_mode** setting to a Helix Core Server instance, set **proxy_mode** in the **serverid** block for that server. This value overrides the global setting in the **p4** block for the Helix Core Server instance.

For more information about configuring Swarm for multiple Helix Core Server instances, see ["Set up the Swarm configuration file for the Helix Core Servers" on page 647](#).

Short links

Short links work with your Swarm installation's current hostname, but you have the option of registering/configuring an even shorter hostname to make shareable file/directory links as short as possible.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the **"Reload Configuration button"** on page 732.

1. Register a short domain name, or if you control your own DNS server, a short domain name for your network.
2. Point the short domain name at your Swarm host.
3. Edit the `SWARM_ROOT/data/config.php` file and add the following configuration block:

```
<?php
    // this block should be a peer of 'p4'
    'short_links' => array(
        'hostname' => 'myho.st',
    ),
```

Replace `myho.st` with the short domain name you registered/configured.

If your Swarm is configured to use **HTTPS**, a **custom port**, a **sub-folder**, or any combination of these custom installation options, the short links configuration block should look like:

```
<?php
    // this block should be a peer of 'p4'
    'short_links' => array(
        'external_url' => 'https://myho.st:port/sub-folder',
    ),
```

Replace `myho.st` with the short domain name you have registered/configured.

If you have not configured Swarm to use HTTPS, replace `https://` with `http://`.

If you have configured Swarm to run on a custom port, replace `:port` with the correct custom port. Otherwise, remove `:port`.

If you have configured Swarm to run in a sub-folder, replace `/sub-folder` with the correct sub-folder name. Otherwise, remove `/sub-folder`.

Important

The `external_url` configuration item is only honored if you have also configured the `"external_url" on page 611` item within the `"Environment" on page 609` configuration item as well. Otherwise, Swarm could generate short links that cannot correctly link to their corresponding full URLs.

When `external_url` is configured, the `hostname` configuration item is ignored.

Single Sign-On PHP configuration

Important

Swarm now supports Helix Authentication Service (HAS) as a Single Sign-On (SSO) provider. This helps to simplify configuration and create a more robust SSO solution. The `SAML` configuration will be removed from Swarm in a later release.

If `SAML` is configured when the `"sso" on page 196` configurable is set to optional, the user can either log in to Swarm using the `Log in with SSO` button or the `Log in with credentials` button. If `SAML` is configured when the `"sso" on page 196` configurable is enabled, the user can only log in to Swarm using the `Log in with SSO` button.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the `"Reload Configuration button" on page 732`.

Helix Authentication Service

Important

- To use Swarm with a Helix Core Server that is configured to use the Helix Authentication Service, you must set the Swarm `sso` configurable to either `optional` or `enabled` in the `"Swarm configuration file" on page 194`.
- From Swarm 2021.1, the `sso_enabled` configurable is deprecated but remains supported. It is replaced with the more flexible `sso` configurable. If the `sso_enabled` configurable and `sso` configurable are both present in the `p4` configuration block, Swarm uses the `sso` configurable value.

The Helix Authentication Service acts as the SAML Identity Provider (IdP) for Swarm's Service Provider (SP) and enables authentication with your external IdP using the protocol the administrator has configured in the Helix Authentication Service.

This section describes how to configure SAML 2.0 in the Swarm `config.php` file to enable Swarm to authenticate with a Helix Core Server that is configured for the Helix Authentication Service. The SAML PHP configuration block in the `SWARM_ROOT/data/config.php` file configures the Helix Authentication Service and Service Provider (SP) connection details for Swarm. This enables Swarm to connect to your Helix Authentication Service and your Helix Authentication Service to connect to Swarm so that you can log in to Swarm using the IdP log in process.

Configuring Swarm for Helix Authentication Service:

- For an overview of the Helix Authentication Service, see [Overview of Helix Authentication Service](#) chapter in the *Helix Authentication Service Administrator Guide*.
- For instructions on configuring the Helix Authentication Service, see the [Configuring Helix Authentication Service](#) chapter in the *Helix Authentication Service Administrator Guide*.
- For instructions on configuring Swarm to use the Helix Authentication Service, see "Swarm SAML 2.0 settings" below.
- For a more in-depth explanation of Swarm configuration for the Helix Authentication Service, see the [Example Helix Swarm configuration](#) chapter in the *Helix Authentication Service Administrator Guide*.

Swarm SAML 2.0 settings

This section describes the minimum settings that you must enter to enable Swarm to connect to a Helix Core Server that is enabled for Helix Authentication Service. The `saml` configuration block must be added to the end of the `SWARM_ROOT/data/config.php` file as shown in the following example.

Tip

The exact content of your `saml` configuration block depends on your SAML configuration, you can add other configurables to the SAML PHP block if they are required by your Helix Core Server SAML configuration. For example the `x509cert` and `privateKey` for your SP (Swarm) might be in the certs folder so you would not need to have them specified in the `sp` part of your `saml` block. It is important to note however that the idp `x509cert` must always be in the `idp` part of your `saml` block.

For an overview of SAML 2.0 , see <https://github.com/onelogin/php-saml>.

Example SAML PHP configuration, follow the underlined links for more information about the configurables:

Warning

While the syntax of this example is correct, it includes configuration values that **cannot work**. Ensure that you adjust the configuration appropriately for your SAML configuration before using the `saml` block in testing or production.

```
<?php
// the saml block should be a peer of 'p4' located at the end of
```

```
// the Swarm configuration block in the config.php file
'saml' => array(
    // If your Helix Server trigger expects a message header so
    that it can
    // easily recognize SAML response messages, add the header
    text
    'header' => 'saml-response: ', // leave empty for no message
    header ''
    // Service Provider Data that we are deploying
    'sp' => array(
        // Identifier of the SP entity (must be a URI)
        'entityId' => '<urn:my_swarm:sp>',
        // Specifies info about where and how the AuthnResponse
        message MUST be
        // returned to the requester, in this case our SP.
        'assertionConsumerService' => array(
            // URL Location where the Response from the IdP will
            be returned, this is the Swarm URL and port
            'url' => '<[http[s]://]<swarm-host>[:<port>]>',
            ),
        // Usually x509cert and privateKey of the SP (Swarm) are
        provided by files placed in
        // the certs folder. These files must be named sp.crt and
        sp.key.
        // Optional: you can also provide them with the following
        parameters
        'x509cert' => '<my_sp_swarm_full_cert_string_including_
        the_BEGIN_CERTIFICATE_and_END_CERTIFICATE_parts>',
        'privateKey' => '<my_sp_swarm_private_key>',
        ),
        // Identity Provider Data that we want to connect to with our
        SP (Swarm)
        'idp' => array(
            // Identifier of the IdP entity (must be a URI)
```

```

        'entityId' => '<my_entityid_provided_by_the_idp>',
        // SSO endpoint info of the IdP. (Authentication Request
protocol)
        'singleSignOnService' => array(
            // URL Target of the IdP where the SP (Swarm) will
send the Authentication Request Message
            'url' => '<full_idp_URL_path_to_send_authentication_
request_message_to>',
        ),

        // The x509cert of the idp is provided by the following
x509cert parameter.
        // Do not add the privateKey parameter.
        // You must use the x509cert parameter, you must not add
the cert file the certs folder.
        'x509cert' => '<my_idp_full_cert_string_including_the_
BEGIN_CERTIFICATE_and_END_CERTIFICATE_parts>',
    ),
),

```

Note

The Swarm configuration file does not include PHP's standard closing tag (`?>`). This is intentional as it prevents unintentional whitespace from being introduced into Swarm's output, which would interfere with Swarm's ability to control HTTP headers. Debugging problems that result from unintentional whitespace can be challenging, since the resulting behavior and error messages often appear to be misleading.

header

Some Helix Core Server triggers need a header (prefix) added to SAML response messages so that the Helix Core Server can easily identify the messages. The header is set in the **header** value.

If a header is not required, set **header** to empty `''`.

```

<?php
    'saml' => array(
        // If your Helix Server trigger expects a message header so
that it can
        // easily recognize SAML response messages, add the header

```

```

text
    'header' => 'saml-response: ', // leave empty for no message
header ''
    ),

```

The default value is `'saml-response: '`.

sp

The **sp** (service provider) section defines the callback destination and identifier information for your IdP (Identity Provider). This tells your IdP how to connect to Swarm.

- **entityId**: this is the identifier your IdP knows your Swarm as. Set your Entity ID here and then use the same value in your IdP's configuration tool. When Swarm connects to your IdP the Entity ID is used to verify your connection. This must be a URI.
- **assertionConsumerService url**: enter your Swarm URL and port number. This sets your Swarm instance as the URL your IdP sends responses to. If you don't enter a port number, port 80 is used.

Important

Do not use `localhost` for the `url`.

- **x509cert** and **privateKey**: enter your Swarm instance security connection details.

```

<?php
    'saml' => array(
        // Service Provider Data that we are deploying
        'sp' => array(
            // Identifier of the SP entity (must be a URI)
            'entityId' => '<urn:my_swarm:sp>',
            // Specifies info about where and how the AuthnResponse
message MUST be
            // returned to the requester, in this case our SP.
            'assertionConsumerService' => array(
                // URL Location where the Response from the IdP will
be returned, this is the Swarm URL and port
                'url' => '<[http[s]://]<swarm-host>[:<port>]>',
            ),

            // Usually x509cert and privateKey of the SP (Swarm) are
provided by files placed in

```



```

        // the certs folder. These files must be named sp.crt and
        sp.key.

        // Optional: you can also provide them with the following
        parameters

        'x509cert' => '<my_sp_swarm_full_cert_string_including_
        the_BEGIN_CERTIFICATE_and_END_CERTIFICATE_parts>',
        'privateKey' => '<my_sp_swarm_private_key>',
    ),
),

```

idp

The Helix Authentication Service acts as the SAML Identity Provider (IdP) for Swarm's Service Provider (SP) and enables authentication with your external IdP using the protocol the administrator has chosen for the Helix Authentication Service. The **idp** (identity provider) section defines the IdP connection and security information, this tells Swarm how to connect to your Helix Authentication Service.

- **entityId**: enter your Entity ID, this is configured in your Helix Authentication Service. This enables Swarm to connect to the Helix Authentication Service. This must be a URI.
- **singleSignOnService url**: enter the URL Swarm sends authentication requests to. This is configured in your Helix Authentication Service.
- **x509cert**: enter your Helix Authentication Service security connection cert.

```

<?php
    'saml' => array(
        // Identity Provider Data that we want to connect to with our
        SP (Swarm)

        'idp' => array(
            // Identifier of the IdP entity (must be a URI)
            'entityId' => '<my_entityid_provided_by_the_idp>',
            // SSO endpoint info of the IdP. (Authentication Request
            protocol)

            'singleSignOnService' => array(
                // URL Target of the IdP where the SP (Swarm) will
                send the Authentication Request Message

                'url' => '<full_idp_URL_path_to_send_authentication_
                request_message_to>',
            ),
        ),
    ),

```

```

        // The x509cert of the idp is provided by the following
x509cert parameter.
        // Do not add the privateKey parameter.
        // You must use the x509cert parameter, you must not add
the cert file the certs folder.
        'x509cert' => '<my_idp_full_cert_string_including_the_
BEGIN_CERTIFICATE_and_END_CERTIFICATE_parts>',
    ),
),

```

Skip activity

Important

We recommend that before enabling the skip activity configurables in the **SWARM_ROOT/data/config.php** file, you must contact Perforce. For contact details, see ["Getting help" on page 1396](#).

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

Swarm activity streams are filtered to display events related to Reviews, Commits (changes), Comments, or Jobs. You can configure Swarm to skip an activity in the **SWARM_ROOT/data/config.php** file, as in the following example:

Note

- By default all skip activity configurables are set to false in the **config.php** file.
- When **skip_groups** and **skip_projects** are set to true, this prevents all associated groups and project members from receiving activity.
- When **skip_groups**, **skip_projects** and **skip_followers** are set to true, this prevents all associated groups, project members, and followers from receiving activity.

```

<?php
// this block should be a peer of 'p4'
'activity' => [
    'skip_activity' => false, //If this is enabled all other skip

```

configurables below default to true.

```
'skip_groups' => false,    // Only skip group activity.
'skip_projects' => false,  // only skip project activity.
'skip_followers' => false, // only skip followers activity.
'skip_review' => false,    // only skip the review activity.

]
```

skip_activity

When **skip_activity** is set to true, all activities in Swarm are disabled. No events for any activity are sent to any user, group, project, or review.

skip_groups

When **skip_groups** is set to true, no groups will receive activity. If you use any groups as members of projects or reviewers, they will not receive any activity.

skip_projects

When **skip_projects** is set to true, none of the projects associated with the review will receive activity.

skip_followers

When **skip_followers** is set to true, all users that are following the user performing the activity will no longer receive activity.

skip_review

When **skip_review** is set to true, none of the participants (users or groups) will receive activity for actions on the review.

Slack integration


Note

- You must be an *admin* or *super* user to integrate Helix Swarm with Slack.
- Slack integration only supports the Swarm project names in lower case letters in the **config.php** file. For example, **myproject**.

This section details how to install, configure, and use the Helix Swarm Slack integration. The Helix Swarm Slack integration lets teams collaborate with Helix Swarm notifications in Slack. The Slack integration posts a message in the configured Swarm project Slack channels each time a change is committed, a review is created, or a review is updated in the project.

To install the Slack integration:

1. Navigate to <https://www.perforce.com/products/helix-swarm/slack-integration>.
2. Click **Add to Slack** button.



3. Select the workspace that you wish to integrate Slack with.
4. Request install permissions from your Application Manager.
5. Once the Application Manager approves your request, Helix Swarm Slack integration is installed into your selected workspace.
6. To enable Swarm to work with Slack, edit the Swarm **config.php** file to include the displayed Slack configuration.

The Swarm **config.php** file is saved in the `SWARM_ROOT/data/` directory.

Note

After the Helix Swarm Slack integration is installed, you must configure the project in the Swarm **config.php** file. For more information on how to configure a project for Slack integration, see ["Setup project configurations" on the facing page](#).

Slack configuration

Add the following Slack configuration to the Swarm **config.php** file. The Swarm **config.php** file is saved in the `SWARM_ROOT/data/` directory.

```
'slack' => array(
    'token' => 'TOKEN',
    'project_channels' => array(), //See "Setup project
    configurations" on the facing page to configure a project for Slack
    integration
    'summary_file_names' => false,
    'bypass_restricted_changelist' => false,
    'summary_file_limit' => 10,
    'user' => array(
        'enabled' => true,
        'name' => 'Helix Swarm',
        'icon' => 'URL',
        'force_user_header' => false,
    ),
),
```

token

The Slack app provides a token that is used by Swarm.

project_channels

Specifies the Swarm projects and Slack channels they post to in the `project_channels` array. You can configure multiple channels to receive a notification for a project. For more information on how to use `project_channels` configurable, see ["Setup project configurations" below](#).

summary_file_names

Attaches the file to the original notification message sent to a Slack channel.

bypass_restricted_changelist

Allows Swarm to post notification messages to a Slack channel when a change is committed or a review is created for a restricted changelist, default value is false. For more information about a restricted changelist, see ["Restricted Changes" on page 703](#).

summary_file_limit

Limits the number of files shown in the original notification message sent to a Slack channel, default value is 10.

user

Specifies details about the Swarm user in a Slack notification.

- **enabled:** Forces the Swarm app to use the custom user name. This overrides the Swarm app details.
- **name:** This is the user name shown in the Slack channel when a notification is posted.
- **icon:** This is the avatar icon shown in the Slack channel when a notification is posted, overrides the avatar set in the Swarm app.
- **force_user_header:** The Slack notification shows the user name and avatar only for the first notification by a user. Slack does not apply the user name and avatar to any more notifications from the same user. By default this is set to false.

When set to true, Slack adds the user name and avatar for every notification.

Setup project configurations

This section details how you can configure a project in Helix Swarm to send Slack notifications to a Slack channel.

Important

After project configuration is complete, you must restart the workers and the configuration cache for the updated configuration to take effect.

- For more information on how to restart workers, see ["Workers" on page 749](#) section.
- For more information on how to restart configuration cache, see ["Reload Configuration button" on page 732](#) section.

Note

For private projects, ensure that the Slack channel you wish to send notifications to is a private channel.

To send notifications to a channel for a project, you must configure the **project_channels** block as follows:

```
<?php
    //This block should be a peer of 'p4'
    'slack' => array(
        'token' => 'xoxb-3875744143189-3925283455653-
pxxZoTKz3JqUeDXGd66XFW4D',
        'project_channels' => array(
            //Project myproject will go into slack channel
myproject-channel
            'myproject' => array(
                'myproject-channel',
            ),
            //Project with no setting will go to no-project-channel
            '*no-project-channel*' => array(
                'no-project-channel',
            ),
            //No project on this review/change
            '*without_project*' => array(
                'no-link-project',
            ),
            //All notification get sent here even if we matched
above
            '*all*' => array(
                'all-notifications',
```

```

    ),
  ),
),

```

project_channels

Specify the Swarm projects and the Slack channels they post to in the `project_channels` array. You can configure multiple channels to receive a notification for a project.

Example

- **Swarm project:** `myproject`
- **Slack channel:** `myproject-channel`

All notifications for the project `myproject` will go to the `myproject-channel` Slack channel.

`*no-project-channel*` (optional)

For a project that has no defined project mapping and with `*no-project-channel*` configurable defined, the Slack notifications are sent to `*no-project-channel*`. For example, all notifications for the project `myproject` goes to `myproject-channel` and similarly all notifications for any other project goes to `*no-project-channel*`.

`*without_project*` (optional)

When `*without_project*` configurable is defined, it means that no notification is sent if the notification is not associated with a project.

`*all*` (optional)

When `*all*` configurable is defined, all notifications are sent here even if `*no-project-channel*` and `*without_project*` options are configured.

Swarm is now configured to send notifications to a Slack channel for a project.

Slack notification trigger

Helix Swarm will post a message on a Slack channel for every commit. This includes each time a change is committed, a review is created, or a review is updated. Any update to a review in Helix Swarm is displayed as a recently occurred activity stream in the associated Slack channel.

Here is a list of actions that will trigger Helix Swarm to send a notification to the associated project Slack channel.

Event	Description
joined	When a user joins the review
left	When a user leaves the review
made their vote required on	When a user's vote has been made required
made their vote optional on	When a user's vote has been made optional
disabled notifications on	When a user disables the notifications for a review
re-enabled notifications on	When a user re-enables the notifications for a review
edited reviewers on	When the reviewers are edited for a review
voted	When a user votes up or down
cleared their vote on	When a user clears their vote
requested further review of	When a user changes the review state to needsReview
requested revisions to	When a user changes the review state to needsRevision
rejected	When a user changes the review state to rejected
archived	When a user changes the review state to archived
approved	When a user changes the review state to approved
updated description of	When a user updates the description of the review
updated files in	When a user updates the files in a review

swarm_root

swarm_root refers to the directory where Swarm lives in the filesystem. Depending on how you installed Swarm, the default location could be:

- For a package installation: `/opt/perforce/swarm`.
- For a tarball installation: wherever you unpacked Swarm. If you are unsure, you could check your web server's configuration to discover where Swarm exists.


System Information

Note

The system information page is available to users with *admin* or *super* privileges.

The system information page provides details that can be useful to Perforce support engineers when you ask for assistance.

To display the system information page, click **System Information** from the **User id** dropdown menu:






**In this section:**

- "Perforce" below tab
- "Log" on the next page tab
- "PHP Info" on page 727 tab
- "Queue Info" on page 727 tab
- "Cache Info" on page 731 tab

Perforce

The **Perforce** tab provides information similar to the `p4 info` command.

System Information







 Perforce	 Log	 PHP Info	 Queue Info	 Cache Info
User Name	admin			
Client Name	*unknown*			
Client Cwd	/			
Client Host	swarm-deploy			
Peer Address	127.0.0.1:56220			
Client Address	127.0.0.1			
Server Address	localhost:4433			
Server Root	/opt/perforce/servers/master			
Server Date	2018/04/27 07:02:59 -0700 PDT			
Server Uptime	242:13:25			
Server Version	P4D/LINUX26X86_64/2018.1/1637071 (2018/03/23)			
Server Services	standard			
Server License	none			
Case Handling	sensitive			

Log

Click the **Log** tab to display the most recent entries, up to 1 megabyte, in Swarm's log file, which resides in `SWARM_ROOT/data/log`. Review the logging levels for [Swarm logs](#) to ensure that the entries you want to see are included.

If your log entries include a critical error, an arrow is displayed on the left of the log entry. Click the error to display the stack trace that accompanies the error.

System Information

 Perforce	 Log	 PHP Info	 Queue Info	 Cache Info	 Refresh Log	 Download Log
Time	Severity	Message				
2023-11-23T05:16:21-08:00	DEBUG	P4 (000000003a840a5500000004adf2b14) start command: Authenticating 10.153.120.243 for user super {"request":"/main/info/log?format=partial&_1700745379858","pid":1389629,"REMOTE_ADDR":"10.153.120.243"}				
2023-11-23T05:16:21-08:00	DEBUG	P4 (00000000704284d1000000005b53d5fb) start command: Authenticating 10.153.120.243 for user super {"request":"/main/info/phpinfo/","pid":1385187,"REMOTE_ADDR":"10.153.120.243"}				
2023-11-23T05:16:21-08:00	DEBUG	P4 (00000000098d83110000000021c9cf7a) start command: spec -o user {"request":"/main/api/v11/users?fields%5B%5D=avatar&fields%5B%5D=User&fields%5B%5D=FullName&fields%5B%5D=Email&fields%5B%5D=JobView&fields%5B%5D=Reviews","pid":1389639,"REMOTE_ADDR":"10.153.120.243"}				
2023-11-23T05:16:21-08:00	DEBUG	P4 (00000000098d83110000000021c9cf7a) start command: info {"request":"/main/api/v11/users?fields%5B%5D=avatar&fields%5B%5D=User&fields%5B%5D=FullName&fields%5B%5D=Email&fields%5B%5D=JobView&fields%5B%5D=Reviews","pid":1389639,"REMOTE_ADDR":"10.153.120.243"}				
2023-11-23T05:16:21-08:00	DEBUG	P4 (000000003f463ef90000000066126b26) start command: info {"request":"/main/api/v11/groups?fields%5B%5D=id&fields%5B%5D=name&expand=false","pid":1385246,"REMOTE_ADDR":"10.153.120.243"}				
2023-11-23T05:16:21-08:00	DEBUG	P4 (000000003f463ef90000000066126b26) start command: spec -o group {"request":"/main/api/v11/groups?fields%5B%5D=id&fields%5B%5D=name&expand=false","pid":1385246,"REMOTE_ADDR":"10.153.120.243"}				
2023-11-23T05:16:21-08:00	INFO	User super making a request for endpoint http://freya.das.perforce.com/main/api/v11/users?fields%5B%5D=avatar&fields%5B%5D=User&fields%5B%5D=FullName&fields%5B%5D=Email&fields%5B%5D=JobView&fields%5B%5D=Reviews for method GET {"request":"/main/api/v11/users?fields%5B%5D=avatar&fields%5B%5D=User&fields%5B%5D=FullName&fields%5B%5D=Email&fields%5B%5D=JobView&fields%5B%5D=Reviews","pid":1389639,"REMOTE_ADDR":"10.153.120.243"}				
2023-11-23T05:16:21-08:00	INFO	User super making a request for endpoint http://freya.das.perforce.com/main/api/v11/groups?fields%5B%5D=id&fields%5B%5D=name&expand=false for method GET {"request":"/main/api/v11/groups?fields%5B%5D=id&fields%5B%5D=name&expand=false","pid":1385246,"REMOTE_ADDR":"10.153.120.243"}				

Log tab buttons:

- **Refresh Log** button: click to load the latest log entries, up to 1 megabyte.
- **Download Log** button: click to download the log and all of its log entries.

PHP Info

Click the **PHP Info** tab to display PHP's own information display generated by executing `phpinfo()`, PHP's internal [diagnostic display](#).

System Information

[Perforce](#)
[Log](#)
[PHP Info](#)
[Queue Info](#)
[Cache Info](#)

PHP Version 7.0.32-0ubuntu0.16.04.1



System	Linux vagrant-swarm-devline-ub1604 4.4.0-116-generic #140-Ubuntu SMP Mon Feb 12 21:23:04 UTC 2018 x86_64
Server API	Apache 2.0 Handler
Virtual Directory Support	disabled
Configuration File (php.ini) Path	/etc/php/7.0/apache2
Loaded Configuration File	/etc/php/7.0/apache2/php.ini
Scan this dir for additional .ini files	/etc/php/7.0/apache2/conf.d
Additional .ini files parsed	/etc/php/7.0/apache2/conf.d/10-opcache.ini, /etc/php/7.0/apache2/conf.d/10-pdo.ini, /etc/php/7.0/apache2/conf.d/15-xml.ini, /etc/php/7.0/apache2/conf.d/20-apcu.ini, /etc/php/7.0/apache2/conf.d/20-calendar.ini, /etc/php/7.0/apache2/conf.d/20-ctype.ini, /etc/php/7.0/apache2/conf.d/20-curl.ini, /etc/php/7.0/apache2/conf.d/20-dom.ini, /etc/php/7.0/apache2/conf.d/20-exif.ini, /etc/php/7.0/apache2/conf.d/20-fileinfo.ini, /etc/php/7.0/apache2/conf.d/20-ftp.ini, /etc/php/7.0/apache2/conf.d/20-gettext.ini, /etc/php/7.0/apache2/conf.d/20-iconv.ini, /etc/php/7.0/apache2/conf.d/20-imagick.ini, /etc/php/7.0/apache2/conf.d/20-intl.ini, /etc/php/7.0/apache2/conf.d/20-json.ini, /etc/php/7.0/apache2/conf.d/20-mbstring.ini, /etc/php/7.0/apache2/conf.d/20-perforce.ini, /etc/php/7.0/apache2/conf.d/20-phar.ini, /etc/php/7.0/apache2/conf.d/20-posix.ini, /etc/php/7.0/apache2/conf.d/20-readline.ini, /etc/php/7.0/apache2/conf.d/20-shmop.ini, /etc/php/7.0/apache2/conf.d/20-simplexml.ini, /etc/php/7.0/apache2/conf.d/20-sockets.ini, /etc/php/7.0/apache2/conf.d/20-sysvmsg.ini, /etc/php/7.0/apache2/conf.d/20-sysvsem.ini, /etc/php/7.0/apache2/conf.d/20-sysvshm.ini, /etc/php/7.0/apache2/conf.d/20-tokenizer.ini, /etc/php/7.0/apache2/conf.d/20-wddx.ini, /etc/php/7.0/apache2/conf.d/20-xdebug.ini, /etc/php/7.0/apache2/conf.d/20-xmlreader.ini, /etc/php/7.0/apache2/conf.d/20-xmlwriter.ini, /etc/php/7.0/apache2/conf.d/20-xsl.ini, /etc/php/7.0/apache2/conf.d/20-zip.ini
PHP API	20151012
PHP Extension	20151012
Zend Extension	320151012
Zend Extension Build	API320151012,NTS
PHP Extension Build	API20151012,NTS
Debug Build	no
Thread Safety	disabled
Zend Signal Handling	disabled


Queue Info

Click the **Queue Info** tab to display the Swarm task queue information.

Note

Workers should be started automatically using a recurring task. For instructions on how to set up a recurring task to spawn workers automatically, see ["Set up a recurring task to spawn workers" on page 220](#).

System Information

 [Perforce](#) [Log](#) [PHP Info](#) [Queue Info](#) [Cache Info](#)

Tasks	0
Future Tasks	4
Workers	1
Max Workers	3
Worker Lifetime	595s
Ping Error	false

[Start a Worker](#)[Refresh Tab](#)[Show Task Queue](#)[Start a Temp Worker](#)

- **Tasks:** displays the number of tasks in the queue.
- **Future Tasks:** displays the number of future tasks in the queue. There are usually at least four future tasks in the queue for checking cache integrity, but there can be more. For example, when a user is making comments, the comment notification is delayed by 30 minutes by default. This is a future task.
- **Workers:** displays the number of workers available to service tasks.

Tip**If the number of Workers is 0 (zero):**

- Your workers are not being created automatically. For instructions on how to set up a recurring task to spawn workers automatically, see "[Set up a recurring task to spawn workers](#)" on page 220.
- To process the current task queue while you get your recurring tasks working again, click **Start a Temp Worker**. This is a quick way to get Swarm running again while you fix the problem.

- **Max Workers:** displays the [maximum number of workers](#) that are allowed. The default is three.
- **Worker Lifetime:** displays the [worker lifetime](#) that is configured. The default is 595 seconds (10 minutes less 5 seconds).
- **Ping Error:** not used.

Queue Info buttons:

- **Start a Worker** button, only available if **Workers** is less than the [Max Workers](#) setting: click to manually start a new Swarm worker. The new worker will only last for the [configured Worker Lifetime](#). If the [number of workers running](#) matches the configured limit, the new worker is not started.

Tip

If your recurring task has stopped working and you need to process the task queue while you get it working again, start a temporary worker by clicking **Start a Temp Worker**. This is a quick way to get Swarm running again while you fix the problem.

- **Refresh Tab** button: click to refresh the queue information in the tab.
- **Show Task Queue** button: click to display the Task Queue information. When the task queue is displayed the button changes to **Hide Task Queue**. The task queue shows more information about the tasks and future tasks in the queue.
- **Start a Temp Worker** button, only available if **Workers** is less than the [Max Workers](#) setting: click to manually start a Swarm temporary worker. Typically you use a temporary worker to process the current tasks in the queue while you fix the recurring task that automatically spawns the workers.

This is a special Swarm worker that is not subject to the worker lifetime setting and will not stop until it has processed all of the tasks in the current queue or one of the tasks fails.

Task Queue

Click the **Show Task Queue** button to display the tasks in the Swarm queue.

System Information

[Perforce](#) [Log](#) [PHP Info](#) [Queue Info](#) [Cache Info](#)

Tasks	0
Future Tasks	4
Workers	0
Max Workers	3
Worker Lifetime	595s
Ping Error	false

[Start a Worker](#) [Refresh Tab](#) [Show Task Queue](#) [Start a Temp Worker](#)

Task Queue

Future task - 1627639200.0000.b7b69a2f99a3b1a1abdb2ad48d9e44d4.0

File /opt/perforce/swarm/data/servers/main/queue/1627639200.0000.b7b69a2f99a3b1a1abdb2ad48d9e44d4.0
Time Fri Jul 30 2021 11:00:00 GMT+0100 (British Summer Time)
Type cache.integrity
Id project
Data swarm, project,

Future task - 1627639200.0000.6b2b3b0dba18ad5b3697bd09a3785391.0

File /opt/perforce/swarm/data/servers/main/queue/1627639200.0000.6b2b3b0dba18ad5b3697bd09a3785391.0
Time Fri Jul 30 2021 11:00:00 GMT+0100 (British Summer Time)
Type cache.integrity
Id workflow
Data swarm, workflow,

Future task - 1627639200.0000.4d90339f34d63ac3dd776502b5ca42e4.0

File /opt/perforce/swarm/data/servers/main/queue/1627639200.0000.4d90339f34d63ac3dd776502b5ca42e4.0
Time Fri Jul 30 2021 11:00:00 GMT+0100 (British Summer Time)
Type cache.integrity
Id group
Data swarm, group,

Future task - 1627639200.0000.33981a32e56f1ad82c10d78b4ef01d39.0

File /opt/perforce/swarm/data/servers/main/queue/1627639200.0000.33981a32e56f1ad82c10d78b4ef01d39.0
Time Fri Jul 30 2021 11:00:00 GMT+0100 (British Summer Time)
Type cache.integrity
Id user
Data swarm, user,

Cache Info

Click the **Cache Info** tab to display the Swarm cache information.

System Information

Perforce
Log
PHP Info
Queue Info
Cache Info

Context	State	Progress
User	Running	Step 4 of 5: Clearing any extraneous Redis cache entries
Group	Queued	Verification complete
Project	Queued	Verification complete
Workflow	Queued	Verification complete

Refresh Status
Verify All
Verify User
Verify Group
Verify Project
Verify Workflow
Reload Configuration

Tip

By default, Swarm runs a regular verification check at 03:00 every day to make sure that the Redis caches and Helix Core Server are in sync, see "[Swarm connection to Redis](#) " on page 674.

The **Cache Info** tab displays the cache verification state and progress for each of the Redis caches:

- Valid states for verify are; **Not Queued**, **Queued**, **Running**, and **Failed**.
- Valid states for progress are, **Step x of y**; (where **x** is the current step and **y** is the total number of steps) followed by:
 - **Getting Redis cache entries**
 - **Calculating checksums for Redis keys**
 - **Calculating checksums for Perforce models**
 - **Clearing any extraneous Redis cache entries**
 - **Indexing any missing Perforce models into the Redis cache**
 - **Verification complete**

Refresh Status button

To refresh the cache information on the tab, click **Refresh Status**.

Verify buttons

Important

On Swarm systems with a large number of users, groups, and projects verification of the caches can take some time and can impact performance.

In normal operation, the Redis cache and Helix Core Server data should stay in sync and there is no need to verify the caches.

If you have a large number of users, groups, and projects, we recommend that you verify the individual user and group caches rather than using **Verify All**.

In normal operation the Redis cache and Helix Core Server data should stay in sync and there is no need to verify the caches. However, in some circumstances, such as when changes are made in the Helix Core Server when Swarm is down or if errors occur during updates, the Redis cache can get out of sync with the Helix Core Server.

If you suspect that you have a cache issue, the verify buttons enable you to manually verify the caches. Before using the **Verify All** button, see the Important note above. If verification finds the cache is out of sync with the Helix Core Server, Swarm automatically updates the data in that cache.

For example, to verify the User cache:

1. Navigate to the **Cache Info** tab on the Swarm **System Information** page.
2. Click **Verify User**.

Swarm responds with a message to say that it has successfully verified the user cache.

Reload Configuration button

The Swarm configuration is stored in the `SWARM_ROOT/data/config.php` file. If you make a configuration change, Swarm will not use it until the configuration cache has been deleted and reloaded, this forces Swarm to use the new configuration.

To delete and reload the Swarm configuration cache:

1. Navigate to the **Cache Info** tab on the Swarm **System Information** page.
2. Click **Reload Configuration**.

Swarm responds with a message to say that it has successfully reloaded the config.

Swarm is now using the updated Swarm `config.php` file.

Tag processor

Including a work-in-progress tag (`#wip` by default) in a pending changelist tells Swarm not to process the changelist when you update files in the shelf. This means that if the changelist already has a review that review will not be updated. If it doesn't have a review, a new review will not be created. For information on using the work-in-progress tag, see "[Work-in-progress tag](#)" on page 393.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button"](#) on the previous page.

wip tag configuration

The work-in-progress tag is configured in the **tags** section of the **tag_processor** block in the **"swarm_root"** on [page 724/data/config.php](#) file. It is enabled by default:

```
<?php
// this block should be a peer of 'p4'
'tag_processor' => array(
    'tags' => array(
        'wip' => '/(^|\s)+#wip($|\s)+/i'
    ),
),
```

Disabling the work-in-progress tag

You can disable the work-in-progress tag completely if you don't want anyone to use it.

To disable the work-in-progress tag, delete the content of the **wip** string in the **"swarm_root"** on [page 724/data/config.php](#) file so that it is empty (**' '**):

```
<?php
// this block should be a peer of 'p4'
'tag_processor' => array(
    'tags' => array(
        'wip' => ''
    ),
),
```

Customizing the work-in-progress tag

By default, the work-in-progress tag is **#wip**, but you can change it to whatever you want to use. Swarm uses Perl Compatible Regular Expressions (PCRE) to process the **wip** tag, ensure that your new tag can be processed by PCRE before changing it. If your **wip** tag cannot be processed by PCRE, Swarm might ignore the tag or stop working correctly.

For example, to change the work-in-progress tag to `#noreview`, update the tag processor `wip` configuration in the `"swarm_root"` on page 724/`data/config.php` file to:

```
<?php
// this block should be a peer of 'p4'
'tag_processor' => array(
    'tags' => array(
        'wip' => '/(^|\s)+#noreview($|\s)+/i'
    ),
),
```

Making the work-in-progress tag case sensitive

To make the work-in-progress tag case sensitive, delete the trailing `i` from the `wip` configuration in the `"swarm_root"` on page 724/`data/config.php` file:

```
<?php
// this block should be a peer of 'p4'
'tag_processor' => array(
    'tags' => array(
        'wip' => '/(^|\s)+#wip($|\s)+/'
    ),
),
```

Test definitions

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button"](#) on page 732.

project_and_branch_separator

Note

The `project_and_branch_separator` configurable character is only applied to tests created from the **Tests** menu, see ["Tests"](#) on page 544.

Specifies the character used to separate the `{projects}` and `{branches}` IDs in the URL that is passed to your CI system. Some CI systems, such as Jenkins, escape the default colon `:` character resulting in a failed test request. If your CI system needs a different separator character, the `project_and_branch_separator` character can be changed to a character that your CI system will accept, for example a dash `-` or an underscore `_`.

Be careful when selecting your separator character, some characters can cause issues when calling tests. For example, a `#` character in a URL call means the branch is treated as an anchor and a `%` character in a JSON body encoding is escaped. Both of these will result in a failed test request.

Configure the `project_and_branch_separator` character with the following configuration block in the `SWARM_ROOT/data/config.php` file:

```
'test_definitions' => array(
    'project_and_branch_separator' => ':',
),
```

The default value for the `project_and_branch_separator` configurable is a colon `:`.

Trigger options

Warning

If you are using Swarm Helix Core Server Triggers, the Swarm Helix Core Server extension must not be installed.

Tip

We recommend you use Helix Core Server Extensions, Helix Core Server Extensions are easier to install and maintain than triggers. See ["Helix Core Server Extensions dependencies" on page 108](#).

The Swarm trigger script, `swarm-trigger.pl`, provides the following command line options and configuration items.

Command-line options

Synopsis

```
swarm-trigger.pl -t type -v ID [-c config_file]
```

```
swarm-trigger.pl -o
```

```
swarm-trigger.pl -h
```

Informational options

The following options perform no processing, and simply provide information useful to a Swarm administrator. These are not intended to be used within trigger entries in the Helix Core Server.

- **-h**

Displays a list of the available options, some guidance on its usage, and a copy of the trigger table entries that should be configured in the Helix Core Server to execute the script in its current path.

- **-o**

Displays a copy of the trigger table entries that should be configured in the Helix Core Server to execute the script in its current path.

Operational options

The following command-line options are used in the trigger table entries in the Helix Core Server to specify how the Swarm trigger script should be executed.

- **-t type**

Specifies the type of processing that the trigger script undertakes. **type** can be one of:

- **job**: this type should be used with Helix Core Server **form-commit** events for job forms, and informs Swarm when a job is created or modified.
- **user**: this type should be used with Helix Core Server **form-commit** events for user forms, which informs Swarm when a user is added or modified.
- **userdel**: this type should be used with Helix Core Server **form-delete** events for user forms, which informs Swarm when a user is deleted.
- **group**: this type should be used with Helix Core Server **form-commit** events for group forms, and informs Swarm when a group is created or modified.
- **groupdel**: this type should be used with Helix Core Server **form-delete** events for group forms, and informs Swarm when a group is deleted.
- **changesave**: this type should be used with Helix Core Server **form-save** events for change forms, and informs Swarm when a changelist is created or modified.
- **shelve**: this type should be used with Helix Core Server **shelve-commit** events, and informs Swarm when a changelist is shelved, which can create or update a review.
- **commit**: this type should be used with Helix Core Server **change-commit** events, and informs Swarm when a changelist is committed.
- **shelvedel**: this type is used with Helix Core Server **shelve-delete** events, and informs Swarm when a file is deleted from a shelf.

Important

Do not use the Swarm user that is configured in the [Swarm configuration file](#) when deleting shelves, or deleting files from shelves. The Swarm logic processes the *shelf-delete* trigger event, if the event is invoked by the Swarm user it is rejected. The delete operations will fail.

- **enforce** comment out if the Workflow feature is [enabled](#) (enabled by default): this type is used to verify that commits to specific depot paths are associated with approved reviews. If a commit includes a file within the specified depot path, and it is not associated with a review (or a review that is not approved), the commit is rejected.

Using the **enforce** type prevents users from committing changes to specific depot paths without those changes being reviewed and approved.

- **strict** comment out if the Workflow feature is [enabled](#) (enabled by default): this type is used to verify that the file content in a commit matches the file content of its associated approved review. If one or more files in a commit do not match the content of the file in its associated review, the commit is rejected.

Using the **strict** type prevents users from making changes to the file content after a review has been approved and then submitting the unapproved changes.

Note

strict implies **enforce**.

- **Workflow trigger types** comment out if the Workflow feature is [disabled](#) (enabled by default). The **enforce** and **strict** triggers above must be commented out if the Workflow feature is [enabled](#):

Important Known limitations

The workflow triggers do not support the following trigger functionality available in Swarm 2018.1 and earlier:

- **EXEMPT_FILE_COUNT**: When set to a positive integer, commits with a file count greater or equal to this value are exempt from **enforce** or **strict** verifications.
- **EXEMPT_EXTENSIONS**: A comma separated list of file extension. Commits with files having only these extensions are exempt from **enforce** or **strict** verifications.

To continue to use this functionality, you must keep your existing **enforce** and **strict** triggers and [disable the Workflow feature](#).

Tip

By default, all group members and users must follow any workflow rules that apply to changelists and reviews. However, your Swarm administrator can configure specific groups and users so that they are excluded from following the workflow rules. This applies to actions taken in the Swarm UI, the Helix Core command-line (P4), a P4D client, and the Swarm API.

The following trigger types are used by Swarm for workflow rules. They are not intended for use in custom trigger scripts.

- **checkenforced**: this type is used by the Swarm **On commit with a review workflow rule** and cannot be configured. **checkenforced** triggers on a Helix Core Server **change-submit** event.

If a commit is made to a branch that has **On commit with a review** set to **Reject unless approved**, and it is not associated with a review (or the review is not approved), the commit is rejected. If the commit is associated with an approved review, a content check is carried out by **checkstrict**.

- **checkstrict**: this type is used by the Swarm **On commit without a review workflow rule** and cannot be configured. **checkstrict** triggers on a Helix Core Server **change-content** event.

This check is performed after **checkenforced**. If the **On commit without a review** rule is set to **Reject**, and one or more files in a commit do not match the content of the file in its associated review, the commit is rejected.

- **checkshelve**: this type is used by the Swarm **On update of a review in an end state workflow rule** and cannot be configured. **checkshelve** triggers on a Helix Core Server **shelve-submit** event.

If a shelf is associated with a review, and it is committed to a branch that has **On update of a review in an end state** set to **Reject**, Swarm checks what state the review is in to make sure it is not in a **protected state**. If the review is in a protected state, the commit is rejected.

Important

You cannot mix Swarm trigger types with unrelated Helix Core Server events; the behavior is undetermined, and the information required for each type of processing may not be available to the trigger.

■ **-v ID**

Specifies **ID**, which is the identifier for the current trigger type.

When the **type** is **job**, **user**, **userdel**, **group**, **groupdel**, or **changesave**, **ID** should be **%formname%** to specify the specific form identifier Swarm should process.

When the **type** is **shelve**, **commit**, **shelvedel**, **checkenforced**, **checkstrict**, or **checkshelve** **ID** should be **%change%** to specify the specific changelist Swarm should process.

■ **-p port (optional) Do not use if the Workflow feature is enabled (default)**

Specifies the Helix Core Server port (**P4PORT**). This value is optional, and is only used for types **enforce** or **strict** as the trigger has to run its own commands against the Helix Core Server during its processing.

- **-r (optional) Do not use if the Workflow feature is enabled (default)**
Specifies that, when types `enforce` or `strict` are being processed, the verifications should only be performed on commits that are currently in review.
- **-g (optional) Do not use if the Workflow feature is enabled (default)**
Specifies a `group` to exclude from `enforce` or `strict` verifications. Members of the group (including sub-groups) are not subject to `enforce` or `strict` verifications.
- **-c config_file (optional)**
Specifies an optional `config_file` which is used to specify configuration items.

Configuration items

The following configuration items are used in the `swarm-trigger.conf` (or another file, if the `-c` option is used in the trigger entries).

- **SWARM_HOST (required)**
Specifies the host URL of your Swarm instance, with the leading `http://` or `https://`. For example: `https://myswarm.url`.
- **SWARM_TOKEN (required)**
A token used when talking to Swarm. To obtain the token, log into Swarm as a user with *super* privileges and select the [About Swarm](#) from the user menu in the main navigation bar.

You can also manually create additional tokens. Tokens are empty files stored within `SWARM_ROOT/data/queue/tokens` (the filename is the token and is reported on the **About Swarm** dialog), that should be readable by the web server.

You might manually create additional tokens to allow other processes to talk to Swarm, such as JIRA build tasks, and to selectively invalidate access to Swarm without interfering with regular Swarm operations.
- **ADMIN_USER (optional) Do not use if the Workflow feature is enabled (default)**
When `enforce` or `strict` verifications are to be performed, you may need specify a username of a user in the Helix Core Server that has *admin* privileges. If you do not specify a username, the trigger script uses the Helix Core Server user set in the environment.
- **ADMIN_TICKET_FILE (optional) Do not use if the Workflow feature is enabled (default)**
When `enforce` or `strict` verifications are to be performed, you may need specify the path to the `.p4tickets` file, if your Helix Core Server tickets file is not the default `$HOME/.p4tickets`.

Important

Ensure that the ticket belongs to a user with *admin* privileges in the Helix Core Server, and is a member of a group with an **unlimited** or very long ticket timeout. If this user's authentication times out, **enforce** and **strict** verifications stop working.

- **P4_PORT (optional) Do not use if the Workflow feature is enabled (default)**

When **enforce** or **strict** verifications are to be performed, you may need to set the port value (**P4PORT**) of the Helix Core Server, particularly if the Helix Core Server is on a non-standard port, or if the Helix Core Server is not using the default hostname.

- **P4 (optional) Do not use if the Workflow feature is enabled (default)**

Specifies the full path to **p4**, the Helix Core Server command-line client. This is only required when **p4** is not found in the **PATH** of the Helix Core Server environment, and when **enforce** or **strict** verifications are to be performed.

- **EXEMPT_FILE_COUNT (optional) Do not use if the Workflow feature is enabled (default)**

When set to a positive integer, commits with a file count greater or equal to this value are exempt from **enforce** or **strict** verifications.

- **EXEMPT_EXTENSIONS (optional) Do not use if the Workflow feature is enabled (default)**

A comma-separated list of file extensions. Commits with files having only these extensions are exempt from **enforce** or **strict** verifications.

- **VERIFY_SSL (optional)**

If HTTPS is used on the Swarm web server, this controls whether the SSL certificate is validated or not.

By default **VERIFY_SSL** is set to **1**, this means any SSL certificates must be valid. If the Swarm web server is using a self signed certificate, **VERIFY_SSL** must be set to **0**.

- **TIMEOUT (optional)**

Sets the number of seconds to wait before returning an error when communicating with the Swarm server.

The default setting is **30** seconds.

- **IGNORE_TIMEOUT (optional)**

Configure the action to take if communication with the Swarm server times out. For example, if you are submitting a large number of files to the Helix Core Server.

Options are:

- **0** if communication with the Swarm server times out, the submit will fail. This is the default setting.

- **1** if communication with the Swarm server times out, the submit will continue. This increases the reliability of submits.
- **IGNORE_NOSERVER (optional)**
- Configure the action to take if communication with the Swarm server fails. For example, the Swarm server is down.
- Options are:
- **0** if communication with the Swarm server fails, the submit will fail. This is the default setting.
 - **1** if communication with the Swarm server fails, the submit will continue. This increases the reliability of submits.

Unapprove modified reviews

By default, when an approved review is committed or updated, Swarm changes the state to **Needs Review** if the files have been modified since the review was approved.

If one or more files in a review has the filetype +k (**ktext**), this behavior is undesirable because the files will appear to be modified as the Helix Core Server replaces RCS keywords with their current values. This behavior can be disabled.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

To disable this behavior, edit the `$SWARM_ROOT/data/config.php` file, and add or update the `unapprove_modified` item to **false**, within the reviews configuration block. For example:

```
<?php
// this block should be a peer of 'p4'
'reviews' => array(
    'unapprove_modified' => false,
),
```

For information on file types, see [File Types](#) in *Helix Core Command-Line (P4) Reference*.

Uninstall Swarm

This section covers the steps required to uninstall Swarm.

Background

The bulk of Swarm's metadata (activity, comments, review records, followers) is stored in p4 keys under **swarm-***. If you are using a 2012.1+ server, Swarm also defines user groups for each project that you define. The names of these groups correspond 1-to-1 with projects, for example **swarm-project-fantastico**. Swarm manages a pool of client workspaces that it uses to shelve and commit files. These clients are named **swarm-{uuid}**, for example **swarm-5ad4a9c0-06e7-20eb-897f-cbd4cc934295**.

Uninstall steps

1. Depending on how you have configured your Helix Core Server events, do one of the following:
 - Uninstall the Swarm Helix Core Server extension. To do this, run the following command on your Helix Core Server:


```
p4 extension --delete Perforce::helix-swarm --yes
```
 - Uninstall the Swarm triggers. As a *super* user, run the **p4 triggers** command from your Helix Core Server and manually remove all the Swarm trigger code lines.
2. Remove your web server's virtual host configuration for Swarm by disabling the Swarm site (**perforce-swarm-site**). If you have made any non-standard changes for Swarm such as SSL configuration or if an alternate web server is used, you must remove these changes for Swarm.

Depending on your OS distribution, do one of the following:

- **Ubuntu** (run the following command as root):
 - a. Disable the Swarm site:


```
a2dissite perforce-swarm-site
```
 - b. Delete the Swarm site:


```
Delete /etc/apache2/sites-available/perforce-swarm-site.conf
```
 - **CentOS/RHEL/Amazon Linux 2** (run the following command as root):

Delete the Swarm site:

```
Delete /etc/httpd/conf.d/perforce-swarm-site.conf
```
3. Restart your web server.
 4. Delete groups/clients/keys that are prefixed with **swarm-***.

Note

The clients could contain shelved files for reviews. Determine how you want to handle those files prior to deleting the clients.

- Additional indexed information is stored in the database file `db.ixtext`. Unfortunately, indexed jobs and other generic indexed information would be lost if this table was simply removed, and modifying the database file can be a dangerous operation in a number of Helix Core Server deployment scenarios.

Important

Contact Perforce support for assistance if you feel the need to remove Swarm's indexed information: support@perforce.com.

- Rebuild the job index. The best approach is to run:

```
$ p4 jobs -R
```

which rebuilds the `db.ixtext` table. There are two caveats that likely require discussion with support@perforce.com:

- If you make use of the unsupported `p4 index` command, you **cannot** use this approach, as it would remove all of your indexes.
- If you have indexing turned on for the domain table, you must also run:

```
$ p4d -xf index.domain.owner
```

- If the `P4.Swarm.URL` or `P4.Swarm.CommitURL` properties were set (for details, see "[Client integration](#)" on page 578 and "[Commit-edge deployment](#)" on page 587 respectively), they should be unset to prevent P4V (and potentially other clients) from attempting Swarm operations. We recommend you view all the properties and delete them all.

To delete each Swarm property:

```
$ p4 property -d -n <P4.Swarm property> -s <sequence>
```

The sequence is provided in the property list. A property with `#none` is **sequence 0**.

To discover all of the definitions of the `P4.Swarm.*` properties and their sequence numbers, as a Perforce **super** user run the following command:

```
$ p4 property -Al P4.Swarm
```

For example, the following property values are returned that apply to all users and groups:

```
$ p4 property -Al P4.Swarm
P4.Swarm.CommitURL = https://myswarm.url/ (any) #none
P4.Swarm.CommitURL = https://myswarm1.url/ (any) #1
P4.Swarm.URL = https://myswarm.url/ (any) #none
P4.Swarm.URL = https://myswarm3.url/ (any) #3
P4.Swarm.URL = https://myswarm2.url/ (any) #2
P4.Swarm.URL = https://myswarm1.url/ (any) #1
```

To delete all of these property definitions, you would run the following commands:

```
$ p4 property -d -n P4.Swarm.CommitURL -s0
$ p4 property -d -n P4.Swarm.CommitURL -s1
$ p4 property -d -n P4.Swarm.URL -s0
$ p4 property -d -n P4.Swarm.URL -s3
$ p4 property -d -n P4.Swarm.URL -s2
$ p4 property -d -n P4.Swarm.URL -s1
```

To verify if all the properties have been deleted, run the following command as a Perforce **super** user:

```
$ p4 property -Al P4.Swarm
```

Upgrade index

The Swarm Upgrade index process can be configured to suit the performance of your Swarm system via the following config items.

Note

If you are upgrading from Swarm version 2017.3 or later, the index upgrade step is not required.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

status_refresh_interval

The `status_refresh_interval` sets the refresh rate of the index upgrade status page. The default is 10 seconds.

```
<?php
    // this block should be a peer of 'p4'
    'upgrade' => array(
        'status_refresh_interval' => 10, //Refresh page every 10
seconds
    ),
```

batch_size

The `batch_size` sets the number of reviews held in memory and processed for each batch. This value can be increased or decreased depending on your Swarm machines system memory. The default is 1000 reviews.

For example:

If the Swarm machine does not have a lot of memory, for instance 128MB, a `batch_size` of 1000 might be too high and may make the upgrade run very slowly. In this instance reduce the `batch_size` to a more manageable size.

```
<?php
    // this block should be a peer of 'p4'
    'upgrade' => array(
        'batch_size' => 1000,    //Fetch 1000 reviews to lower memory
usage
    ),
```

Users

This section describes the configurables that can be set for users.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button"](#) on page 732.

Dashboard refresh interval

By default, when your dashboard is open, Swarm checks for new content every five minutes. If Swarm finds new content for a user's dashboard, the **Refresh** button is displayed on their dashboard. For more information about using the **Refresh** button, see [Refresh button](#).

Configure the dashboard refresh interval in milliseconds with the following configuration block in the `SWARM_ROOT/data/config.php` file:

```
<?php
    // this block should be a peer of 'p4'
    'users' => array(
        'dashboard_refresh_interval' => 300000, // Default 300000
milliseconds
```

```
    ),
  ),
```

Default value is **300000** milliseconds (300 seconds, 5 minutes)

Display full names

Note

For Swarm 2022.1 and later, the **display_fullname** value defaults to **true**. For earlier Swarm releases, the default was **false**. Upgrading your Swarm version to 2022.1 and later will not change your **display_fullname** setting.

By default, Swarm displays full user names. Set the **display_fullname** configurable to **false** to display *userids* on the:

- [Project Settings](#) page
- [Group Settings](#) page
- [Activity streams](#)
- [Commits](#) page
- [Edit reviewers](#) dialog

To display *userids* in Swarm, set the following configuration block in the [SWARM_ROOT/data/config.php](#) file to **false**:

```
<?php
    // this block should be a peer of 'p4'
    'users' => array(
        'display_fullname' => false,
    ),
),
```

Default value is **true**, full user names are displayed.

Review preferences

The review preferences configure the default way that [diffs](#) are initially displayed to users when they view them in changelists and code reviews. The review preference settings are used for a user until they change their user preferences, see user profile [Settings](#).

Configure the default user diff view settings with the following configuration block in the [SWARM_ROOT/data/config.php](#) file:

```
<?php
    // this block should be a peer of 'p4'
```

```
'users' => array(
  'settings' => array(
    'review_preferences' => array(
      'show_comments_in_files' => true,
      'view_diffs_side_by_side' => true,
      'show_space_and_new_line_characters' => false,
      'ignore_whitespace' => false,
    ),
  ),
),
```

show_comments_in_files:

- **true**: display comments in files and in the **Comments** tab. This is the default value.
- **false**: display comments only in the **Comments** tab.

view_diffs_side_by_side:

- **true**: display diffs in files side by side in two panes. This is the default value.
- **false**: display diffs in files inline in a single pane.

show_space_and_new_line_characters:

- **true**: display whitespace characters as dots, tabs as arrows that point to a bar, and line endings as arrows that point down.
- **false**: whitespace, tabs and line endings are not displayed as special characters. This is the default value.

ignore_whitespace:

- **false**: changes made to whitespace are highlighted in file diffs. This makes it easier to identify changes in file types where whitespace is important. This is the default value.
- **true**: changes made to whitespace are not highlighted in file diffs. This makes it easier to see the important changes in file types where whitespace changes are not important.

Note

ignore_whitespace replaces **ignore_whitespace_default** that was used in Swarm version 2017.4 and earlier. By default **ignore_whitespace** is set to false, this is the same original default value set for **ignore_whitespace_default**.

Time preferences

The time preferences configure the default way that time is initially displayed to users. The time preference settings are used for a user until they change their user preferences, see [user profile Settings](#).

Configure the default user time settings with the following configuration block in the [SWARM_ROOT/data/config.php](#) file:

```
<?php
    // this block should be a peer of 'p4'
    'users' => array(
        'settings' => array(
            'time_preferences' => array(
                'display' => 'Timeago', // Default to 'Timeago' but
can be set to 'Timestamp'
            ),
        ),
    ),
```

display:

- **'Timeago'**: time is displayed as how long ago the event happened. This is the default value.
- **'Timestamp'**: time is displayed as the date and time the event happened using the local machine browser time.

Tip

If timestamp is selected, the date format used by Swarm matches the date format configuration of the local machine browser.

Workers

Helix Swarm uses background processes, called *workers*, to respond to events in the Helix Core Server. The default number of workers is three, and each worker processes events for up to 10 minutes. When a worker terminates, a new one is spawned.

Note

Each worker maintains a connection to the Helix Core Server for the duration of its lifetime. This may impact your Helix Core Server management practices.

Tip

To view Swarm task queue information, as an *admin* or *super* user, navigate to the **User id** dropdown menu, select **System Information**, and open the "Queue Info" on page 727 tab.

Worker status

To determine the current status of workers, do one of the following:

- As an *admin* or *super* user, navigate to the **User id** dropdown menu, select **System Information**, and open the "Queue Info" on page 727 tab.

or

- Visit the URL: `https://myswarm.url/queue/status`.

The response is formatted in JSON, and looks like this:

```
{ "tasks": 0, "futureTasks": 1, "workers": 3, "maxWorkers": 3, "workerLifetime": "595s" }
```

During normal use of Swarm, the following error message appears for logged-in users when Swarm detects that no workers are running:

Hmm... no queue workers? Ask your administrator to check the [worker setup](#).

Worker configuration

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the "Reload Configuration button" on page 732.

To adjust the configuration for workers, add a configuration block to the `SWARM_ROOT/data/config.php` file:

```
<?php
// this block should be a peer of 'p4'
'queue' => array(
    'workers'           => 3,      // defaults to 3
    'worker_lifetime'   => 595,    // defaults to 10 minutes (less
5 seconds)
    'worker_task_timeout' => 1800, // defaults to 30 minutes
    'worker_memory_limit' => '1G', // defaults to 1 gigabyte
),
```

where:

- **workers** specifies the number of worker processes that should be available. The default is 3. The **cron job** ensures that new worker processes are started when necessary. If the limit is reached or exceeded, new worker processes are not started.
- **worker_lifetime** specifies the amount of time in seconds that a worker process should run for. The default is 595 seconds (10 minutes less 5 seconds). If a worker process exceeds this limit while processing a task, it will complete the active task and then terminate. **worker_lifetime** does not cause tasks to terminate mid-processing.
- **worker_task_timeout** specifies the maximum amount of time in seconds that a worker process can spend processing a single task. The default is 1800 seconds (30 minutes). This is useful for terminating workers that might get stalled in a variety of situations.
- **worker_memory_limit** specifies the maximum amount of memory that a worker process is allowed to use while processing a task. The default is 1G (1 gigabyte).

Manually start workers

To kick off a new worker process, do one of the following:

- As an *admin* or *super* user, navigate to the **User id** dropdown menu, select **System Information**, open the "Queue Info" on page 727 tab, and click **Start a worker**.
- or
- Visit the URL: **`https://myswarm.url/queue/worker`**.

If the number of workers running matches the configured limit, the requested worker process is not started.

Note

This technique does start a worker, but it lasts only for its configured lifetime. Typically, you would always want at least one worker running. See "Set up a recurring task to spawn workers" on page 220 for details.

Manually restart workers

To restart an idle worker process, remove its lock file:

```
rm data/queue/workers/worker_id
```

A worker process that is busy processing a task will continue operation until its task is complete. Immediately afterwards, if the worker notices that its lock file is missing it exits.

If you have a recurring task to start workers, the recurring task starts a fresh worker, if necessary. See "Set up a recurring task to spawn workers" on page 220 for details.

Workflow global rules

This section provides information on how to enable workflow and configure the global workflow rules for Swarm.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

Workflow prerequisites

For Swarm 2019.2 and later, the workflow feature is enabled by default.

Tip

If you are upgrading from an earlier version you will need to update your triggers, see ["Upgrading Swarm" on page 241](#).

The workflow feature requires:

- Workflow must be enabled, see the ["workflow " below](#) system configurable.
- Your `swarm-trigger.conf` file must be configured correctly, see ["Using triggers to push events to Swarm" on page 204](#).
- Your Helix Core Server trigger table must contain the workflow trigger lines, see ["Using triggers to push events to Swarm" on page 204](#).

workflow

From Swarm 2019.2, `workflow` is enabled by default.

To disable the workflow feature, edit the `SWARM_ROOT/data/config.php` file and set the `enabled` configurable to `false` in the `workflow` section of the codeblock:

```
'workflow' => array(
    'enabled' => false, // Switches the workflow feature on.
    Default is true
),
```

The default value is `true`.

Tip

If you disable the workflow feature in the Swarm `config.php` file, workflow will not be processed by Swarm but a small overhead is still incurred by the Helix Core Server each time it runs a workflow trigger script. This overhead can be eliminated by commenting out the `swarm.enforce change-submit`, `swarm.strict change-content`, and `swarm.shelvesub shelf-submit` workflow triggers.

Global workflow

Note

- The **Global Workflow** page can be edited by a Swarm user with *super* or *admin* privileges, or users that have been made an owner.
- In earlier versions of Swarm, global workflow rules were set in the Swarm `config.php` file. From Swarm 2020.1, they are set in the global workflow on the Swarm **Workflows** page. If you upgrade from an earlier version of Swarm, your global workflow settings are automatically migrated to a workflow called **Global Workflow**.
- In earlier versions of Swarm, global tests were set in the Swarm `config.php` file. From Swarm 2020.1, tests are added to the global workflow on the Swarm **Workflows** page so that they operate as global tests. If you upgrade from an earlier version of Swarm, your global tests are automatically migrated. Each global test is migrated as follows: the owner is set as the Swarm admin user (not shared), the name is set to the original test name, the description is set to the original test title, and it is added to the **Global Workflow**.

Global workflow rules enable you apply rules globally if required, ensuring basic company policies are followed by every project. Each rule in the global workflow can be set with Enforce off (default) or Enforce on, this determines how that rule is used by Swarm:

- **Enforce off:** applies the workflow rule setting to projects and project branches that don't have an associated [workflow](#). If a project or project branch has an associated Swarm workflow, the global rule is ignored.
- **Enforce on:** applies a minimum workflow rule setting to all projects and project branches. If a project or project branch has an associated [workflow](#), the global rule is merged with the workflow rule and the most restrictive setting is used.

To edit the global workflow rules:

1. On the Swarm **Workflows** page.
2. Click **Global Workflow** to open the settings page:

Tip

By default, the global workflow is called **Global Workflow**, but it can be changed. To help you identify the global workflow it is always shown as the first workflow in the list and it has a different background color to the other workflows.

Global Workflow ✕

Name

Description

Description

Owners

✎ Add an Owner

Individuals:

swarm ✕

Global Rules [?]

		Enforce [?]
On commit without a review	Allow	<input type="checkbox"/>
On commit with a review	Reject unless approved	<input checked="" type="checkbox"/>
On update of a review in an end state	Reject	<input checked="" type="checkbox"/>
Count votes up from	Anyone	<input type="checkbox"/>
Automatically approve reviews	Never	<input type="checkbox"/>

Members or groups who can ignore ALL workflow rules

✎

Groups:

swarm-test ✕

Tests [?]	When	Blocks
Code sniff	On Update	Nothing ✎ ✖
Build tests	On Submit	Nothing ✎ ✖
Security check	On Update	Approve ✎ ✖

[+Add Test](#)

Save

Cancel



3. **Optional:** change the name if required, we recommend that you leave it set to **Global Workflow** if possible for ease of identification.
4. **Optional:** provide a description for the global workflow.
5. **Optional:** add more owners if required. This field auto-suggests groups, and users within Helix Core Server as you type (up to a combined limit of 20 entries).

Important

- The global workflow must have at least one owner.
- If you remove yourself as an owner, you cannot edit this workflow configuration later unless you have *super* user rights.

6. **Global Rules:**

Each rule in the global workflow can be set with Enforce off (default) or Enforce on using the **Enforce** slider next to the rule, this determines how that rule is used by Swarm:

-  **Enforce off:** applies the workflow rule setting to projects and project branches that don't have an associated [workflow](#). If a project or project branch has an associated Swarm workflow, the global rule is ignored.
-  **Enforce on:** applies a minimum workflow rule setting to all projects and project branches. If a project or project branch has an associated [workflow](#), the global rule is merged with the workflow rule and the most restrictive setting is used.

Important

- Changes to global workflow rules that are set with **Enforce off** will change the workflow for projects and project branches that do not have an associated workflow.
- Changes to global workflow rules that are set with **Enforce on** will change the workflow for all projects and project branches.

On commit without a review:

This rule is applied when a changelist without an associated review is submitted from outside of Swarm.

Select one of the following options:

- **Allow:** the changelist is not checked, the changelist is submitted without a review.
- **Create a review:** the changelist is submitted and a review is created automatically for the changelist.
- **Reject:** the changelist submit is rejected.

Tip

The selected rule is also applied when a changelist is submitted with **#review** in the description. For more information about creating a review by including a keyword in the changelist description, see "[Create a review](#)" on page 682.

On commit with a review: This rule is applied when:

- A Swarm review is committed.
- A changelist with an associated review is submitted from outside of Swarm.

Select one of the following options:

- **Allow:** the changelist review state is not checked. The changelist is committed even if its associated review is not approved.

- **Reject unless approved:** the changelist is only submitted if its associated review is approved and the content of the changelist is identical to the content of the approved review.

Tip

The selected rule is also applied when a changelist is submitted with `#review-nnnnn`, `#replace-nnnnn`, or `#append-nnnnn` in the description (**nnnnn** = review ID). For more information about adding a changelist to a review by including a keyword in the changelist description, see ["Add a changelist to a review" on page 682](#).

On update of a review in an end state:

Used to stop review content being automatically changed for reviews that are in specific states. By default, the protected end states are **Archived**, **Rejected**, and **Approve:Commit**. The end states are set by the Swarm administrator, see [end_states](#).

Tip

When a review is in a protected end state, it can still be updated manually by a user from the Swarm UI.

This rule is applied when a changelist is added to a review.

Select one of the following options

- **Allow:** the review is not checked. The changelist is added to the review no matter what the review state is. This is the default setting.
- **Reject:** if the review is in one of the states specified in the `end_state` configurable:
 - The **Add Change** button is disabled for the review.
 - The changelist is rejected if it is added outside of Swarm using `#review-nnnnn`, `#replace-nnnnn`, or `#append-nnnnn` in the description (**nnnnn** = review ID).

Count votes up from:

By default, all of the up votes on a review are counted for the **Minimum up votes** value set on the project/branch the review is associated with. Limit the up votes that are counted to just the members of the project the review is associated with by using this rule.

This rule is applied when a user votes on a review.

Select one of the following options:

- **Anyone:** votes are counted for all reviewers on a review. This is the default setting.
- **Members:** only the up votes of members of the project the review is associated with are counted for the **Minimum up votes** set on projects/branches.

Tip

For instructions on how to set **Minimum up votes** for projects and branches, see [Project minimum up votes](#) and [Branch minimum up votes](#).

Automatically approve reviews:

By default, reviews must be manually approved. Enable automatic approval of reviews with this rule.

This rule is applied when a user votes on a review.

Select one of the following options:

- **Never:** reviews are not automatically approved. This is the default setting.
- **Based on vote count:** reviews are automatically approved if:
 - There are no down votes on the review.
 - There are no moderators on the review. If a review has moderators it cannot be automatically approved.
 - All of the [Required reviewers](#) on the review have voted up.
 - The **Minimum up votes** on the review has been satisfied for **each** of the projects and branches the review spans.
 - All of the tests configured to block approval have passed.

Important

Moderators prevent the automatic approval of reviews. For more information about moderators, see ["Moderators" on page 477](#).

Tip

After a review has been automatically approved it needs to be manually committed.

Members or groups who can ignore ALL workflow rules

Optional: add users and groups to enable the users and groups (including any sub-groups) to ignore all of the workflow rules across the entire depot. This also ignores the `disable_commit` and `disable_self_approve` settings in the `SWARM_ROOT/data/config.php` file.

This field auto-suggests users, and groups within Helix Core Server as you type (up to a combined limit of 20 entries).

Typically, these permissions are given to a small set of trusted users and groups, for example project owners and administrators.

Tip

This applies to actions taken in the Swarm UI, the Helix Core command-line (P4), a P4D client, and the Swarm API.

7. Tests

Optional: add tests to the workflow, the tests are run when a review associated with the workflow is either created/updated or submitted. Selected from the **When** dropdown for the test.

Tip

- Tests are only saved on the workflow when you click the **Save** button.
- If a review is associated with multiple workflows, all of the tests on all of those workflows are run for the review. If the same test is on more than one of the associated workflows, it is only run once.

Add a test:

- a. Click **+Add Test** in the Tests area.
- b. Select the test to run from the **Tests** dropdown list.
- c. Select when you want the test to run from the **When** dropdown list.

Tip

- Review content change is when the files or content of files in a review change. A change to the review description does not trigger a test.
- **On Demand** tests are unaffected by the review content change check.


- **On Update** the test will run when:

- a review is created.
- a review is submitted and the review content has changed compared to the previous review version.
- a review is updated and the review content has changed since the previous review version.



Tip

If the review content has not changed since a test last successfully reported a pass, it will not be re-run when a review is versioned or submitted. If a test is still running or has not updated Swarm with a pass state, it will be re-run, even on a submit.

- **On Submit** the test will run when a pre-commit review is submitted or a post-commit review is created.
- **On Demand** the test will not run automatically, but you can run it manually from the ["Tests" on page 451](#) dialog on the review page.

- d. Select whether Swarm blocks approval if the test fails from the **Blocks** dropdown list:
 - **Nothing** if the test fails it will not block approval.
 - **Approve** if the test fails it will prevent the review from being approved.
- e. Click the **Accept**  button to add the test to the workflow.

Edit a test on the workflow:

- a. Click the **Edit**  button next to the test you want to edit.
- b. The **Test**, **When**, and **Blocks** dropdown lists are displayed for the test.
- c. Make changes as required.
- d. Click the **Accept**  button to confirm your changes.

Remove a test from the workflow:

Click the **Delete**  button next to the test to remove it from the workflow.

The test is removed immediately, no confirmation is requested. You must save the workflow to complete the test removal.

8. Click **Save**.

Note

The **Save** button is disabled if any required fields are empty.

12 | Extending Swarm

Helix Swarm is built with open source technologies, using modern development methods, resulting in a platform that is capable, modular, and can be extended in a variety of ways. While this section is currently incomplete, it does identify the technologies that make up Swarm and provides notes regarding some of the available extension points.

Resources	761
Development mode	762
Modules	763
CSS & JavaScript	788

Resources

Swarm is built with a number of different technologies. As such it can be a bit of a challenge to ramp up on if you haven't done web development before. This section lists some of the best resources we've found for learning how to develop for Swarm.

jQuery

A free three hour course on jQuery basics. Highly recommended if you haven't used jQuery before. It also covers some JavaScript and CSS basics:

<http://try.jquery.com>

JavaScript resources

A free, in-browser JavaScript course:

<https://www.w3schools.com/js/default.asp>

PHP resources

A free PHP tutorial:

<https://www.w3schools.com/php7/>

Laminas Framework resources

Note

Swarm supports the Laminas component versions in the `LICENSE.txt` file, features and functions in the Laminas documentation that were introduced in later versions of Laminas will not work with Swarm. The `LICENSE.txt` file is in the `readme` folder of your Swarm installation.

- The documentation overview page for Laminas Framework:

<https://docs.laminas.dev/>

- Laminas tutorial that covers most of the basics:

<https://docs.laminas.dev/tutorials/>

- Laminas Quick Start documentation:

<https://docs.laminas.dev/laminas-mvc/quick-start/>

Development mode

Swarm has a *development* mode that, when enabled, changes Swarm's behavior in the following ways:

- CSS and JavaScript code is not aggregated. Each CSS or JavaScript file is fetched individually, which makes it much easier to identify where styles or functions exist and how they apply to Swarm. Due to the additional HTTP requests, development mode should not be used in production environments.
- Errors and exceptions that may occur are displayed in Swarm's UI. This is particularly useful to developers of Swarm modules. As the error information might disclose system paths or configuration, development mode should not be used in production environments.

Tip

If you make a configuration change, Swarm will not use it until the configuration cache has been reloaded, this forces Swarm to use the new configuration. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the "Reload Configuration button" on page 732.

To enable development mode:

Adjust the `SWARM_ROOT/data/config.php` file to include:

```
<?php
    return array(
        'p4' => ...
        'environment' => array(
```

```

        'mode' => 'development'
    )
);

```

To disable development mode:

Adjust the `SWARM_ROOT/data/config.php` file to exclude the `'mode' => 'development'` line.

Modules

You can design and implement your own custom Swarm modules to modify the behavior of Swarm. This section covers the design and implementation of Swarm custom modules.

Important

Editing custom modules from the Swarm UI will not work once the new review page replaces the original review page in a later Swarm release.

Important

The operation and testing of custom modules is the responsibility of the module creator. Perforce Software, Inc. is not responsible for the operation of your custom modules, nor does Perforce Software, Inc. make any representations or warranties regarding the interoperability of your custom modules with Helix Swarm.

Tip

You must test your custom modules on a test system before transferring them to your production system. This avoids any negative impact on the operation of your production system. If you have more than one custom module, the modules should all be tested at the same time on the same test system as this ensures that the modules operate correctly with each other and with Helix Swarm.

Tip

If you add or edit a module, Swarm will not use that change until the config cache has been reloaded, this forces Swarm to use the module change. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

Module overview

Important

The operation and testing of custom modules is the responsibility of the module creator. Perforce Software, Inc. is not responsible for the operation of your custom modules, nor does Perforce Software, Inc. make any representations or warranties regarding the interoperability of your custom modules with Helix Swarm.

Tip

You must test your custom modules on a test system before transferring them to your production system. This avoids any negative impact on the operation of your production system. If you have more than one custom module, the modules should all be tested at the same time on the same test system as this ensures that the modules operate correctly with each other and with Helix Swarm.

Tip

If you add or edit a module, Swarm will not use that change until the config cache has been reloaded, this forces Swarm to use the module change. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

A Swarm custom module is created in a folder in the **modules** folder of your Swarm installation. The folder name must match the custom module name and must contain at least a **Module.php** file and a **module.config.php** file. A custom module will not work in Swarm until you enable it. Enable your custom modules in Swarm by adding them to the **custom.modules.config.php** file in the **config** directory (a peer of the **modules** directory).

This chapter provides basic information about how custom modules integrate with Swarm and assumes a basic knowledge of the Laminas framework. For more information about the Laminas framework, see [Laminas Quick Start documentation](#).

This section includes:

- ["Migrate existing custom modules to the Laminas framework" on the facing page](#)
- ["Influence activity events, emails, etc." on page 766](#)
- ["Templates" on page 767](#)
- ["Custom module file locations" on page 768](#)
- ["IndexControllers" on page 768](#)
- ["View helpers" on page 769](#)
- ["Event listeners" on page 769](#)
- ["Enabling your custom modules" on page 772](#)
- ["Further reading" on page 773](#) , includes links to example custom modules and the official Laminas documentation

Migrate existing custom modules to the Laminas framework

Swarm 2020.1 and later uses the Laminas framework, previous Swarm versions used the Zend framework. The move to Laminas is part of our commitment to move away from using versions of platforms that have reached End-of-Life (EOL).

If you have any custom Swarm modules that were created for Swarm 2019.3 or earlier, you must modify them so that they will work in Swarm 2020.1 and later. The modifications required depend on the version of Swarm you were using the custom modules with before your upgraded Swarm:

- "Custom modules on Swarm 2019.1, 2019.2, or 2019.3 " below
- "Custom modules on Swarm 2018.3 and earlier" below

Custom modules on Swarm 2019.1, 2019.2, or 2019.3

Note

Migrate your custom modules on a test system before transferring them to your production system. This avoids any negative impact on the operation of your production system. If you have more than one custom module, the modules should all be tested at the same time on the same test system as this ensures that the modules operate correctly with each other and with Helix Swarm.

Custom modules used with Swarm 2019.1, 2019.2, and 2019.3 use the Zend 3 framework, they must be migrated to the Laminas framework before they can be used with Swarm 2020.1 and later:

1. Migrate your Zend 3 custom modules to the Laminas framework by pointing the Laminas migration tool at the folder you store your custom modules in. For the migration tool and full instructions on migrating your custom Zend 3 modules to Laminas, see <https://docs.laminas.dev/migration/>.
2. Swarm will not use the custom modules until the Swarm config cache has been deleted. For instructions on deleting the config cache, see "[Swarm config cache file delete](#)" on page 813.
3. Check that your custom modules work as expected before putting them on your production system.

Custom modules on Swarm 2018.3 and earlier

Note

Update and migrate your custom modules on a test system before transferring them to your production system. This avoids any negative impact on the operation of your production system. If you have more than one custom module, the modules should all be tested at the same time on the same test system as this ensures that the modules operate correctly with each other and with Helix Swarm.

Custom modules used with Swarm 2018.3 and earlier use the Zend 2 framework, they must be migrated to the Laminas framework before they can be used with Swarm 2020.1 and later:

1. Convert your Zend 2 custom modules to the Zend 3 framework, see [Upgrade custom modules to work with Zend 3](#).

Tip

Laminas is based on the Zend 3 framework, this makes it is relatively simple to convert your Zend 3 custom modules to the Laminas framework:

- If you only have one or two custom modules, you can manually change all of the **Zend** references in your modules to **Laminas** while you are converting them from Zend 2. This means you can skip the Laminas migration step.
- If you have a larger number of custom modules, it is probably quicker to automatically change all of the references in your custom modules by running the Laminas migration tool as described in the next step.

2. Migrate your Zend 3 custom modules to the Laminas framework by pointing the Laminas migration tool at the folder you store your custom modules in. For the migration tool and full instructions on migrating your custom Zend 3 modules to Laminas, see <https://docs.laminas.dev/migration/>.
3. Swarm will not use the custom modules until the Swarm config cache has been deleted. For instructions on deleting the config cache, see ["Swarm config cache file delete" on page 813](#).
4. Check that your custom modules work as expected before putting them on your production system.

Influence activity events, emails, etc.

When something happens in Helix Core Server (change submitted, files shelved, job added/edited), or state changes within Swarm (comment added, review state changed, etc.), the event is pushed onto the Swarm task queue. A background worker process pulls events off of the queue and publishes an event alerting modules about activity they may be interested in processing. This architecture allows the Swarm user interface to be fairly quick while accommodating tasks that might require notable processing time, or need to wait for related information to become available.

Subscribers to the worker event flesh the item out (fetch the change/job details, for example) and indicate if it should result in an activity entry, email notification, etc. By subscribing to various event topics, your custom module can pay attention to specific types of events. While your custom module is processing an event, it can modify the text of activity events, change the contents of emails, drop things entirely from activity, etc.

When your custom module subscribes to an event, set the priority to influence how early or late in the process it runs. You will likely want your module to run after most other modules have done their work to flesh out the event, but before Swarm's events module processes it. The events module is a good example of subscribing to these events:

```
swarm_install/module/Events/config/module.config.php
```

Tip

Note that its priority is set to `-100`. Select a value before that for your own module (for example, `0` would be neutral and `-90` would indicate that you are interested in being last).

Event priority is from the highest positive number to the lowest negative number with `0` in the middle. For a full list of tasks and their event priorities, see ["Task details" on page 773](#).

Task types

For details of each of the following tasks, see ["Task details" on page 773](#).

- `task.commit`
- `task.shelve`
- `task.review`
- `task.change`
- `task.mail`
- `task.cleanup.attachment`
- `task.cleanup.archive`
- `task.shelvedel`
- `task.changesave`
- `task.changesaved`
- `task.comment`
- `task.comment.batch`
- `task.commentSendDelay`
- `task.group`
- `task.groupdel`
- `task.job`
- `task.user`
- `task.userdel`
- `task.workflow.created`
- `task.workflow.updated`

Templates

Override existing view templates using your custom module. Have a look at this [example module](#) that demonstrates how to customize the email templates Swarm uses for comment notifications.

For more information about views, see the [Laminas-View Quick Start](#).

Note

Swarm supports the Laminas component versions in the **LICENSE.txt** file, features and functions in the Laminas documentation that were introduced in later versions of Laminas will not work with Swarm. The **LICENSE.txt** file is in the **readme** folder of your Swarm installation.

Custom module file locations

The custom module files must be stored using the PSR-4 standard locations shown below:

```
config/
    custom.modules.config.php (required)
module/
    ModuleName/ (required)
        config/ (required)
            module.config.php (required)
        src/ (php files here, required)
        Controller/
            MyIndexController.php
        Listener/ (optional)
            MyListener.php (optional)
        ...
        other directories (non-php files here, optional)
        ...
    Module.php (required)
```

IndexControllers

Create your **IndexControllers** using factories. The **IndexControllers** should extend **Application\Controller\AbstractIndexController** and use the Swarm **IndexControllerFactory**, this means that the **IndexControllerFactory** can automatically inject the services .

Tip

We recommend you follow the `::class` description to avoid errors in string representation.

```
'controllers' => array(
    'factories' => [
        MyModuleName\Controller\IndexController::class =>
        Application\Controller\IndexControllerFactory::class,
```

```
],
),
```

View helpers

Create your view helpers using factories. The factories that create the view helpers must inject any services that are required.

Tip

We recommend using `::class` rather than arbitrary strings as this can limit the possibility of errors. The use of classes also makes your factory simpler because reflection can be used to construct the helpers if they all extend a common parent class.

```
'view_helpers' => array(
    'factories' => array_fill_keys(
        [
            MyModuleName\View\Helper\Helper1::class,
            MyModuleName\View\Helper\Helper2::class,
        ],
        <YourHelperFactory>::class
    )
)
```

Set options on existing helpers

It is possible to influence the behavior of existing view helpers by setting options on them; for an example see: [swarm_install/module/Application/Module.php](#).

Register new helpers

It is also possible to register new view helpers by placing them within your module's hierarchy, for example, `MyModule/src/View/Helper/Foo.php`. Use the following Swarm view helper for inspiration: [swarm_install/module/Activity/src/View/Helper/Activity.php](#).

Then register your view helper with the view manager via your `ModuleConfig`: `swarm_root/module/Activity/config/module.config.php`.

Event listeners

Your event listeners should extend `Events\Listener\AbstractEventListener`. This means that the `ListenerFactory` can create them without you having to write your own factory.

The `Listener.php` file contains the implementation of your event listener and the `module.config.php` file configures your event listeners.

Tip

- Look at the code in the `AbstractEventListener.php` file in `module/events/src/Listener` to see the functions that are available to you.
- We strongly recommend that you create your event listeners to use the declarative model because it has a number of advantages over the non-declarative model:
 - `AbstractEventListener` provides common code for listeners to use
 - The declarative model offers better debug options (logging)
 - The declarative model is neater, easier to read, easier to support, and easier to maintain

Example `module.config.php` file that uses the declarative model for event listeners:

Tip

`Listener::class` example details:

- `Events\Listener\ListenerFactory::ALL => [`
 We are listening to **ALL** here for convenience because we are listening for mail events but this means it will trigger on every event. Usually is better to just listen for the events you are interested in. For example, if you are interested in commits and shelves you would listen on **COMMIT** and **SHELVE** events.
- `Events\Listener\ListenerFactory::PRIORITY => -199,`
 Declares an event priority of **-199** for the event listener because email delivery events are processed with a priority of **-200** and the example event needs to run just before the email delivery event.
- `Events\Listener\ListenerFactory::CALLBACK => 'handleEmail',`
 Declares the function name within the listener class that is called.
- `Events\Listener\ListenerFactory::MANAGER_CONTEXT => 'queue'`
 Triggers your custom listener when Swarm processes the Swarm event queue.

```
<?php
/**
 * Perforce Swarm
 *
 * @copyright 2019 Perforce Software. All rights reserved.
 * @license   Please see LICENSE.txt in top-level folder of this
distribution.
 * @version   <release>/<patch>
 */
```

```

$listeners = [EmailExample\Listener\Listener::class];
return [
    'listeners' => $listeners,
    'service_manager' =>[
        'factories' => array_fill_keys(
            $listeners,
            Events\Listener\ListenerFactory::class
        )
    ],
    Events\Listener\ListenerFactory::EVENT_LISTENER_CONFIG => [
        Events\Listener\ListenerFactory::ALL => [
            EmailExample\Listener\Listener::class => [
                [
                    Events\Listener\ListenerFactory::PRIORITY => -199,
                    Events\Listener\ListenerFactory::CALLBACK =>
'handleEmail',
                    Events\Listener\ListenerFactory::MANAGER_CONTEXT
=> 'queue'
                ]
            ]
        ]
    ]
];

```

Example of an email Listener.php file:

```

namespace MyModuleName\Listener;

use Events\Listener\AbstractEventListener;
use Laminas\EventManager\Event;

class Listener extends AbstractEventListener
{
    public function handleEmail(Event $event)
    {
        $mail = $event->getParam('mail');
    }
}

```

```

        if (!$mail || !isset($mail['htmlTemplate'], $mail
['textTemplate'])) {
            return;
        }
    }
}
}

```

Enabling your custom modules

Swarm uses the **custom.modules.config.php** file to check which custom modules it should load. This gives you control over which modules Swarm loads and prevents modules from being loaded by mistake.

Create the **custom.modules.config.php** file in the **config** directory (a peer of the **modules** directory) if it does not already exist. Edit the file so that it contains the following details for each of your modules:

- **namespaces** array: enter your custom module names and paths
- **return** array : enter your custom module names

For example, if you have three modules called **MyModuleName**, **AnotherModuleName**, and **NewModuleName** the file content would look similar to:

```

<?php
\Laminas\Loader\AutoloaderFactory::factory(
    array(
        'Laminas\Loader\StandardAutoloader' => array(
            'namespaces' => array(
                'MyModuleName'          => BASE_PATH .
'/module/MyModuleName/src',
                'AnotherModuleName'    => BASE_PATH .
'/module/AnotherModuleName/src',
                'NewModuleName'        => BASE_PATH .
'/module/NewModuleName/src',
            )
        )
    )
);
return [
    'MyModuleName',
    'AnotherModuleName',

```



```
'NewModuleName',
];
```

Further reading

For detailed information about the Laminas framework and examples see the following:

- The [Laminas Quick Start documentation](#) and [Laminas framework documentation portal](#), it is useful to have a basic knowledge of the Laminas framework before you create your own modules. .

Note

Swarm supports the Laminas component versions in the **LICENSE.txt** file, features and functions in the Laminas documentation that were introduced in later versions of Laminas will not work with Swarm. The **LICENSE.txt** file is in the **readme** folder of your Swarm installation.

- "Example email module" on page 776, a simple custom module that makes Swarm use a custom email template when it sends out when a comment notification is sent out.
- Swarm JIRA module, a simple module implementation within Swarm: `swarm_install/module/Jira`

Task details

Swarm has lots of events that can be listened to and it uses these to process the tasks to generate data for Swarm.

This section lists the Swarm tasks available for use with your modules. They are listed along with the priorities of the task events. Events for each task are listed in highest to lowest priority order.

task.commit

- **Priority 200:** the listener searches for all of the projects that are impacted by the change and checks for any user mentions that need to be linked. It sets up the email list here and queues the activity to be processed.
- **Priority 100:** the listener determines whether it should update or create any reviews for the changelist.
- **Priority -100:** the listener checks if any users have the Review daemon setup and attaches those users to the email to list.
- **Priority -200:** has two Listeners:
 - One listener checks the `disable_change_emails` configurable setting: `true`, the email will not be sent.

false, Swarm adds any reviews that might not be part of the *to list* because they are not a member of the project.

- The other listener creates the activity for all users/projects that are linked to this event.
- **Priority -300**: the listener sends out the email that is related to the committed change.
- **Priority -400**: the listener links to JIRA for this change list.

task.shelve

- **Priority 100**: the listener processes the shelve task to determine whether it should create or updates any reviews.
- **Priority -200**: the listener creates the activity for all users/projects that are linked to this event.
- **Priority -300**: the listener sends out the email that is related to the shelved change.

task.review

- **Priority 100**: the listener processes the review and updates any records that are linked to the review. For example, @mention, new author, new participants, activity, and requirement level.
- **Priority -200**: the listener creates the activity for all users/project that are linked to this event.
- **Priority -300**: has two Listeners:
 - One listener sends out the email that is related to the review change.
 - The other listener fetches any associated changes and ensures they are linked to JIRA.

task.change

Priority 1: the listener ensures the change and any mentions in the description are linked to JIRA issues.

task.mail

Priority 1: this event is manually triggered from the **review task**. Other events do not manually trigger this task, they wait for it to run at the end. Once we are here Swarm will validate all of the email settings and send the email out.

task.cleanup.attachment

Priority 1: the listener will delete the attachment if a cleanup attachment task has been queued.

task.cleanup.archive

Priority 1: the listener will clean up any archive files if their life time has expired.

task.shelvedel

- **Priority 1:** the listener processes any changelist that has had files deleted from it.
- **Priority -200:** the listener creates the activity for all users/projects that are linked to this event.

task.changesave

Priority 1: the listener processes the change to ensure that it is valid and puts a future task in the queue to do the `task.changesaved` tasks.

task.changesaved

Priority 1: the listener takes the changelist description and updates the review description if `sync description` is enabled.

task.comment

- **Priority 1:** the listener processes the comment being created.
- **Priority -200:** the listener creates the activity for all users/projects that are linked to this event.
- **Priority -300:** the listener sends out an email related to the committed change.

task.comment.batch

- **Priority 1:** If comments are being batched, the listener will process the comment(s) and put a task into the queue to be processed when the `notification_delay_time` has elapsed.
- **Priority -200:** the listener creates the activity for all users/projects that are linked to this event.
- **Priority -300:** the listener sends out an email related to the committed change.

task.commentSendDelay

- **Priority 1:** if the user or system has forced a batched comment to be sent the listener will setup the email here.
- **Priority -200:** the listener creates the activity for all users/projects that are linked to this event.
- **Priority -300:** the listener sends out an email related to the committed change.

task.group

Priority 1: the listener invalidates the group cache if the group is changed.

task.groupdel

Priority 1: the listener invalidates the group cache if the group has been deleted.

task.job

- **Priority 1:** the listener sets up the Activity for the job and attaches any users/projects to the activity. It also checks if any users want to be notified about Job changes.
- **Priority -200:** the listener creates the activity for all users/projects that are linked to this event.
- **Priority -300:** has two Listeners:
 - One listener sends out an email related to the committed change.
 - The other listener fetches associated changes and ensures they are linked to JIRA.

task.user

Priority 1: the listener invalidates the user cache if the user is changed.

task.userdel

Priority 1: the listener invalidates the user cache if the user has been deleted.

task.workflow.created

- **Priority 1:** the listener sets up the Activity for the new workflow and attaches any users/projects to the activity.
- **Priority -200:** the listener creates the activity for all users/projects that are linked to this event.

task.workflow.updated

- **Priority 1:** the listener sets up the Activity for the updated workflow and attaches any users/projects to the activity.
- **Priority -200:** the listener creates the activity for all users/projects that are linked to this event.

Example email module

The following example module demonstrates how to customize the email template Swarm uses when sending notifications for comments.

Note

Swarm supports the Laminas component versions in the **LICENSE.txt** file, features and functions in the Laminas documentation that were introduced in later versions of Laminas will not work with Swarm. The **LICENSE.txt** file is in the **readme** folder of your Swarm installation.

Tip

You must test your custom modules on a test system before transferring them to your production system. This avoids any negative impact on the operation of your production system. If you have more than one custom module, the modules should all be tested at the same time on the same test system as this ensures that the modules operate correctly with each other and with Helix Swarm.

Tip

If you add or edit a module, Swarm will not use that change until the config cache has been reloaded, this forces Swarm to use the module change. You must be an *admin* or *super* user to reload the Swarm config cache. Navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button"](#) on page 732.

Basic steps needed to create the Email Example module are:

- ["Create the Module.php file" on the next page](#)
- ["Create the module.config.php file" on page 779](#)
- ["Create the Listener.php file" on page 781](#)
- ["Create the comment-html.phtml file" on page 784](#)
- ["Create the comment-text.phtml file" on page 786](#)
- ["Enable the EmailExample module for Swarm in custom.modules.config.php" on page 787](#)

Further reading:

["Extending the EmailExample module to other email templates" on page 788](#)

File locations

For reference, the EmailExample module uses the following filenames and locations:

```
config/
    custom.modules.config.php
module/
    EmailExample/
        config/
            module.config.php
        src/
            Listener/
                Listener.php
        view/
            mail/
                commit-html.phtml
```

```

        commit-text.phtml
        comment-html.phtml
        comment-text.phtml
        review-html.phtml
        review-text.phtml
    Module.php

```

Create the Module.php file

Module.php will:

- Declare the module namespace, this must match the directory name of the module
- Provide the `getConfig()` function

Tip

Optional: You can also implement the `onBootstrap (Event $event)` function if you want to do some early setup when the module is loaded:

```

public function onBootstrap(Event $event)
{
}

```

Event is a Laminas class and is added after namespace:

```

namespace EmailExample
use Laminas\EventManager\Event

```

Create the Module.php file:

1. Create a directory called **EmailExample** in the module directory.
2. Create the file **Module.php** in **module/EmailExample**.
3. Edit **Module.php** to contain:

```

<?php
/**
 * Performer Swarm
 *
 * @copyright 2019 Performer Software. All rights reserved.
 * @license   Please see LICENSE.txt in top-level folder of
this distribution.
 * @version   <release>/<patch>

```

```

*/

namespace EmailExample;

class Module
{
    public function getConfig()
    {
        return include __DIR__ . '/config/module.config.php';
    }
}

```

4. Now "Create the module.config.php file" below.

Create the module.config.php file

The `module.config.php` file will:

- Configure your module including event listeners
- Declare an event priority of `-199` for the event listener because email delivery events are processed with a priority of `-200` and the example event needs to run just before the email delivery event.

Tip

`Listener::class` example details:

- `Events\Listener\ListenerFactory::ALL => [`
We are listening to **ALL** here for convenience because we are listening for mail events but this means it will trigger on every event. Usually is better to just listen for the events you are interested in. For example, if you are interested in commits and shelves you would listen on **COMMIT** and **SHELVE** events.
- `Events\Listener\ListenerFactory::PRIORITY => -199,`
Declares an event priority of `-199` for the event listener because email delivery events are processed with a priority of `-200` and the example event needs to run just before the email delivery event.
- `Events\Listener\ListenerFactory::CALLBACK => 'handleEmail',`
Declares the function name within the listener class that is called.
- `Events\Listener\ListenerFactory::MANAGER_CONTEXT => 'queue'`
Triggers your custom listener when Swarm processes the Swarm event queue.

Create the `module.config.php` file:

1. Create a directory called **config** in the **EmailExample** directory.
2. Create a file called **module.config.php** in **config**.
3. Edit **module.config.php** to contain:

```
<?php
/**
 * Perforce Swarm
 *
 * @copyright 2019 Perforce Software. All rights reserved.
 * @license   Please see LICENSE.txt in top-level folder of
this distribution.
 * @version   <release>/<patch>
 */

$listeners = [EmailExample\Listener\Listener::class];
return [
    'listeners' => $listeners,
    'service_manager' =>[
        'factories' => array_fill_keys(
            $listeners,
            Events\Listener\ListenerFactory::class
        )
    ],
    Events\Listener\ListenerFactory::EVENT_LISTENER_CONFIG => [
        Events\Listener\ListenerFactory::ALL => [
            EmailExample\Listener\Listener::class => [
                [
                    Events\Listener\ListenerFactory::PRIORITY =>
-199,
                    Events\Listener\ListenerFactory::CALLBACK =>
'handleEmail',
                    Events\Listener\ListenerFactory::MANAGER_
CONTEXT => 'queue'
                ]
            ]
        ]
    ]
]
```



```
        ]  
    ]  
];
```

4. Now "Create the Listener.php file" below.

Create the Listener.php file

The **Listener.php** file will:

- Contains the implementation of your event listener
- Allow logging

Create the Listener.php file:

1. Create a directory called **src** in the **EmailExample** directory.
2. Create a directory called **Listener** in the **src** directory.
3. Create a file called **Listener.php** in **Listener**.
4. Edit **Listener.php** to contain:

```
<?php
/**
 * Perforce Swarm
 *
 * @copyright 2019 Perforce Software. All rights reserved.
 * @license   Please see LICENSE.txt in top-level folder of
this distribution.
 * @version   <release>/<patch>
 */

namespace EmailExample\Listener;

use Events\Listener\AbstractEventListener;
use Laminas\EventManager\Event;

class Listener extends AbstractEventListener
{
    /**
     * Automatically uses any custom email templates found under
this
     * module's view/mail folder (e.g. Example/view/mail/commit-
html.phtml).
     *
     * Valid templates include:
     *
     * commit-html.phtml (HTML version of commit notification)
     * commit-text.phtml (text version of commit notification)
     * comment-html.phtml (HTML version of comment notification)
     * comment-text.phtml (text version of comment notification)
     * review-html.phtml (HTML version of review notification)
     * review-text.phtml (text version of review notification)
     *
     * Note: you need to provide custom templates for both HTML
and text;
```

```

    * if you do not provide both, it is possible that the search
for
    * customized templates only finds the non-customized
versions, making
    * it appear that this module is not working.
    */

/**
 * Handle the Email and set the new templates.
 *
 * @param Event $event
 */
public function handleEmail(Event $event)
{
    $logger = $this->services->get('logger');
    $logger->info("EmailExample: handleEmail");
    $mail = $event->getParam('mail');
    if (!$mail || !isset($mail['htmlTemplate'], $mail
['textTemplate'])) {
        return;
    }

    $html = __DIR__ . '/view/mail/' . basename($mail
['htmlTemplate']);
    $text = __DIR__ . '/view/mail/' . basename($mail
['textTemplate']);

    if (file_exists($html)) {
        $mail['htmlTemplate'] = $html;
    }
    if (file_exists($text)) {
        $mail['textTemplate'] = $text;
    }
}

```

```

        $event->setParam('mail', $mail);
        $logger->info("EmailExample: handleEmail end.");
    }
}

```

5. Now "Create the comment-html.phtml file" below.

Create the comment-html.phtml file

The **comment-html.phtml** file is a view script that provides the content for the HTML portion of the comment notification email.

Tip

It is considered best practice to use inline CSS for styling emails.

Create the comment-html.phtml file:

1. Create a directory called **view** in the **module/Example** directory.
2. Create a directory called **mail** in the **module/Example/view** directory.
3. Create the **comment-html.phtml** file in **module/Example/view/mail**.
4. Edit **comment-html.phtml** to contain:

```

<?php
    $user          = $activity->get('user');
    $userLink      = $user
                    ? $this->qualifiedUrl('user', array('user' =>
$user))
                    : null;
    $targetLink = $activity->getUrl($this->plugin
('qualifiedUrl'));
?>
<html>
    <body style="font-family: sans-serif; background-color: #eee;
padding: 1em;">
        <div style="background-color: #fff; border: 1px solid #ccc;
padding: 1em;">
            <div style="font-size: 115%;">
                <?php if ($user): ?>
                    <a style="text-decoration: none;" href="<?php echo

```

```

$userLink ?>">
    <?php echo $this->escapeHtml($user) ?>
    </a>
    <?php endif; ?>
    <?php echo $this->escapeHtml($activity->get('action')) ?>
    <a style="text-decoration: none;" href="<?php echo
$targetLink ?>">
        <?php echo $this->escapeHtml($activity->get
('target'))?>
    </a>
</div>
<br/>
<?php
    // if the comment has file context, show it.
    $comment = $event->getParam('comment');
    $context = $comment
        ? $comment->getFileContext()
        : array('content' => null, 'line' => null);
    if (is_array($context['content']) && $context['line'])
    {
        $line = $context['line'] - count($context
['content']) + 1;
        echo '<div style="font-family: monospace; white-
space: nowrap;'
            . ' padding: .5em 1em; overflow-x: auto; color:
#444;'
            . ' border: 1px solid #ddd; background-color:
#f7f7f7;">';
        foreach ((array) $context['content'] as $i =>
$content) {
            echo '<div><span style="color: #999;">'
                . str_pad($line + $i,
                    strlen($context['line']),
                    "0",
                    STR_PAD_LEFT

```

```
        )
        . ' </span>&nbsp;'
        . $this->preformat($content)
            ->setLinkify(false)
            ->setEmojify(false)
            ->setWordWrap(900)
        . "</div>\n";
    }
    echo '</div><br/>';
}
?>
<div style="padding-bottom: .5em;">
<?php
    echo $this->preformat($activity->get('description'))
        ->setBaseUrl($this->qualifiedUrl())
        ->setEmojify(false)
        ->setWordWrap(900)
?>
</div>
</div>
</body>
</html>
```

5. Now ["Create the comment-text.phtml file"](#) below.

Create the comment-text.phtml file

The **comment-text.phtml** is a view script that provides the content for the text-only portion of the comment notification email.

Create the comment-text.phtml file:

1. Create the file **comment-text.phtml** in the **module/Example/view/mail** directory.
2. Edit **comment-text.phtml** to contain:

```
<?php
    echo trim($activity->get('user'))
        . ' commented on '
        . $activity->get('target'));
```

```

// if the comment has file context, show it.
$comment = $event->getParam('comment');
$context = $comment
    ? $comment->getFileContext()
    : array('content' => null);
if (is_array($context['content'])) {
    echo "\n\n> " . $this->wordWrap(
        implode("\n> ", $context['content']), 900
    );
}

echo "\n\n" . trim($this->wordWrap($activity->get(
    'description'), 900));
echo "\n\n" . $activity->getUrl($this->plugin(
    'qualifiedUrl'));
?>

```

3. Now ["Enable the EmailExample module for Swarm in custom.modules.config.php"](#) below.

Enable the EmailExample module for Swarm in custom.modules.config.php

Swarm uses the `custom.modules.config.php` file to auto-load classes and to check which custom modules it should run. This gives you control over which modules Swarm loads and prevents modules from being loaded by mistake.

Create the `custom.modules.config.php` file:

1. Create the `config` directory at the same level as the `module` directory if it does not exist.
2. Create the `custom.modules.config.php` file in the `config` directory if it does not exist.
3. Edit the `custom.modules.config.php` file so that it contains the auto-loader and the **EmailExample** module details:

Tip

If you already have one or more custom modules enabled for Swarm, the auto-loader and the existing module details will already be in the file.

Just add **EmailExample** to the `namespaces` and `return` arrays of the `custom.modules.config.php` file.

```
<?php
\Laminas\Loader\AutoloaderFactory::factory(
    array(
        'Laminas\Loader\StandardAutoloader' => array(
            'namespaces' => array(
                'EmailExample' => BASE_PATH .
                '/module/EmailExample/src',
            )
        )
    )
);
return [
    'EmailExample'
];
```

4. The Swarm config cache must be reloaded so that Swarm can see your new module. As an *admin* or *super* user, navigate to the **User id** dropdown menu, select **System Information**, click the **Cache Info** tab, and click the ["Reload Configuration button" on page 732](#).

The Email Example module is now enabled for Swarm. Swarm will use the new custom email template whenever a comment notification is sent out.

5. Check that the module works correctly before moving it to your production server.

Extending the EmailExample module to other email templates

If you need to customize any other types of Swarm notification email messages, locate the view scripts (both HTML and text) and copy them into `module/Example/view/mail`, maintaining the existing filenames, then modify the new files as desired.

Note

If you do not copy both the HTML and text templates, it is possible for the search for customized templates to only find non-customized versions, making it appear that your module is not working.

CSS & JavaScript

Custom CSS and JavaScript files will be loaded automatically, following a browser refresh, if you place them under either:

- `SWARM_ROOT/public/custom/*.css|js`
- or
- `SWARM_ROOT/public/custom/sub-folder/*.css|js`

Note

- Swarm only supports customizations placed directly within the `SWARM_ROOT/public/custom` folder or one sub-folder down. Customizations are added after all standard CSS and JavaScript. If more than one custom file is present, they are added in alphabetical order.
- Before creating any customizations, make sure that the `SWARM_ROOT/public/custom` folder exists. Swarm does not ship with or create this folder.
- When you add custom CSS or JavaScript to Swarm, only a browser refresh is required. You do not need to restart `apachectl`.
- If you have checked the custom CSS or JavaScript code for accuracy and it still doesn't work after a browser refresh, check the folder permissions for the folder the customization is in.

Sample JavaScript extensions

The following is an example JavaScript customization that you might wish to apply to your Swarm installation. JavaScript customizations can be implemented separately, but ideally, you would apply the JavaScript customizations in a single file to reduce the number of web requests required.

Warning

Coding errors in your custom JavaScript files could cause the Swarm UI to stop working. If this occurs, use your browser's development tools to identify which file contains the problem, and move that file out of the `SWARM_ROOT/public/custom` folder. When the problem has been resolved, the file can be returned.

Add text to the Log in dialog

This customization can only be used if `require_login` is set to `true`, see ["Require login" on page 698](#).

```
// Custom js to put text into the login box
$(function() {
    setTimeout(function () {
        if ($('.login-form')) {
            $('.form-body')
                .prepend("<b style='padding-bottom: 20px; display: inline-block'>Please use your Helix server login credentials.</b>");
        }
    })
})
```

```
    },
    500);
});
```

Replace the *Please use your Helix Core Server login credentials* with the text you want to add to the top of the Swarm log in dialog.

Save this line in a file, perhaps `customize-login-text.js`, within the `SWARM_ROOT/public/custom`. Swarm automatically includes the JavaScript file, adding the text to the top of the log in dialog immediately for subsequent log ins.

Sample CSS customizations

The following are example CSS customizations that you might wish to apply to your Swarm installation. Each example can be implemented separately. Ideally, you would apply the CSS customizations in a single file to reduce the number of web requests required.

Tip

If your Swarm administrator has configured Swarm to connect to more than one Helix Core Server instance you can customize Swarm for each server instance it is connected to. For instructions on customizing Swarm for individual Helix Core Server instances, see ["CSS customization when Swarm is connected to multiple Helix Core Server instances" on page 794](#).

Adjust the default tab size

```
body {
    tab-size: 4;
}
```

Replace the `4` with the tab size you prefer.

Save these lines in a file, perhaps `tab-size.css`, within the `SWARM_ROOT/public/custom` folder. Swarm automatically includes the CSS file, adjusting the tab size immediately for subsequent page views.

Apply a custom background to the login screen

```
.session-container .modal-overlay.login::after {
    background-image: url('/custom/login_background.jpg');
    background-position: top center;
    background-size: 100%;
}
```

Replace the `/custom/login_background.jpg` URL fragment with the image you want to use. If you do not specify a full URL, the image you specify must exist within the `SWARM_ROOT/public` folder, preferably within the `SWARM_ROOT/public/custom` folder.

Save these lines in a file, perhaps `login-background.css`, within the `SWARM_ROOT/public/custom` folder. Swarm automatically includes the CSS file, immediately replacing the login screen's background for subsequent page views.

Replace Swarm's logo in the main navigation bar

```
#react-swarm-app-container .primary.navigation .swarmLogo a svg{
    display:none;
}
#react-swarm-app-container .primary.navigation .swarmLogo a {
    background-repeat: no-repeat;
    background-image: url('/custom/navbar_logo.jpg');
    background-size: 116px 25px;
    background-position: 15px center;
    background-clip: content-box;
}
```

Replace the `/custom/navbar_logo.jpg` URL fragment with the image you want to use. If you do not specify a full URL, the image you specify must exist within the `SWARM_ROOT/public` folder, preferably within the `SWARM_ROOT/public/custom` folder.

Save these lines in a file, perhaps `navbar-logo.css`, within the `SWARM_ROOT/public/custom` folder. Swarm automatically includes the CSS file, immediately replacing the Swarm logo for subsequent page views.

Note

Swarm's navbar design supports logos up to 24 pixels tall. Even if your logo fits within that height, you may need to also adjust the width, height, margins, or padding to suit your logo.

Adjust the appearance of avatars

The Swarm UI is changing and moving toward using React, currently there is a mix of React and legacy PHP controlling the UI. This means that there are currently two methods of rendering avatars and you need to customize both.

- For avatars rendered by Swarm using PHP, see ["Adjust avatar appearance when rendered by PHP" on the next page](#)
- For avatars rendered by Swarm using React, see ["Adjust avatar appearance when rendered by React" on the next page](#)

For instructions on adding custom avatars, see ["Add custom avatars" on the next page](#)

Adjust avatar appearance when rendered by PHP

```
img.avatar {  
  border: 2px dashed red;  
  border-radius: 10%;  
}
```

You can make a number of adjustments to the way Swarm presents avatars rendered by PHP, such as adding a border and adjusting the border radius, as the example above demonstrates. You should avoid attempting to set specific sizes because Swarm uses different sizes depending on where the avatar is displayed.

Save these lines in a file, perhaps `avatars.css`, within the `SWARM_ROOT/public/custom` folder. Swarm automatically includes the CSS file, immediately changing the appearance of avatars for subsequent page views.

Adjust avatar appearance when rendered by React

```
#react-swarm-app-container .userAvatar img{  
  border: 2px dashed red;  
  border-radius: 10%;  
}
```

You can make a number of adjustments to the way Swarm presents avatars rendered by React, such as adding a border and adjusting the border radius, as the example above demonstrates. You should avoid attempting to set specific sizes because Swarm uses different sizes depending on where the avatar is displayed.

Save these lines in a file, perhaps `avatars-react.css`, within the `SWARM_ROOT/public/custom` folder. Swarm automatically includes the CSS file, immediately changing the appearance of avatars for subsequent page views.

Add custom avatars

Custom avatars for individual users or groups are added to Swarm using custom CSS. Custom avatars are intended for use with system users and groups.

To add a custom avatar for a user:

```
img.avatar[data-user="<userid>"], .userAvatar img[alt="<user  
name>"] {  
  content: url(<path to avatar>);  
}
```

To add a custom avatar for a group:

```
img.avatar[data-user="<groupid>"], .userAvatar img
[alt="<group name>"]{
    content: url(<path to avatar>);
}
```

Save these lines in a file, perhaps **avatars-custom.css**, within the **SWARM_ROOT/public/custom** folder. Swarm automatically includes the CSS file, immediately using your custom avatars for subsequent page views.

Add a custom icon to a menu item

When you create a custom menu item in the [menu_helpers](#) configuration block, you can specify the menu icon to use in the Swarm CSS.

Add some custom CSS to Swarm to replace the default icon with a custom icon for the menu item. For example, the CSS below replaces the default icon with the **double-speech-bubbles.svg** icon for the **custom01** menu item by modifying the **svg.svgIcon.custom01MenuIcon** and **.menuItem-custom01** classes:

```
.swarmMenu .menuItem.menuItem-tests .svgIcon{
    display:none;
}.swarmMenu .menuItem.menuItem-tests a::before{
    content: '';
    background-image: url('/swarm/img/icons/line/double-speech-
bubbles.svg');
    background-size: 16px;
    background-repeat: no-repeat;
    padding-left: 20px;
    padding-right: 4px;
    margin-left: 20px;
}
```

Save these lines in a file, perhaps **menu_custom01.css** in a folder called **custom_menus**, in the **SWARM_ROOT/public/custom** folder. Swarm automatically includes the CSS file, immediately using your custom menus for subsequent page views. If you are using your own custom images, we recommend you store them in the same folder as your custom css.

CSS customization when Swarm is connected to multiple Helix Core Server instances

If Swarm is configured to connect to more than one Helix Core Server instance, Swarm can be customized for each of the Helix Core Server instances it is connected to. This gives users an simple visual prompt as to which Helix Core Server instance they are currently viewing in Swarm. All of the customizations described in the section above are available for multiple Helix Core Server instances.

Note

- The Swarm Global Dashboard cannot be customized.
- For instructions on how to configure Swarm to connect to multiple Helix Core Server instances, see ["Multiple Helix Core Server instances" on page 646](#).

To customize Swarm for a Helix Core Server instance, include the `[data-server-id="<server-name>"]` data attribute in the CSS selector you are customizing. Replace `<server-name>` with the Helix Core Server name the customization is for.

Tip

To apply a customization to all Swarm Helix Core Server instances, don't include the data attribute in the CSS selector you are customizing.

See the examples below to see how the data attribute is used to customize Swarm for the individual Helix Core Server instances. As with the customization of Swarm when it is connected to a single Helix Core Server, you should ideally apply the CSS customizations in a single file to reduce the number of web requests required.

Apply a custom background color to the navigation bar of each of the Helix Core Server instances

```
body[data-server-id="serverA"] #react-swarm-app-container .swarm-
navigation {
  background-color: #d21544;
}
body[data-server-id="serverB"] #react-swarm-app-container .swarm-
navigation {
  background-color: pink;
}
```

Apply custom background images to the login screen of each of the Helix Core Server instances

```
body.route-login[data-server-id="serverA"] .session-container .modal-  
overlay.login::after {  
    background-position: top center;  
    background-size: 100%;  
    background-image: url(/custom/login_background_A.jpg);  
}  
body.route-login[data-server-id="serverB"] .session-container .modal-  
overlay.login::after {  
    background-position: top center;  
    background-size: 100%;  
    background-image: url(/custom/login_background_B.jpg);  
}
```

Swarm API

Important

From Swarm 2022.2, Swarm no longer supports APIs older than v9.

This chapter describes the REST-like API provided by Swarm, which can be used to automate common Swarm interactions or integrate with external systems.

API versions	796
Swarm API basics	798
API endpoints (v9)	802
API endpoints (v10)	1003

API versions

From Swarm 2022.1, we added a new set of v10 APIs to Swarm that replace the v10 APIs. The v10 APIs refine and extend the v10 APIs, and further standardize the endpoint and response pattern.

The v9 and v10 APIs will continue to be available for some time to come, but their functionality will be duplicated and expanded in the v10 APIs over the next few releases. v10 does not currently provide a full set of features, so it is recommended that you use v9 for now if the features that you need are not present in v10.

Important

- From Swarm 2022.2, Swarm no longer supports APIs older than v9. Swarm will continue to support the v9 APIs and you can continue to use them if you prefer.
- Any improvements made in the migration to v11 will not be backported to v9 or v10.
- New API endpoints will be created as v11 and will not be backported to v9 or v10.

Current API versions

API version	Swarm Release	Date	Description
v11	2022.1	March 2022	Extended the v10 APIs and further standardized the endpoint and response pattern for the new rich User Interface.
v10	2019.3	October 2019	Include support for integration with CI tools.

API version	Swarm Release	Date	Description
v9	2018.1	April 2018	Include support for append and replace changelist to a review, 2 Factor Authentication and mark a comment as read.

Unsupported API versions

From Swarm 2022.2, Swarm no longer supports APIs older than v9.

API version	Swarm Release	Date	Description
v8	2017.4	December 2017	Include support for default reviewers on a project or project branch.
v7	2017.3	October 2017	Include support for groups as participants of a review and groups as moderators of a project branch.
v6	2017.1	May 2017	Include support for activity dashboard, archiving of inactive reviews, cleaning completed reviews and for voting reviews up and down.
v5	2016.3	December 2016	Include support for limiting comments to a specific review version.
v4	2016.2	October 2016	Include support for private projects, as well as file-level and line-level inline comments.
v3	2016.1 SP1	September 2016	Include new endpoint for comments.
v2	2016.1	May 2016	Include new endpoints for projects, groups, etc.
v1.2	2015.3	October 2015	Add author filter to the list reviews endpoint.
v1.1	2014.4	January 2015	Addition of required reviewers, and apiVersions .
v1	2014.3	July 2014	Initial release.

Get Swarm version information

The **version** endpoint lets you see the currently-installed Swarm version, the year of release, and the included API versions.

GET /api/version

Get Swarm version and API versions:

`https://myswarm.url/api/version`

Example response:

```
{ "apiVersions":
[1,1.1,1.2,2,3,4,5,6,7,8,9,10,11], "version": "SWARM\2022.1\2139325
(2022\01\28)", "year": "2022" }
```

Note

year refers to the year of the Swarm release, not necessarily the current year.

Swarm API basics

Important

From Swarm 2022.2, Swarm no longer supports APIs older than v9.

This section describes the API endpoints that are commonly used when constructing an API call to Swarm.

Authentication

Swarm's API requires an authenticated connection for all data-modifying endpoints. Authenticated connections are achieved using HTTP Basic Access Authentication.

Note

API endpoints require authentication unless `require_login` is set to **false**.

For example:

```
curl -u "apiuser:ticket" https://myswarm.url/api/v9/projects
```

Swarm accepts a ticket from the Helix Core Server (previous versions of Swarm required this ticket to be host-unlocked, this is no longer true since Swarm 2017.1). It might also be possible to use a password in place of the ticket if your Helix Core Server allows it.

To acquire a ticket, run the following command:

```
p4 -p myp4host:1666 -u apiuser login -p
```

Important

For a Helix Core Server that has been configured for security level 3, passwords are not accepted. For more information on security levels, see on [Server security levels](#) in the *Helix Core Server Administrator Guide*.

Note

If you use a ticket to authenticate against the Swarm API and the ticket expires, you need to acquire a new ticket to continue using the API.

If you make a request that requires authentication and you have not authenticated, the response is:

```
HTTP/1.1 200 OK

{
  "error": "Unauthorized"
}
```

Requests

Important

Some API endpoints have undocumented metadata parameters, these are for internal Swarm use only. The metadata returned is subject to change without notice and is not suitable for use in customer API calls.

Swarm's API includes endpoints that provide, create, and update information within Swarm.

If you make a request against an endpoint with a method that is not supported, the response is:

```
HTTP/1.1 200 OK

{
  "error": "Method Not Allowed"
}
```

GET information

Tip

GET requests require authentication unless `require_login` is set to **false**.

Use HTTP **GET** requests to ask for information from the API.

For example, to get the list of reviews using the v10 API:

```
curl -u "username:ticket" "https://my-swarm-host/api/v10/reviews"
```

Certain API calls support a **fields** parameter that allows you to specify which fields to include in the response, enabling more compact data sets.

Fields can be expressed as a comma-separated list, or using array-notation. For example:

```
curl -u "username:ticket" "https://my-swarm-host/api/v10/projects/jam?fields=id,description,members"
```

The following endpoints support fields:

API (v9)

- `/api/v9/activity`
- `/api/v9/comments`
- `/api/v9/groups`
- `/api/v9/groups/{id}`
- `/api/v9/projects`
- `/api/v9/reviews`
- `/api/v9/reviews/{id}`
- `/api/v9/users`
- `/api/v9/workflows`
- `/api/v9/workflows/{id}`

API (v10)

- `/api/v10/files/{id}`
- `/api/v10/projects`
- `/api/v10/projects/{id}`
- `/api/v10/reviews`
- `/api/v10/reviews/{id}`
- `/api/v10/reviews/{id}/files`
- `/api/v10/reviews/{id}/testruns`
- `/api/v10/search`
- `/api/v10/specs/{spec}/fields`
- `/api/v10/testdefinitions`
- `/api/v10/workflows`
- `/api/v10/workflows/{id}`

POST new information

Use HTTP **POST** requests to create information via the API.

For example, to create a review using form-encoded values:

```
curl -u "apiuser:ticket" -X POST -d"change=12345"
https://myswarm.url/api/v9/reviews
```

The response should be similar to:

```
HTTP/1.1 200 OK

{
  "isValid": true,
  "id": 12206
}
```

To create a review using JSON:

```
curl -u "apiuser:ticket" -H "Content-type: application/json" -X
POST -d'{"change": 12345}' https://myswarm.url/api/v9/reviews
```

Update

Use HTTP **PATCH** requests to update information via the API.

If your HTTP client does not support **PATCH** requests, you can emulate this behavior by submitting an HTTP **POST** with a "**?_method=PATCH**" parameter.

Pagination

Most Swarm endpoints that provide data include the ability to paginate their results.

Each time data is requested, up to **max** results are included in the response, as is a value called **lastSeen**. **lastSeen** identifies the **id** of the last entry included in the results. If there are no further results, **lastSeen** is **null**.

To get the next set of results, include **after** set to the value of **lastSeen** in the API request. Entries up to and including the **id** specified by **after** are excluded from the response, and the next **max** entries are included.

See the [Activity endpoint](#) for example usage that demonstrates pagination.

Responses

Swarm's API responses are JSON formatted.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

API endpoints (v9)

Important

From Swarm 2022.2, Swarm no longer supports APIs older than v9.

This section includes coverage for the endpoints provided by the API (v9).

Activity : Swarm Activity List

Important

From Swarm 2022.2, Swarm no longer supports APIs older than v9.

List Activity Entries

Summary

List Activity Entries

GET /api/v9/activity

Description

Retrieve the Activity List.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on](#)

[page 387](#).

- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required	Default Value
<code>change</code>	Optionally filter activity entries by associated Changelist ID. This only includes records for which there is an activity entry in Swarm.	integer	form	No	
<code>stream</code>	Optional activity stream to query for entries. This can include user-initiated actions (<code>user-alice</code>), activity relating to a user's followed projects/users (<code>personal-alice</code>), review streams (<code>review-1234</code>), and project streams (<code>project-exampleproject</code>).	string	form	No	
<code>type</code>	Type of activity, for example, <code>change</code> , <code>comment</code> , <code>job</code> , or <code>review</code> .	string	form	No	

Parameter	Description	Type	Parameter Type	Required	Default Value
after	An activity ID to seek to. Activity entries up to and including the specified ID are excluded from the results and do not count towards max . Useful for pagination. Commonly set to the lastSeen property from a previous query.	integer	query	No	
max	Maximum number of activity entries to return. This does not guarantee that max entries are returned. It does guarantee that the number of entries returned won't exceed max . Server-side filtering may exclude some activity entries for permissions reasons.	integer	query	No	100
fields	An optional comma-separated list (or array) of fields to show. Omitting this parameter or passing an empty value shows all fields.	string	query	No	

Example response

Successful Response:

To get the latest activity entry on a review:

```
curl -u "username:password"
http://myswarm.url/api/v9/activity?stream=review-290&max=1
```


Swarm responds with an array of activity entities, and a **lastSeen** value that can be used for pagination:

```
HTTP/1.1 200 OK
```

```
{
  "activity": [
    {
      "id": 619,
      "action": "updated files in",
      "behalfOf": null,
      "behalfOfExists": false,
      "behalfOfFullName": "",
      "change": 290,
      "comments": [
        0,
        0
      ],
      "date": "2019-01-23T02:46:59-08:00",
      "depotFile": null,
      "description": "Start of Project Blue Book.",
      "details": [

      ],
      "followers": [

      ],
      "link": [
        "review",
        {
          "review": "290",
          "version": 2
        }
      ],
      "preposition": "for",
      "projectList": {
        "blue-book": [
```

```

        "main"
      ]
    },
    "projects": {
      "blue-book": [
        "main"
      ]
    },
    "streams": {
      "0": "review-290",
      "1": "user-allison.clayborne",
      "2": "personal-allison.clayborne",
      "3": "project-blue-book",
      "7": "group-triage",
      "8": "personal-alex.randolph",
      "9": "personal-jack.boone"
    },
    "target": "review 290 (revision 2)",
    "time": 1548240419,
    "topic": "reviews/290",
    "type": "review",
    "url": "/main/reviews/290/v2/",
    "user": "allison.clayborne",
    "userExists": true,
    "userFullName": "Allison Clayborne"
  }
],
"lastSeen": 618
}

```

Examples of usage

Fetching review activity

To get the latest activity entries on a review:

Tip

In this example, the review has more than 2 activities but the request has been made to limit the response to a **max** of 2 activities. This will return the most recent activities unless **lastseen** is included in the request. This allows for pagination of your results, see ["Activity pagination"](#) below.

```
curl -u "username:password"  
http://myswarm.url/api/v9/activity?stream=review-  
1234&fields=id,action,date,description,type&max=2
```

You can tweak **max** and **fields** to fetch the data that works best for you.

Swarm responds with an array of activity entities, and a **lastSeen** value that can be used for pagination:

```
HTTP/1.1 200 OK  
  
{  
  "activity": [  
    {  
      "id": 619,  
      "action": "updated files in",  
      "date": "2019-01-23T02:46:59-08:00",  
      "description": "Start of Project Blue Book.",  
      "type": "review"  
    },  
    {  
      "id": 618,  
      "action": "commented on",  
      "date": "2019-01-23T02:46:19-08:00",  
      "description": "This is a short comment.",  
      "type": "comment"  
    }  
  ],  
  "lastSeen": 618  
}
```

Activity pagination

To get the second page of activity entries for a review (based on the previous example):

```
curl -u "username:password"  
"https://myswarm.url/api/v9/activity?stream=review-  
1234&fields=id,date,description,type&max=2&after=618"
```

Swarm again responds with a list of activity entities and an **lastSeen** value:

```
HTTP/1.1 200 OK  
  
{  
  "activity": [  
    {  
      "id": 617,  
      "action": "commented on",  
      "date": "2018-12-30T12:12:12-07:00",  
      "description": "This is the first test comment.",  
      "type": "comment"  
    },  
    {  
      "id": 616,  
      "action": "commented on",  
      "date": "2018-12-27T12:13:14-07:00",  
      "description": "Start of Project Blue Book.",  
      "type": "review"  
    }  
  ],  
  "lastSeen": 616  
}
```

Create Activity Entry

Summary

Create Activity Entry

POST /api/v9/activity

Description

Creates an entry in the Activity List.

Important

The authenticated user must match the user provided unless the authenticated user has minimum admin-level privileges in which case any user may be specified.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>type</code>	Type of activity, used for filtering activity streams (values can include <code>change</code> , <code>comment</code> , <code>job</code> , <code>review</code>).	string	form	Yes
<code>user</code>	User who performed the action.	string	form	Yes
<code>action</code>	Action that was performed - past-tense, for example, <code>created</code> or <code>commented on</code> .	string	form	Yes
<code>target</code>	Target that the action was performed on, for example, <code>issue 1234</code> .	string	form	Yes
<code>topic</code>	Optional topic for the activity entry. Topics are essentially comment thread IDs. Examples: <code>reviews/1234</code> or <code>jobs/job001234</code> .	string	form	No

Parameter	Description	Type	Parameter Type	Required
description	Optional description of object or activity to provide context.	string	form	No
change	Optional changelist ID this activity is related to. Used to filter activity related to restricted changes.	integer	form	No
streams []	Optional array of streams to display on. This can include user-initiated actions (user-alice), activity relating to a user's followed projects/users (personal-alice), review streams (review-1234), and project streams (project-exampleproject).	array (of strings)	form	No
link	Optional URL for target .	string	form	No

Examples of usage

Creating an activity entry

To create a plain activity entry:

```
curl -u "username:password"
  -X POST
  -d "type=job"
  -d "user=jira"
  -d "action=punted"
  -d "target=review 123" \
  "https://myswarm.url/api/v9/activity"
```

JSON Response:

HTTP/1.1 200 OK

```
{
  "activity":{
    "id":620,
    "action":"punted",
    "behalfOf":null,
    "change":null,
    "depotFile":null,
    "description":null,
    "details":[

    ],
    "followers":[

    ],
    "link":null,
    "preposition":"for",
    "projects":[

    ],
    "streams":[

    ],
    "target":"review 123",
    "time":1560783425,
    "topic":null,
    "type":"job",
    "user":"jira"
  }
}
```

Linking an activity entry to a review

Linking activity entries to reviews is useful. This involves providing `link`, `streams`, and `topic` fields in the activity data. The `link` field is used to make the `review 123` string in the activity entry clickable. The `stream` field is needed so that the activity entry can be attached to the review in the Swarm interface. The `topic` field is used to link the activity entry to the comment thread for that topic, in the event that a user wants to comment on the activity.

To create a fully linked activity entry:

```
curl -u "username:password"
  -X POST
  -d "type=job"
  -d "user=jira"
  -d "action=punted"
  -d "target=review 123" \
  -d "streams[]=review-123" \
  -d "link=reviews/123" \
  -d "topic=reviews/123" \
  "https://myswarm.url/api/v9/activity"
```

Swarm responds with an activity entity:

```
HTTP/1.1 200 OK

{
  "activity": {
    "id": 1375,
    "action": "punted",
    "behalfOf": null,
    "change": null,
    "depotFile": null,
    "description": null,
    "details": [

    ],
    "followers": [

    ],
    "link": "reviews/123",
    "preposition": "for",
```



```
"projects": [

],
"streams": [
  "review-123"
],
"target": "review 123",
"time": 1461607739,
"topic": "reviews/123",
"type": "job",
"user": "jira"
}
}
```

Cache: Swarm Cache

Important

From Swarm 2022.2, Swarm no longer supports APIs older than v9.

Swarm config cache file delete

Tip

Alternatively, use the **Cache Info** tab of the **System Information** page. Navigate to the **User id** dropdown menu, select **System Information**, and click the **Cache Info** tab. For more information, see ["Cache Info" on page 731](#).

Summary

Introduced for Swarm 2019.2 and later, deletes the Swarm configuration cache files and restarts the workers.

DELETE /api/v9/cache/config/

Description

Used to delete the Swarm configuration cache files after the Swarm configuration or custom modules have changed. When workers complete their current tasks they are automatically restarted so that they use the new Swarm configuration and custom modules for their next tasks.

Important

Only *admin* users can delete the Swarm configuration cache.

Example response

When cache files are not deleted and workers are not shutdown:

For example, this can happen if you make the call and then make the same call again before the cache files have rebuilt.

```
HTTP/1.1 200 OK

{
  "isValid":true,
  "messages":[
    "File <SWARM_ROOT>/data/cache/module-config-cache.php was not
deleted",
    "File <SWARM_ROOT>/data/cache/module-classmap-cache.php was
not deleted"
  ]
}
```

Example of usage

Tip

It might also be possible to use a password in place of the ticket if your Helix Core Server allows it.

```
curl -u "username:ticket" -X DELETE
"https://myswarm.url/api/v9/cache/config/"
```

Assuming that the authenticated user has permission, Swarm responds with the cache deleted and workers shutdown messages:

```
HTTP/1.1 200 OK

{
  "isValid":true,
  "messages":[
    "File <SWARM_ROOT>/data/cache/module-config-cache.php
deleted",
```

```
"File <SWARM_ROOT>/data/cache/module-classmap-cache.php
deleted",
    "Workers are shutdown. They will start automatically by
recurring task"
]
}
```

Verify Redis cache

Tip

Alternatively, use the **Cache Info** tab of the **System Information** page. Navigate to the **User id** dropdown menu, select **System Information**, and click the **Cache Info** tab. For more information, see ["Cache Info" on page 731](#).

Summary

Introduced for Swarm 2019.2 and later, verifies the Swarm Redis cache.

POST /serverid/api/v9/cache/redis/verify?context=user,group,project,workflow

Description

The data in the Redis cache for groups, users, projects, and workflows should be the same as the data held by the Helix Core Server. However, in some circumstances, such as when changes are made in the Helix Core Server when Swarm is down or if errors occur during updates, the Redis cache can get out of sync with the Helix Core Server.

Used to verify the content of the Redis cache. Use the context parameter to verify specific areas of the Redis cache, if no context is specified the entire Redis cache is verified. If a verification finds an out of sync cache file, Swarm automatically updates the data in that cache.

Important

Only *admin* users can verify the Redis cache.

Parameters

Parameter	Description	Type	Parameter Type	Required	Default Value
<code>serverid</code>	If Swarm is connected to multiple Helix Core Servers , <code>serverid</code> specifies which Helix Core Server the cache should be verified for.	string	path	See description	

Parameter	Description	Type	Parameter Type	Required	Default Value
context	<p>Add context to the call to verify a specific area of the Redis cache:</p> <ul style="list-style-type: none"> ▪ No context specified, verifies all of the Redis caches. ▪ user verifies the Redis user cache. ▪ group verifies the Redis group cache. ▪ project verifies the Redis project cache. ▪ workflow verifies the Redis workflow cache. <p>To verify multiple areas of cache, separate the entries with a comma ,</p>	string	path	No	

Example usage

Tip

It might also be possible to use a password in place of the ticket if your Helix Core Server allows it.

To start verification of the user and group caches:

```
curl -u "username:ticket" -X POST
"https://myswarm.url/api/v9/cache/redis/verify?context=user,group"
```

JSON response

```
HTTP/1.1 200 OK

{
  isValid: true,
  data: {
    user: {
      state: "Queued",
      progress: false
    },
    group: {
      state: "Queued",
      progress: false
    }
  },
  messages: {
    "Verifying Redis cache for [user,group]"
  }
}
```

Check progress of Redis cache verification

Tip

Alternatively, use the **Cache Info** tab of the **System Information** page. Navigate to the **User id** dropdown menu, select **System Information**, and click the **Cache Info** tab. For more information, see ["Cache Info" on page 731](#).

Summary

Introduced for Swarm 2019.2 and later, check progress of the Swarm Redis cache verification.

GET /serverid/api/v9/cache/redis/verify?context=user,group,project,workflow

Description

Used to check the progress of Redis cache verification process. Use the context parameter to check verification of specific areas of the Redis cache, if no context is specified progress is checked for all areas of the entire Redis cache.

Parameters

Parameter	Description	Type	Parameter Type	Required	Default Value
<code>serverid</code>	If Swarm is connected to multiple Helix Core Servers , <code>serverid</code> specifies which Helix Core Server cache verification is being checked for.	string	path	See description	

Parameter	Description	Type	Parameter Type	Required	Default Value
context	<p>Add context to the check verification progress for a specific area of the Redis cache:</p> <ul style="list-style-type: none"> ■ No context specified, checks verification progress for all of the Redis caches. ■ user checks verification progress for the Redis user cache. ■ group checks verification progress for the Redis group cache. ■ project checks verification progress for the Redis project cache. ■ workflow checks verification progress for the Redis workflow cache. 	string	path	No	

Parameter	Description	Type	Parameter Type	Required	Default Value
	To check the verification progress for multiple areas of cache, separate the entries with a comma ,				

Example usage

Tip

It might also be possible to use a password in place of the ticket if your Helix Core Server allows it.

To check verification progress for the user and group caches:

```
curl -u "username:ticket"
"http://myswarm.url/api/v9/cache/redis/verify?context=user,group"
```

JSON response

```
HTTP/1.1 200 OK

{
  isValid: true,
  data: {
    user: {
      state: "Running",
      progress: "Step 2 of 5; Calculating checksums for Redis
keys"
    },
    group: {
      state: "Running",
      progress: "Step 1 of 5; Getting Redis cache entries"
    }
  }
}
```

Tip

- Valid states for verify are; **Not Queued**, **Queued**, **Running**, and **Failed**.
- Valid states for progress are, **Step x of y**; (where **x** is the current step and **y** is the total number of steps) followed by:
 - **Getting Redis cache entries**
 - **Calculating checksums for Redis keys**
 - **Calculating checksums for Perforce models**
 - **Clearing any extraneous Redis cache entries**
 - **Indexing any missing Perforce models into the Redis cache**
 - **Verification complete**

Redis cache delete

Summary

Introduced for Swarm 2019.2 and later, deletes the Swarm Redis cache.

DELETE /serverid/api/v9/cache/redis?context=user,group,project,workflow

Description

Used to delete the Redis cache. Use the context parameter to delete specific areas of the Redis cache, if no context is specified the entire Redis cache is deleted.

Important

Only *admin* users can delete the Redis cache.

Parameters

Parameter	Description	Type	Parameter Type	Required	Default Value
<code>serverid</code>	If Swarm is connected to multiple Helix Core Servers , <code>serverid</code> specifies which Helix Core Server the cache should be deleted for.	string	path	See description	

Parameter	Description	Type	Parameter Type	Required	Default Value
context	<p>Add context to the call to delete a specific area of the Redis cache:</p> <ul style="list-style-type: none"> ▪ No context specified, deletes all of the Redis caches. ▪ user deletes the Redis user cache. ▪ group deletes the Redis group cache. ▪ project deletes the Redis project cache. ▪ workflow deletes the Redis workflow cache. <p>To delete multiple areas of cache, separate the entries with a comma ,</p>	string	path	No	

Example response

When an invalid Redis cache context is specified:

If an invalid context is specified in the call:

```
HTTP/1.1 200 OK
```

```
{
  "isValid":true,
  "messages":[
    "Invalid context [usr], must contain only [user, group,
project, workflow]",
  ]
}
```

Example of usage

Tip

It might also be possible to use a password in place of the ticket if your Helix Core Server allows it.

```
curl -u "username:ticket" -X DELETE
"https://myswarm.url/serverA/api/v9/cache/redis?context=user,group"
```

Assuming that the authenticated user has permission, Swarm responds with the cache deleted message:

```
HTTP/1.1 200 OK

{
  "isValid":true,
  "messages":[
    "Deleted key(s) [Swarm^serverA:user-populated-status,
Swarm^serverA:group-populated-status]",
  ]
}
```

Changes : API controller providing a service for changes

Important

From Swarm 2022.2, Swarm no longer supports APIs older than v9.

Get projects and branches affected by a given change id

Summary

Get projects, and branches, affected by a given change id.

GET /api/v9/changes/{change}/affectsprojects

Description

All authenticated users are able to use this API.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Example response

Successful Response:

HTTP/1.1 200 OK

```
{
  "change": {
    "id": "1050",
    "projects": {
      "jam": [
        "live",
        "main"
      ]
    }
  }
}
```

Get default reviewers for a given change id

Summary

Get default reviewers for a given change id.

GET /api/v9/changes/{change}/defaultreviewers

Description

All authenticated users are able to use this API.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Example response

Successful Response:

```
HTTP/1.1 200 OK

{
  "change": {
    "id": "1050",
    "defaultReviewers": {
      "groups": {
        "group1": {"required": "1"},
        "group2": {}
      },
      "users": {
        "user1": {},
        "user2": {"required": "true"}
      }
    }
  }
}
```

Perform checks on a change if enabled

Summary

Performs checks on the change if enabled

GET /api/v9/changes/{id}/check

Description

Performs checks on the change if workflow configuration requires it.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Change to check	string	form	Yes
<code>type</code>	The type of check. Must have a value of enforced, strict or shelve	string	form	Yes

Example response

Successful Response:

HTTP/1.1 200 OK

```
{
  "status":    "OK",
  "isValid":   "true",
  "messages": []
}
```


Examples of usage

Carry out enforced checks on a change if configured

To check change 42:

```
curl -u "username:password" "https://my-swarm-host/api/v9/changes/42/check?type=enforced"
```

JSON Response:

```
HTTP/1.1 200 OK
```

```
{
  "status":    "OK",
  "isValid":   "true",
  "messages":  []
}
```

Carry out enforced checks on a change if configured. Example of a failed 'requires review'

To check change 42:

```
curl -u "username:password" "https://my-swarm-host/api/v9/changes/42/check?type=enforced"
```

JSON Response:

```
HTTP/1.1 200 OK
```

```
{
  "status":    "NO_REVIEW",
  "isValid":   "false",
  "messages":  ["Change [42] must be associated with a review"]
}
```

Carry out strict checks on a change if configured

To check change 42:

```
curl -u "username:password" "https://my-swarm-host/api/v9/changes/42/check?type=strict"
```

JSON Response:

```
HTTP/1.1 200 OK
```

```
{
  "status":    "OK",
  "isValid":   "true",
  "messages":  []
}
```

```
HTTP/1.1 200 OK

{
  "status":    "OK",
  "isValid":   "true",
  "messages": []
}
```

Comments : Swarm Comments

Important

From Swarm 2022.2, Swarm no longer supports APIs older than v9.

Get List of Comments

GET /api/v9/comments/

Summary

Get List of Comments

Description

List comments.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required	Default Value
<code>after</code>	A comment ID to seek to. Comments up to and including the specified ID are excluded from the results and do not count towards <code>max</code> . Useful for pagination. Commonly set to the <code>lastSeen</code> property from a previous query.	integer	query	No	
<code>max</code>	Maximum number of comments to return. This does not guarantee that <code>max</code> comments are returned. It does guarantee that the number of comments returned won't exceed <code>max</code> .	integer	query	No	100
<code>topic</code>	Only comments for given topic are returned. Examples: <code>reviews/1234</code> , <code>changes/1234</code> or <code>jobs/job001234</code> .	string	query	No	

Parameter	Description	Type	Parameter Type	Required	Default Value
context [version]	If a reviews/1234 topic is provided, limit returned comments to a specific version of the provided review.	integer	query	No	
ignoreArchived	Prevents archived comments from being returned. (v5+)	boolean	query	No	
tasksOnly	Returns only comments that have been flagged as tasks. (v5+)	boolean	query	No	
taskStates	Limit the returned comments to ones that match the provided task state (one or more of open , closed , verified , or comment). (v5+)	array (of strings)	query	No	
fields	An optional comma-separated list (or array) of fields to show for each comment. Omitting this parameter or passing an empty value shows all fields.	string	query	No	

Examples of successful responses

Successful Response:

```
HTTP/1.1 200 OK

{
  "topic": "",
  "comments": {
    "51": {
      "id": 51,
      "attachments": [],
      "body": "Short loin ground round sin reprehensible, venison west
participle triple.",
      "context": [],
      "edited": null,
      "flags": [],
      "likes": [],
      "taskState": "comment",
      "time": 1461164347,
      "topic": "reviews/885",
      "updated": 1461164347,
      "user": "bruno"
    }
  },
  "lastSeen": 51
}
```

Note

lastSeen can often be used as an offset for pagination, by using the value in the **after** parameter of subsequent requests.

When no results are found, the `comments` array is empty:

```
HTTP/1.1 200 OK

{
```

```
"topic": "jobs/job000011",
"comments": [],
"lastSeen": null
}
```

Examples of usage

Listing comments

To list comments:

```
curl -u "username:password" "https://my-swarm-
host/api/v9/comments?topic=reviews/911&max=2&fields=id,body,time,user"
```

Swarm responds with a list of the first two comments for review 911 and a **lastSeen** value for pagination:

```
HTTP/1.1 200 OK

{
  "topic": "reviews/911",
  "comments": {
    "35": {
      "id": 35,
      "body": "Excitation thunder cats intelligent man braid organic
bitters.",
      "time": 1461164347,
      "user": "bruno"
    },
    "39": {
      "id": 39,
      "body": "Chamber tote bag butcher, shirk truffle mode shabby
chic single-origin coffee.",
      "time": 1461164347,
      "user": "swarm_user"
    }
  },
  "lastSeen": 39
}
```

Paginating a comment listing

To obtain the next page of a comments list (based on the previous example):

```
curl -u "username:password" "https://my-swarm-  
host/api/  
v9/comments?topic=reviews/911&max=2&fields=id,body,time,user&after=39"
```

Swarm responds with the second page of results, if any comments are present after the last seen comment:

```
HTTP/1.1 200 OK  
  
{  
  "topic": "reviews/911",  
  "comments": {  
    "260": {  
      "id": 260,  
      "body": "Reprehensible do lore flank ham hock.",  
      "time": 1461164349,  
      "user": "bruno"  
    },  
    "324": {  
      "id": 324,  
      "body": "Sinter lo-fi temporary, nihilist tote bag mustache swag  
consequence interest flexible.",  
      "time": 1461164349,  
      "user": "bruno"  
    }  
  },  
  "lastSeen": 324  
}
```

Send notification for comments

Summary

Sends notification for comments

POST /api/v9/comments/notify

Description

Sends notification for comments.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>topic</code>	This is going to send a single notification for any comments that were not notified immediately for the user authenticated for a given topic they are posting for. Example: <code>reviews/1234</code> .	string	form	Yes
<code>preview</code>	Set to <code>true</code> to check if there are comments waiting for a notification to be sent, the notification is not sent.	string	form	No

Examples of successful responses

Comment notification sent.

```
HTTP/1.1 200 OK
```

```
{
  "isValid": true,
  "message": "Sending a notification for 3 comments",
  "code": 200
}
```


No comments are waiting for a notification.

```
HTTP/1.1 200 OK

{
  "isValid": true,
  "message": "No comment notifications to send",
  "code": 200
}
```

Examples of usage

Sending notification for comments

To send a notification for all delayed comments:

```
curl -u "username:password" \
  -X POST \
  -d "topic=reviews/911" \
  "https://my-swarm-host/api/v9/comments/notify"
```

Swarm responds with a message detailing a notification sent with the number of comments it applied to.

```
HTTP/1.1 200 OK

{
  "isValid": true,
  "message": "Sending a notification for 3 comments",
  "code": 200
}
```

Checking to see if any comments are waiting for a notification to be sent

To check if there are any comments awaiting a notification without sending the notification, include **"preview=true"** in your request.

```
curl -u "username:password" \
  -X POST \
  -d "topic=reviews/911" \
```

```
-d "preview=true" \  
"https://my-swarm-host/api/v9/comments/notify"
```

Swarm responds with a message without sending a notification.

```
HTTP/1.1 200 OK
```

```
{  
  "isValid": true,  
  "message": "Sending a notification for 3 comments",  
  "code": 200  
  "preview": true  
}
```

Add a Comment

Summary

Add a Comment

POST /api/v9/comments/

Description

Add a comment to a topic (such as a review or a job).

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required	Default Value
<code>topic</code>	Topic to comment on. Examples: <code>reviews/1234</code> , <code>changes/1234</code> or <code>jobs/job001234</code> .	string	form	Yes	
<code>body</code>	Content of the comment.	string	form	Yes	
<code>silenceNotification</code>	If set to 'true' no notifications will ever be sent for this created comment	string	form	No	
<code>delayNotification</code>	If set to 'true' notifications will be delayed	string	form	No	
<code>taskState</code>	Optional task state of the comment. Valid values when adding a comment are <code>comment</code> and <code>open</code> . This creates a plain comment or opens a task, respectively.	string	form	No	comment
<code>flags[]</code>	Optional flags on the comment. Typically set to <code>closed</code> to archive a comment.	array (of strings)	form	No	
<code>context[file]</code>	File to comment on. Valid only for <code>changes</code> and <code>reviews</code> topics. Example: <code>//depot/main/READEME.txt</code> .	string	form	No	
<code>context[leftLine]</code>	Left-side diff line to attach the inline comment to. Valid only for <code>changes</code> and <code>reviews</code> topics. If this is specified, <code>context[file]</code> must also be specified.	integer	form	No	

Parameter	Description	Type	Parameter Type	Required	Default Value
<code>context</code> <code>[rightLine]</code>	Right-side diff line to attach the inline comment to. Valid only for <code>changes</code> and <code>reviews</code> topics. If this is specified, <code>context</code> <code>[file]</code> must also be specified.	integer	form	No	
<code>context</code> <code>[content]</code>	Optionally provide content of the specified line and its four preceding lines. This is used to specify a short excerpt of context in case the lines being commented on change during the review. When not provided, Swarm makes an effort to build the content on its own - as this involves file operations, it could become slow.	array (of strings)	form	No	
<code>context</code> <code>[version]</code>	With a <code>reviews</code> topic, this field specifies which version to attach the comment to.	integer	form	No	

Example of response

Successful Response contains Comment entity:

```
HTTP/1.1 200 OK
```

```
{
  "comment": {
    "id": 42,
    "attachments": [],
```

```
{
  "body": "Best. Comment. EVER!",
  "context": [],
  "edited": null,
  "flags": [],
  "likes": [],
  "taskState": "comment",
  "time": 123456789,
  "topic": "reviews/2",
  "updated": 123456790,
  "user": "bruno"
}
```

Examples of usage

Create a comment on a review

To create a comment on a review:

```
curl -u "username:password" \
  -X POST \
  -d "topic=reviews/2" \
  -d "body=This is my comment. It is an excellent comment. It
contains a beginning, a middle, and an end." \
  "https://my-swarm-host/api/v9/comments"
```

JSON Response:

```
HTTP/1.1 200 OK

{
  "comment": {
    "id": 42,
    "attachments": [],
    "body": "This is my comment. It is an excellent comment. It
contains a beginning, a middle, and an end.",
    "context": [],
    "edited": null,
    "flags": [],
```

```

    "likes": [],
    "taskState": "comment",
    "time": 123456789,
    "topic": "reviews/2",
    "updated": 123456790,
    "user": "username"
  }
}

```

Open a task on a review

To create a comment on a review, and flag it as an open task:

```

curl -u "username:password" \
  -X POST \
  -d "topic=reviews/2" \
  -d "taskState=open" \
  -d "body=If you could go ahead and attach a cover page to your
TPS report, that would be great." \
  "https://my-swarm-host/api/v9/comments"

```

JSON Response:

```

HTTP/1.1 200 OK

{
  "comment": {
    "id": 43,
    "attachments": [],
    "body": "If you could go ahead and attach a cover page to your TPS
report, that would be great.",
    "context": [],
    "edited": null,
    "flags": [],
    "likes": [],
    "taskState": "open",
    "time": 123456789,
    "topic": "reviews/2",
    "updated": 123456790,

```

```
"user": "username"
}
```

Edit a Comment

Summary

Edit a Comment

PATCH /api/v9/comments/{id}

Description

Edit a comment.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	ID of the comment to be edited	integer	path	Yes
<code>topic</code>	Topic to comment on. Examples: <code>reviews/1234</code> , <code>changes/1234</code> or <code>jobs/job001234</code> .	string	form	No
<code>body</code>	Content of the comment.	string	form	Yes

Parameter	Description	Type	Parameter Type	Required
taskState	Optional task state of the comment. Note that certain transitions (such as moving from open to verified) are not possible without an intermediate step (addressed , in this case). Examples: comment (not a task), open , addressed , verified .	string	form	No
flags[]	Optional flags on the comment. Typically set to closed to archive a comment.	array (of strings)	form	No
silenceNotification	If set to 'true' no notifications will ever be sent for this edited comment	string	form	No
delayNotification	If set to 'true' notifications will be delayed	string	form	No

Example response

Successful Response contains Comment entity:

HTTP/1.1 200 OK

```
{
  "comment": {
    "id": 1,
    "attachments": [],
    "body": "Best. Comment. EVER!",
    "context": [],
```



```
"edited": 123466790,
"flags": [],
"likes": [],
"taskState": "comment",
"time": 123456789,
"topic": "reviews/42",
"updated": 123456790,
"user": "bruno"
}
}
```

Examples of usage

Edit and archive a comment on a review

To edit and archive a comment on a review:

```
curl -u "username:password" \
  -X PATCH \
  -d "flags[]=closed" \
  -d "body=This comment wasn't as excellent as I may have lead you
to believe. A thousand apologies." \
  "https://my-swarm-host/api/v9/comments/42"
```

JSON Response:

```
HTTP/1.1 200 OK

{
  "comment": {
    "id": 42,
    "attachments": [],
    "body": "This comment wasn't as excellent as I may have lead you
to believe. A thousand apologies.",
    "context": [],
    "edited": 123466790,
    "flags": ["closed"],
    "likes": [],
    "taskState": "comment",
```

```

    "time": 123456789,
    "topic": "reviews/2",
    "updated": 123456790,
    "user": "username"
  }
}

```

Flag a task as addressed on a review

To flag an open task as addressed on a review:

```

curl -u "username:password" \
  -X PATCH \
  -d "taskState=addressed" \
  "https://my-swarm-host/api/v9/comments/43"

```

JSON Response:

```

HTTP/1.1 200 OK

{
  "comment": {
    "id": 43,
    "attachments": [],
    "body": "If you could go ahead and attach a cover page to your TPS
report, that would be great.",
    "context": [],
    "edited": 123466790,
    "flags": ["closed"],
    "likes": [],
    "taskState": "addressed",
    "time": 123456789,
    "topic": "reviews/2",
    "updated": 123456790,
    "user": "username"
  }
}

```

Groups : Swarm Groups

Important

From Swarm 2022.2, Swarm no longer supports APIs older than v9.

Get List of Groups

Summary

Get List of Groups

GET /api/v9/groups/

Description

Returns the complete list of groups in Swarm.

Parameters

Parameter	Description	Type	Parameter Type	Required	Default Value
after	A group ID to seek to. Groups prior to and including the specified ID are excluded from the results and do not count towards max . Useful for pagination. Commonly set to the lastSeen property from a previous query.	string	query	No	

Parameter	Description	Type	Parameter Type	Required	Default Value
max	Maximum number of groups to return. This does not guarantee that max groups are returned. It does guarantee that the number of groups returned won't exceed max .	integer	query	No	100
fields	An optional comma-separated list (or array) of fields to show for each group. Omitting this parameter or passing an empty value shows all fields.	string	query	No	

Parameter	Description	Type	Parameter Type	Required	Default Value
keywords	Keywords to limit groups on. Only groups where the group ID, group name (if set), or description contain the specified keywords are returned.	string	query	No	
ignoreExcludeList	Determines if the list of groups has the group_exclude_list filter applied or not. Add the parameter to ignore the group_exclude_list filter.	boolean	query	No	

Example response

Successful Response:

```
HTTP/1.1 200 OK
```

```
{
  "groups": [
    {
      "Group": "test-group",
      "MaxLockTime": null,
      "MaxResults": null,
      "MaxScanRows": null,
```

```

    "Owners": [],
    "PasswordTimeout": null,
    "Subgroups": [],
    "Timeout": 43200,
    "Users": ["bruno"],
    "config": {
      "description": "Our testing group",
      "emailAddress": "test-group3@host.domain",
      "emailFlags": [],
      "name": "Test Group",
      "useMailingList": true
    }
  }
]
}

```

Examples of usage

Listing groups

To list groups:

```

curl -u "username:password"
"https://myswarm.url/api/v9/groups?keywords=test-
group&fields=Group,Owners,Users,config&max=2"

```

Swarm responds with a list of groups:

```

HTTP/1.1 200 OK

{
  "groups": [
    {
      "Group": "test-group",
      "Owners": [],
      "Users": ["bruno"],
      "config": {
        "description": "Our testing group",
        "emailAddress": "test-group@host.domain",

```

```
    "emailFlags": {
      "reviews": "1",
      "commits": "0"
    },
    "name": "Test Group",
    "useMailingList": true
  }
},
{
  "Group": "test-group2",
  "Owners": [],
  "Users": ["bruno"],
  "config": {
    "description": "Our second testing group",
    "emailAddress": "test-group2@host.domain",
    "emailFlags": [],
    "name": "Test Group 2",
    "useMailingList": true
  }
}
],
"lastSeen": "test-group2"
}
```

Paginating the groups list

Based on the previous example, we can pass a lastSeen value of **test-group2** to see if there are any subsequent groups in Swarm.

```
curl -u "username:password"
"https://myswarm.url/api/v9/groups?keywords=test-
group&fields=Group,config&max=2&lastSeen=test-group2"
```

Swarm responds with a list of groups (minus the Owners and Users fields, as we haven't requested them):

```
HTTP/1.1 200 OK
```

```
{
```

```
"groups": [  
  {  
    "Group": "test-group3",  
    "config": {  
      "description": "Our 3rd testing group",  
      "emailAddress": "test-group3@host.domain",  
      "emailFlags": [],  
      "name": "Test Group 3",  
      "useMailingList": true  
    }  
  }  
],  
"lastSeen": "test-group3"  
}
```

Get Group Information

Summary

Get Group Information

GET /api/v9/groups/{id}

Description

Retrieve information about a group.

Parameters

Parameter	Description	Type	Parameter Type	Required
id	Group ID	string	path	Yes

Parameter	Description	Type	Parameter Type	Required
<code>fields</code>	An optional comma-separated list (or array) of fields to show for each group. Omitting this parameter or passing an empty value shows all fields.	string	query	No
<code>ignoreExcludeList</code>	Determines if the <code>group_exclude_list</code> filter is applied or not when getting the group information. If the filter is applied and the group specified is present in the <code>group_exclude_list</code> the group information is not included in the response. Add the parameter to ignore the <code>group_exclude_list</code> filter.	boolean	query	No

Example response

Successful Response:

```
HTTP/1.1 200 OK

{
  "group": {
    "Group": "test-group",
    "MaxLockTime": null,
    "MaxResults": null,
    "MaxScanRows": null,
    "Owners": [],
    "PasswordTimeout": null,
    "Subgroups": [],
    "Timeout": 43200,
    "Users": ["bruno"],
    "config": {
      "description": "Our testing group",
      "emailFlags": [],
      "name": "Test Group"
    }
  }
}
```

Examples of usage

Fetching a group

To fetch an individual group:

```
curl -u "username:password" "https://myswarm.url/api/v9/groups/my-group"
```

Swarm responds with the group entity:

```
HTTP/1.1 200 OK

{
  "group": {
```

```
"Group": "test-group",
"LdapConfig": null,
"LdapSearchQuery": null,
"LdapUserAttribute": null,
"MaxLockTime": null,
"MaxResults": null,
"MaxScanRows": null,
"Owners": [],
"Users": ["bruno"],
"config": {
  "description": "Our testing group",
  "emailAddress": "test-group@host.domain",
  "emailFlags": {
    "reviews": "1",
    "commits": "0"
  },
  "name": "Test Group",
  "useMailingList": true
}
}
```

Limiting returned fields

To limit the returned fields when fetching an individual group:

```
curl -u "username:password" "https://myswarm.url/api/v9/groups/my-
group?fields=Group,Owners,Users,config"
```

Swarm responds with the group entity:

```
HTTP/1.1 200 OK

{
  "group": {
    "Group": "test-group",
    "Owners": [],
    "Users": ["bruno"],
    "config": {
```

```

    "description": "Our testing group",
    "emailFlags": [],
    "name": "Test Group"
  }
}
}

```

Create a new Group

Summary

Create a new Group

POST /api/v9/groups/

Description

Creates a new group in Swarm.

Parameters

Parameter	Description	Type	Parameter Type	Required
Group	Group identifier string.	string	form	Yes
Users	<p>An optional array of group users.</p> <p>*At least one of Users, Owners, or Subgroups is required.</p>	array	form	<p>No</p> <p>*See Description</p>
Owners	<p>An optional array of group owners.</p> <p>*At least one of Users, Owners, or Subgroups is required.</p>	array	form	<p>No</p> <p>*See Description</p>

Parameter	Description	Type	Parameter Type	Required
Subgroups	An optional array of subgroups. *At least one of Users, Owners, or Subgroups is required.	array	form	No *See Description
config[name]	An optional full name for the group.	string	form	No
config[description]	An optional group description.	string	form	No
config[emailAddress]	The email address for this group.	string	form	No
config[emailFlags][reviews]	Email members when a new review is requested.	boolean	form	No
config[emailFlags][commits]	Email members when a change is committed.	boolean	form	No
config[useMailingList]	Whether to use the configured email address or expand individual members addresses.	boolean	form	No

Example response

Successful Response:

```
HTTP/1.1 200 OK
```

```
{
```

```

"group": {
  "Group": "test-group",
  "MaxLockTime": null,
  "MaxResults": null,
  "MaxScanRows": null,
  "Owners": [],
  "PasswordTimeout": null,
  "Subgroups": [],
  "Timeout": null,
  "Users": ["alice"],
  "config": {
    "description": "Test test test",
    "emailAddress": "test-group@host.domain",
    "emailFlags": [],
    "name": "TestGroup",
    "useMailingList": true
  }
}
}

```

Example of usage

Create a group

Important

- Only users with *super* privileges in the Helix Core Server (**p4d**), or users with *admin* privileges in **p4d** versions 2012.1 or newer, can create groups.
- This API version is only capable of setting specific fields: **Group**, **Users**, **Owners**, **Subgroups**, **config**. Any other fields specified in the creation request are ignored.

To create a new group:

```

curl -u "username:password" \
  -X POST \
  -d "Group=my-group" \
  -d "Owners[]=alice" \
  -d "Owners[]=bob" \
  -d "Users[]=bruno" \

```

```
-d "Users[]=user2" \  
-d "config[description]=This group is special to me." \  
-d "config[name]=My Group" \  
-d "config[emailFlags][reviews]=1" \  
-d "config[emailFlags][commits]=0" \  
-d "config[emailAddress]=my-group@host.domain" \  
-d "config[useMailingList]=false" \  
"https://myswarm.url/api/v9/groups"
```

Assuming that the authenticated user has permission, Swarm responds with the new group entity:

HTTP/1.1 200 OK

```
{  
  "group": {  
    "Group": "my-group",  
    "MaxLockTime": null,  
    "MaxResults": null,  
    "MaxScanRows": null,  
    "Owners": ["username"],  
    "PasswordTimeout": null,  
    "Subgroups": [],  
    "Timeout": null,  
    "Users": [],  
    "config": {  
      "description": "This group is special to me.",  
      "emailFlags": {  
        "reviews": "1",  
        "commits": "0"  
      },  
      "name": "My Group",  
      "useMailingList": true  
    }  
  }  
}
```

Edit a Group

Summary

Edit a Group

PATCH /api/v9/groups/{id}

Description

Change the settings of a group in Swarm. Only super users and group owners can perform this action.

Parameters

Parameter	Description	Type	Parameter Type	Required
id	Group ID	string	path	Yes
Users	An optional array of group users.	array	form	No
Owners	An optional array of group owners.	array	form	No
Subgroups	An optional array of group subgroups.	array	form	No
config[name]	An optional full name for the group.	string	form	No
config[description]	An optional group description.	string	form	No
config[emailAddress]	The email address for this group.	string	form	No
config[emailFlags][commits]	Email members when a change is committed.	boolean	form	No

Parameter	Description	Type	Parameter Type	Required
<code>config[emailFlags][reviews]</code>	Email members when a new review is requested.	boolean	form	No
<code>config[useMailingList]</code>	Whether to use the configured email address or expand individual members addresses.	boolean	form	No

Example response

Successful Response:

HTTP/1.1 200 OK

```
{
  "group": {
    "Group": "test-group",
    "Users": [],
    "Owners": [],
    "Subgroups": [],
    "config": {
      "description": "New Group Description",
      "name": "TestGroup",
      "useMailingList": true
    }
  }
}
```

Example of usage

Editing a group

Important

- Only users with *super* privileges in the Helix Core Server, or group owners, can edit groups.
- This API version is only capable of modifying specific fields: **Users**, **Owners**, **Subgroups**, **config**. Any other fields specified in the edit request are ignored.

Here is how to update the **name**, **description**, and **emailFlags** configuration of the group **my-group**:

```
curl -u "username:password" \  
  -X PATCH \  
  -d "config[description]=This group is special to me." \  
  -d "config[name]=My Group" \  
  -d "config[emailFlags][commits]=1" \  
  "https://myswarm.url/api/v9/groups/my-group"
```

Assuming that the authenticated user has permission, Swarm responds with the modified group entity:

```
HTTP/1.1 200 OK  
  
{  
  "group": {  
    "Group": "my-group",  
    "Users": [],  
    "Owners": [],  
    "Subgroups": [],  
    "config": {  
      "description": "This group is special to me.",  
      "emailAddress": "test-group@host.domain",  
      "emailFlags": {  
        "reviews": "1",  
        "commits": "1"  
      },  
      "name": "My Group",  
      "useMailingList": true  
    }  
  }  
}
```

```
}  
}
```

Delete a Group

Summary

Delete a Group

DELETE /api/v9/groups/{id}

Description

Delete a group. Only super users and group owners can perform this action.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Group ID.	string	path	Yes

Example response

Successful Response:

```
HTTP/1.1 200 OK  
  
{  
  "id": "test-group"  
}
```

Example of usage

Deleting a group

Important

Only *super* users and group owners can delete groups.

```
curl -u "username:password" -X DELETE  
"https://myswarm.url/api/v9/groups/my-group"
```

Assuming that the authenticated user has permission, Swarm responds with the `id` of the deleted group:

```
HTTP/1.1 200 OK

{
  "id": "my-group"
}
```

Login : Swarm Login API

Important

From Swarm 2022.2, Swarm no longer supports APIs older than v9.

Checking the 2FA authentication

Summary

Checking the 2FA authentication

GET /api/v9/checkauth/

Description

Checking the 2FA authentication

Example response

Successful Response:

```
HTTP/1.1 200 OK

{
  "results": {
    "trigger": "GAuth says yes!",
    "successMsg": "Second factor authentication approved."
  },
  "code": 200
}
```

Usage example

Checking the 2FA authentication

To Check User prompt input for 2FA

```
curl -u "username:password"  
"https://myswarm.url/api/v9/login/checkauth"
```

Assuming that the authenticated user has permission, Swarm responds with the next step in the 2FA login:

```
HTTP/1.1 200 OK  
  
{  
  "results": {  
    "trigger": "GAuth says yes!",  
    "successMsg": "Second factor authentication approved."  
  },  
  "code": 200  
}
```

Get List of 2FA Methods

Summary

Get List of 2FA Methods

GET /api/v9/listmethods/

Description

Returns the complete list of methods of 2FA.

Example response

Successful Response:

```
HTTP/1.1 200 OK  
  
{  
  "results": {  
    "methods": {
```

```

    "1": {
      "methodName": "Method Name will be here",
      "methodDesc": "Method Description will be here"
    },
    "2": {
      "methodName": "Method Name will be here",
      "methodDesc": "Method Description will be here"
    },
    "3": {
      "methodName": "Method Name will be here",
      "methodDesc": "Method Description will be here"
    },
    "4": {
      "methodName": "Method Name will be here",
      "methodDesc": "Method Description will be here"
    }
  },
  "option": {
    "persist": "option",
    "nextState": "init-auth"
  },
  "code": 200
}

```

Usage example

Listing 2FA Methods

To list the 2FA methods:

```

curl -u "username:password"
"https://myswarm.url/api/v9/login/listmethods"

```

Swarm responds with a list of 2FA methods:

```

HTTP/1.1 200 OK

{

```

```
"results": {
  "methods": {
    "1": {
      "methodName": "Method Name will be here",
      "methodDesc": "Method Description will be here"
    },
    "2": {
      "methodName": "Method Name will be here",
      "methodDesc": "Method Description will be here"
    },
    "3": {
      "methodName": "Method Name will be here",
      "methodDesc": "Method Description will be here"
    },
    "4": {
      "methodName": "Method Name will be here",
      "methodDesc": "Method Description will be here"
    }
  },
  "option": {
    "persist": "option",
    "nextState": "init-auth"
  },
  "code": 200
}
```

Get the current effective user details

Summary

Get the current effective user details

GET /api/v9/session

Description

Get user logged in

Example response

Successful Response:

```
HTTP/1.1 200 OK

{
  "isValid": true,
  "messages": [],
  "user": {
    "User": "reviewer",
    "FullName": "Code Reviewer",
    "Email": "reviewer@swarm.local",
    "Type": "standard",
    "Password": "enabled"
  }
}
```

Usage example

Getting the currently effective user details

To get session details:

```
curl -u "<username>:<ticket>" "http://myswarm.url/api/v9/session"
```

Checking the 2FA authentication

POST /api/v9/checkauth/

Summary

Checking the 2FA authentication

Description

Checking the 2FA authentication

Parameters

Parameter	Description	Type	Parameter Type	Required
token	The token from the user for their 2FA prompt.	string	form	Yes

Example response

Successful Response:

```
HTTP/1.1 200 OK

{
  "results": {
    "trigger": "otp-generated||GAAuth says yes!",
    "successMsg": "Second factor authentication approved."
  },
  "code": 200
}
```

Example usage

Checking the 2FA authentication

To Check User prompt input for 2FA

```
curl -u "username:password" \
  -X POST \
  -d "token=TOKEN" \
  "https://myswarm.url/api/v9/login/checkauth"
```

Assuming that the authenticated user has permission, Swarm responds with the next step in the 2FA login:

```
HTTP/1.1 200 OK

{
  "results": {
    "trigger": "GAAuth says yes!",
```

```
    "successMsg": "Second factor authentication approved."
  },
  "code": 200
}
```

Initiating the 2FA authentication

Summary

Initiating the 2FA authentication

POST /api/v9/initauth/

Description

Initiating the 2FA authentication

Parameters

Parameter	Description	Type	Parameter Type	Required
method	The Method in which you want to use.	string	form	Yes

Example response

Successful Response:

HTTP/1.1 200 OK

```
{
  "results": {
    "trigger": "TriggerName",
    "successMsg": "Message from Authentication method"
  },
  "option": {
    "prompt": true,
    "nextState": "check-auth"
  },
}
```

```
"code": 200
}
```

Example usage

Initiating the 2FA authentication

To Initiate the user 2FA login:

```
curl -u "username:password" \
  -X POST
  -d "method=METHOD" \
  "https://myswarm.url/api/v9/login/initauth"
```

Assuming that the authenticated user has permission, Swarm responds with the next step in the 2FA login:

```
HTTP/1.1 200 OK

{
  "results": {
    "trigger": "TriggerName",
    "successMsg": "Message from Authentication method"
  },
  "option": {
    "prompt": true,
    "nextState": "check-auth"
  },
  "code": 200
}
```

Login to Swarm

Summary

Login to Swarm

POST /api/v9/login/

Description

Login to Swarm

Example response

Successful Response:

```
HTTP/1.1 200 OK

{
  "isValid": true,
  "messages": [],
  "user": {
    "User": "swarm.user",
    "FullName": "Swarm User",
    "Email": "swarm.user@mydomain.com",
    "Type": "standard",
    "Password": "enabled",
    "isAdmin": false,
    "isSuper": false
  }
}
```

Note

In the event of a failed login attempt Swarm responds with:

Example usage

Logging in to swarm

To login:

```
curl -H "Content-Type: application/json" \
  -X POST \
  -u "super:<ticket>" \
  -d '{"username":"swarm.user","password":"1234"}' \
  "http://myswarm.url/api/v9/login"
```

Login to Swarm with SAML

Important

Swarm now supports Helix Authentication Service (HAS) as a Single Sign-On (SSO) provider. This helps to simplify configuration and create a more robust SSO solution. The [SAML](#) configuration will be removed from Swarm in a later release.

If [SAML](#) is configured when the "sso" on [page 196](#) configurable is set to optional, the user can either log in to Swarm using the [Log in with SSO](#) button or the [Log in with credentials](#) button. If [SAML](#) is configured when the "sso" on [page 196](#) configurable is enabled, the user can only log in to Swarm using the [Log in with SSO](#) button.

Summary

Login to Swarm with SAML

POST /api/v9/login/saml

Description

Login to Swarm with SAML

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>redirect</code>	<p>Options are:</p> <ul style="list-style-type: none">▪ <code>redirect=true</code> or not specified: Swarm redirects the user to the <code>HTTP_REFERER</code> url or to the specified custom "logout_url" on page 612 if it has been set.▪ <code>redirect=false</code>: Swarm does not redirect the user.	string	query	No

Example usage

Logging in to swarm with SAML

```
curl -u "super:<ticket>" \
  -X POST \
  "http://myswarm.url/api/v9/login/saml"
```

JSON response:

```
HTTP/1.1 302 OK

{
  "isValid": "true"
  "url": "<url to redirect to>"
}
```

Logging in to swarm with SAML and redirect=false

```
curl -u "super:<ticket>" \
  -X POST \
  "http://myswarm.url/api/v9/login/saml?redirect=false"
```

JSON response:

```
HTTP/1.1 200 OK

{
  "isValid": "true",
}
```

Logout of Swarm with optional redirect

Summary

Logout of Swarm with optional redirect

POST /api/v9/logout/

Description

Logout of Swarm

Examples responses

Successful Response:

```
HTTP/1.1 302 OK

{
  "isValid": true,
  "messages": []
}
```

Successful Response:

```
HTTP/1.1 200 OK

{
  "isValid": true,
  "messages": []
}
```

Note

In the event of a failed login attempt Swarm responds with:

```
HTTP/1.1 200 OK

{
  "isValid": false,
  "messages": ["Error message."]
}
```

Example usage

Logout of Swarm

To logout:

```
curl -X POST \
  -u "super:<ticket>" \
  "http://myswarm.url/api/v9/logout"
```

Logout of Swarm without redirect

To logout without any redirect:

```
curl -X POST \  
  -u "super:<ticket>" \  
  "http://myswarm.url/api/v9/logout?stay=true"
```

Create a new Swarm session using the given credentials

Summary

Create a new Swarm session using the given credentials

POST /api/v9/session

Description

Login to swarm

Example response

Successful Response:

```
HTTP/1.1 200 OK  
  
{  
  "isValid": true,  
  "messages": [],  
  "user": {  
    "User": "reviewer",  
    "FullName": "Code Reviewer",  
    "Email": "reviewer@swarm.local",  
    "Type": "standard",  
    "Password": "enabled"  
  }  
}
```


Example usage

Logging in to Swarm

To login:

```
curl -H "Content-Type: application/json" \
  -X POST \
  -d '{"username": "<username>", "password":
"<password>", "remember": "false"}' \
  -X POST http://localhost/api/v9/session
```

Destroy the current session, for instance logout

Summary

Destroy the current session, for instance logout

DELETE /api/v9/session

Description

Logout of Swarm

Example response

Successful Response:

```
HTTP/1.1 200 OK

{
  "isValid": true,
  "messages": []
}
```

Example usage

Logging out of swarm

To login:

```
curl -u "<username>:<ticket>" -X DELETE  
"http://myswarm.url/api/v9/session"
```

Projects : Swarm Projects

Important

From Swarm 2022.2, Swarm no longer supports APIs older than v9.

Get List of Projects

Summary

Get List of Projects

GET /api/v9/projects/

Description

Returns a list of projects in Swarm that are visible to the current user. Administrators will see all projects, including private ones.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>fields</code>	An optional comma-separated list (or array) of fields to show for each project. Omitting this parameter or passing an empty value shows all fields.	string	query	No

Parameter	Description	Type	Parameter Type	Required
workflow	An optional parameter to only list projects using a workflow.	string	query	No

Example response

Successful Response:

HTTP/1.1 200 OK

```
{
  "project": {
    "id": "jplugin",
    "branches": [],
    "defaults": {
      "reviewers": {
        "users": {
          "allison.clayborne": {
            "required": true
          },
          "jack.boone": [ ],
          "steve.russell": [ ]
        }
      }
    },
    "deleted": false,
    "deploy": {
      "enabled": false,
      "url": ""
    },
    "description": "A Java plugin for continuous integration.",
    "emailFlags": {
      "change_email_project_users": "1",
      "review_email_project_members": "1"
    },
  },
}
```

```

    "jobview": "",
    "members": [
      "allison.clayborne",
      "jack.boone",
      "steve.russell"
    ],
    "minimumUpVotes": 1,
    "name": "JPlugin",
    "owners": [
      "allison.clayborne"
    ],
    "private": false,
    "retainDefaultReviewers": false,
    "subgroups": [ ],
    "tests": {
      "enabled": false,
      "url": "",
      "postBody": "",
      "postFormat": "URL"
    },
    "workflow": null
  },
  "readme":<h1>P4 Plugin</h1>\n<p>Jenkins plugin for a Perforce Helix
Versioning Engine (P4D).</p>\n<h2>Contents</h2>\n<ul>\n<li><a
href=\"https://github.com/jenkinsci/p4-
plugin/blob/main/RELEASE.md\">Release notes</a></li>\n\">
}

```

Example usage

Listing projects

To list all projects:

```
curl -u "username:password" "https://my-swarm-
host/api/v9/projects?fields=id,description,members,name"
```

Pagination is not currently supported by this endpoint. Swarm responds with a list of all projects:

HTTP/1.1 200 OK

```
{
  "projects": [
    {
      "id": "blue-book",
      "description": "This is a private project for use only by users
with sufficient clearance.",
      "members": [
        "alex.randolph",
        "allison.clayborne"
      ],
      "name": "Blue Book"
    },
    {
      "id": "jplugin",
      "description": "A Java plugin for continuous integration.",
      "members": [
        "allison.clayborne",
        "jack.boone",
        "steve.russell"
      ],
      "name": "JPlugin"
    },
    {
      "id": "mercury",
      "description": "Mercury project for code management.",
      "members": [
        "alex.randolph",
        "claire.brevia",
        "jack.boone"
      ],
      "name": "Mercury"
    },
    {
```

```

    "id": "test-data",
    "description": "",
    "members": [
      "steve.russell"
    ],
    "name": "Test Data"
  },
  {
    "id": "test-project-for-workflow",
    "description": "This project has been created to test Swarm
workflow.",
    "members": [
      "alex.randolph",
      "allison.clayborne",
      "claire.brevia",
      "jack.boone",
      "paula.boyle",
      "steve.russell"
    ],
    "name": "Test project for workflow"
  }
]
}

```

Project administrators wishing to see the **tests** and **deploy** fields must fetch projects individually.

Get Project Information

Summary

Get Project Information

GET /api/v9/projects/{id}

Description

Retrieve information about a project.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Project ID	string	path	Yes
<code>fields</code>	An optional comma-separated list (or array) of fields to show for each project. Omitting this parameter or passing an empty value shows all fields.	string	query	No

Example response**Successful Response:**

HTTP/1.1 200 OK

```
{
  "project": {
    "id": "jplugin",
    "branches": [
      {
        "id": "main",
        "name": "MAIN",
        "workflow": null,
        "paths": [
          "//projects/jplugin/main/..."
        ],
        "defaults": {
```

```
    "reviewers": {
      "users": {
        "steve.russell": [ ]
      }
    },
    "minimumUpVotes": null,
    "retainDefaultReviewers": false,
    "moderators": [
      "claire.brevia"
    ],
    "moderators-groups": [ ]
  },
  {
    "id": "candidate",
    "name": "CANDIDATE",
    "workflow": null,
    "paths": [
      "//projects/jplugin/candidate/..."
    ],
    "defaults": {
      "reviewers": [ ]
    },
    "minimumUpVotes": null,
    "retainDefaultReviewers": false,
    "moderators": [ ],
    "moderators-groups": [ ]
  }
],
"defaults": {
  "reviewers": {
    "users": {
      "allison.clayborne": {
        "required": true
      },

```



```
        "jack.boone": [ ],
        "steve.russell": [ ]
    }
}
},
"deleted": false,
"deploy": {
    "enabled": false,
    "url": ""
},
"description": "A Java plugin for continuous integration.",
"emailFlags": {
    "change_email_project_users": "1",
    "review_email_project_members": "1"
},
"jobview": "",
"members": [
    "allison.clayborne",
    "jack.boone",
    "steve.russell"
],
"minimumUpVotes": 1,
"name": "JPlugin",
"owners": [
    "allison.clayborne"
],
"private": false,
"retainDefaultReviewers": false,
"subgroups": [ ],
"tests": {
    "enabled": false,
    "url": "",
    "postBody": "",
    "postFormat": "URL"
},
}
```

```

    "workflow": null
  },
  "readme":<h1>P4 Plugin</h1>\n<p>Jenkins plugin for a Perforce Helix
Versioning Engine (P4D).</p>\n<h2>Contents</h2>\n<ul>\n<li><a
href=\"https://github.com/jenkinsci/p4-
plugin/blob/main/RELEASE.md\">Release notes</a></li>\n
}

```

Example usage

Fetching a project

To fetch an individual project:

```
curl -u "username:password" "https://my-swarm-
host/api/v9/projects/testproject2?fields=id,description,members,name"
```

Swarm responds with a project entity:

```

HTTP/1.1 200 OK

{
  "project": {
    "id": "jplugin",
    "description": "A Java plugin for continuous integration.",
    "members": [
      "allison.clayborne",
      "jack.boone",
      "steve.russell"
    ],
    "name": "JPlugin"
  }
}

```

Project administrators have access to additional fields (**tests** and **deploy**) when fetching individual projects using this endpoint.

Create a new Project

Summary

Create a new Project

POST /api/v9/projects/

Description

Creates a new project in Swarm.

Parameters

Parameter	Description	Type	Parameter Type	Required
name	Project Name (is also used to generate the Project ID)	string	form	Yes
members	An array of project members.	array	form	Yes
subgroups	An optional array of project subgroups.	array	form	No
owners	An optional array of project owners.	array	form	No
description	An optional project description.	string	form	No
private	Private projects are visible only to Members, Moderators, Owners, and Administrators. (Default: false)	boolean	form	No
deploy	Configuration for automated deployment. Example: {"enabled": true, "url": "http://localhost/?change={change}"}	array	form	No

Parameter	Description	Type	Parameter Type	Required
tests	Configuration for testing/continuous integration.	array	form	No
branches	Optional branch definitions for this project.	array	form	No
jobview	An optional jobview for associating certain jobs with this project.	string	form	No
emailFlags[change_email_project_users]	Email members, moderators and followers when a change is committed.	boolean	form	No
emailFlags[review_email_project_members]	Email members and moderators when a new review is requested.	boolean	form	No
defaults	An optional array of defaults at a project level (for example default reviewers).	array	form	No
retainDefaultReviewers	An optional to retain the default reviewers.	boolean	form	No
minimumUpVotes	An optional to set the minimum number of up votes required.	string	form	No

Example responses

Successful Response:

HTTP/1.1 200 OK

```
{
  "project": {
    "id": "testproject",
    "branches": [ ],
```

```
"defaults": {
  "reviewers": [ ]
},
"deleted": false,
"deploy": {
  "enabled": false,
  "url": null
},
"description": "My project is great",
"emailFlags": [ ],
"jobview": null,
"members": [
  "bruno"
],
"minimumUpVotes": null,
"name": "testProject",
"owners": [ ],
"private": false,
"retainDefaultReviewers": false,
"subgroups": [ ],
"tests": {
  "enabled": false,
  "url": null
},
"workflow": null
},
"readme": "",
"mode": "add"
}
```

Successful Response:

```
HTTP/1.1 200 OK
```

```
{
  "project": {
```

```
"id": "project-x",
"branches": [ ],
"defaults": {
  "reviewers": {
    "users": {
      "bruno": {
        "required": "1"
      }
    }
  }
},
"deleted": false,
"deploy": {
  "enabled": false,
  "url": null
},
"description": "This project is temporary until project Z replaces
it",
"emailFlags": [ ],
"jobview": null,
"members": [
  "swarm"
],
"minimumUpVotes": null,
"name": "project x",
"owners": [ ],
"private": false,
"retainDefaultReviewers": false,
"subgroups": [ ],
"tests": {
  "enabled": false,
  "url": null
},
"workflow": null
},
```

```
"readme": "",  
"mode": "add"  
}
```

Example usage

Creating a new project

To create a project:

```
curl -u "username:password" \  
  -X POST \  
  -d "name=TestProject 3" \  
  -d "description=The third iteration of our test project." \  
  -d "members[]=alice" \  
  -d "members[]=bob" \  
  "https://my-swarm-host/api/v9/projects/"
```

Swarm responds with the new project entity:

```
HTTP/1.1 200 OK  
  
{  
  "project": {  
    "id": "testproject-3",  
    "branches": [ ],  
    "defaults": {  
      "reviewers": [ ]  
    },  
    "deleted": false,  
    "deploy": {  
      "enabled": false,  
      "url": null  
    },  
    "description": "The third iteration of our test project.",  
    "emailFlags": [ ],  
    "jobview": null,  
    "members": [  
      "bruno",
```

```

    "swarm"
  ],
  "minimumUpVotes": null,
  "name": "TestProject 3",
  "owners": [ ],
  "private": false,
  "retainDefaultReviewers": false,
  "subgroups": [ ],
  "tests": {
    "enabled": false,
    "url": null
  },
  "workflow": null
},
"readme": "",
"mode": "add"
}

```

Creating a private project with branches

Specifying a branch requires using array notation and providing at least two fields (**name** and **paths**) for each branch you wish to create. Creating more than one branch involves incrementing the **branches [0]** specifier for each branch - an example of this accompanies the PATCH endpoint documentation.

Projects are public by default. To mark a project as private, set the **private** parameter to **true**.

```

curl -u "username:password" \
  -X POST \
  -d "name=TestProject 4" \
  -d "description=The 4th iteration of our test project." \
  -d "private=true" \
  -d "members[]=bob" \
  -d "branches[0][name]=Branch One" \
  -d "branches[0][paths][]=//depot/main/TestProject/..." \
  "https://my-swarm-host/api/v9/projects"

```

Swarm responds with the new project entity:

HTTP/1.1 200 OK

```
{
  "project": {
    "id": "testproject-4",
    "branches": [ ],
    "defaults": {
      "reviewers": [ ]
    },
    "deleted": false,
    "deploy": {
      "enabled": false,
      "url": null
    },
    "description": "The 4th iteration of our test project.",
    "emailFlags": [ ],
    "jobview": null,
    "members": [
      "bruno"
    ],
    "minimumUpVotes": null,
    "name": "TestProject 4",
    "owners": [ ],
    "private": true,
    "retainDefaultReviewers": false,
    "subgroups": [ ],
    "tests": {
      "enabled": false,
      "url": null
    },
    "workflow": null
  },
  "readme": "",
  "mode": "add"
}
```

Edit a Project

Summary

Edit a Project

PATCH /api/v9/projects/{id}

Description

Change the settings of a project in Swarm. If a project has owners set, only the owners can perform this action.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Project ID	string	path	Yes
<code>name</code>	Project Name (changing the project name does not change the project ID)	string	form	No
<code>members</code>	An array of project members.	array	form	No
<code>subgroups</code>	An optional array of project subgroups.	array	form	No
<code>owners</code>	An optional array of project owners.	array	form	No
<code>description</code>	Your project description.	string	form	No

Parameter	Description	Type	Parameter Type	Required
private	Private projects are visible only to Members, Moderators, Owners, and Administrators. (Default: false)	boolean	form	No
deploy	Configuration for automated deployment. Example: {"enabled": true, "url": "http://localhost/?change={change}"}	array	form	No
tests	Configuration for testing/continuous integration.	array	form	No
branches	Optional branch definitions for this project.	array	form	No
jobview	A jobview for associating certain jobs with this project.	string	form	No
emailFlags[change_email_project_users]	Email members, moderators and followers when a change is committed.	boolean	form	No
emailFlags[review_email_project_members]	Email members and moderators when a new review is requested.	boolean	form	No
defaults	An optional array of defaults at a project level (for example default reviewers).	array	form	No
retainDefaultReviewers	An optional to retain the default reviewers.	boolean	form	No

Parameter	Description	Type	Parameter Type	Required
minimumUpVotes	An optional to set the minimum number of up votes required.	string	form	No

Example response

Successful Response:

```
HTTP/1.1 200 OK

{
  "project":{
    "id":"project1",
    "branches":[
      {
        "name":"main",
        "paths":[
          "//depot/dev/APIProject/..."
        ],
        "id":"main",
        "moderators":[ ],
        "moderators-groups":[ ],
        "defaults":{
          "reviewers":[ ]
        },
        "workflow":null,
        "retainDefaultReviewers":false,
        "minimumUpVotes":null
      }
    ],
    "defaults":{
      "reviewers":[ ]
    },
    "deleted":false,
    "deploy":{
```

```
    "enabled":false,
    "url":""
  },
  "description":"Updated description",
  "emailFlags":[ ],
  "jobview":"","
  "members":[
    "swarm-admin"
  ],
  "minimumUpVotes":null,
  "name":"API Project",
  "owners":[ ],
  "private":false,
  "retainDefaultReviewers":false,
  "subgroups":[ ],
  "tests":{
    "enabled":false,
    "url":""
  },
  "workflow":null
},
"readme":"","
"mode":"edit"
}
```

Example usage

Editing a project

To edit a project:

Note

It is safe to edit a project without specifying branches, but the instructions for adding branches contain important information for modifying branch configuration.

```
curl -u "username:password" \
  -X PATCH
  -d "description=Witness the power of a fully operational Swarm"
```

```
project." \
  "https://my-swarm-host/api/v9/projects/jam"
```

Swarm responds with the updated project entity:

```
HTTP/1.1 200 OK

{
  "project":{
    "id":"jam",
    "branches":[
      {
        "id":"main",
        "name":"Main",
        "workflow":null,
        "paths":[
          "\\depot\\Jam\\MAIN\\..."
        ],
        "defaults":{
          "reviewers":[ ]
        },
        "minimumUpVotes":null,
        "retainDefaultReviewers":false,
        "moderators":[ ],
        "moderators-groups":[ ]
      },
      {
        "id":"release-2-1",
        "name":"Release 2.1",
        "workflow":null,
        "paths":[
          "\\depot\\Jam\\REL2.1\\..."
        ],
        "defaults":{
          "reviewers":[ ]
        },
        "minimumUpVotes":null,
```

```
    "retainDefaultReviewers":false,
    "moderators":[ ],
    "moderators-groups":[ ]
  },
  {
    "id":"release-2-2",
    "name":"Release 2.2",
    "workflow":null,
    "paths":[
      "\\depot\\Jam\\REL2.2\\..."
    ],
    "defaults":{
      "reviewers":[ ]
    },
    "minimumUpVotes":null,
    "retainDefaultReviewers":false,
    "moderators":[ ],
    "moderators-groups":[ ]
  }
],
"defaults":{
  "reviewers":[ ]
},
"deleted":false,
"deploy":{
  "enabled":false,
  "url":""
},
"description":"Witness the power of a fully operational Swarm
project.",
"emailFlags":{
  "change_email_project_users":"1",
  "review_email_project_members":"1"
},
"jobview":"","
```

```

    "members": [
      "alex_gc",
      "pubuser"
    ],
    "minimumUpVotes": null,
    "name": "Jam",
    "owners": [
      "bruno"
    ],
    "private": false,
    "retainDefaultReviewers": false,
    "subgroups": [
      "Administrators"
    ],
    "tests": {
      "enabled": false,
      "url": "",
      "postBody": "",
      "postFormat": "URL"
    },
    "workflow": "1"
  },
  "readme": "",
  "mode": "edit"
}

```

Editing a project to add a moderated branch and make the project public

Specifying a branch requires using array notation and providing at least two fields (**name** and **paths**) for each branch you wish to create. Creating more than one branch involves incrementing the **branches [0]** specifier for each branch.

Important

If you have existing branches, you must specify all of them in the query to avoid data loss. This operation sets the value of the entire **branches** property to match the provided input.

To change a private project to public, set the **private** parameter to **false**.


```
curl -u "username:password" \  
  -X PATCH \  
  -d "private=false" \  
  -d "branches[0][name]=Branch One" \  
  -d "branches[0][paths][]=//depot/main/TestProject/..." \  
  -d "branches[1][name]=Branch Two" \  
  -d "branches[1][paths][]=//depot/main/SecondBranch/..." \  
  -d "branches[1][moderators][]=bob" \  
  -d "branches[1][moderators-groups][]=group1" \  
  "https://my-swarm-host/api/v9/projects/testproject-4"
```

Swarm responds with the new project entity:

```
HTTP/1.1 200 OK  
  
{  
  "project":{  
    "id":"testproject-4",  
    "branches":[  
      {  
        "name":"Branch One",  
        "paths":[  
          "//depot/main/TestProject/..."  
        ],  
        "id":"branch-one",  
        "moderators":[ ],  
        "moderators-groups":[ ],  
        "defaults":{  
          "reviewers":[ ]  
        },  
        "workflow":null,  
        "retainDefaultReviewers":false,  
        "minimumUpVotes":null  
      },  
      {  
        "name":"Branch Two",  
        "paths":[]  
      }  
    ]  
  }  
}
```

```
        "//depot/main/SecondBranch/..."
    ],
    "moderators":[
        "non-admin"
    ],
    "moderators-groups":[
        "group1"
    ],
    "id":"branch-two",
    "defaults":{
        "reviewers":[ ]
    },
    "workflow":null,
    "retainDefaultReviewers":false,
    "minimumUpVotes":null
    }
],
"defaults":{
    "reviewers":[ ]
},
"deleted":false,
"deploy":{
    "enabled":false,
    "url":""
},
"description":"Updated description",
"emailFlags":[ ],
"jobview":null,
"members":[
    "non-admin"
],
"minimumUpVotes":null,
"name":"testproject-4",
"owners":[ ],
"private":false,
```

```
"retainDefaultReviewers":false,
"subgroups":[ ],
"tests":{
  "enabled":false,
  "url":""
},
"workflow":null
},
"readme":"","
"mode":"edit"
}
```

Editing a project to add default reviewers

```
curl -u "<user>:<password>" \
  -H "Content-Type: application/json" \
  -X PATCH \
  -d '{"defaults":{"reviewers":{"swarm_admin":{"required":true}} \
    "branches":[{" \
      "name":"dev", \
      "paths":{"//depot/dev/API Project/..."} \
      "defaults": { \
        "reviewers":{" \
          "groups":{"group1":{"required":"1"}}, \
          "users":{"non_admin":{"required":true}} \
        } \
      } \
    ]}' \
  "https://my-swarm-host/api/v9/projects/API Project"
```

Or without JSON content:

```
curl -u "username:password" \
  -X PATCH \
  -d "name=Udated description" \
  -d "defaults[reviewers][swarm_admin][required]=true" \
  -d "branches[0][name]=dev" \
  -d "branches[0][paths][]=//depot/dev/API Project/..." \
  -d "branches[0][defaults][reviewers][groups][group1][required]=1"
```

```
\
  -d "branches[0][defaults][reviewers][users][non_admin]
[required]=true" \
  "https://my-swarm-host/api/v9/projects/API Project"
```

Swarm responds with project entity similar to:

```
HTTP/1.1 200 OK

{
  "project":{
    "id":"api-project",
    "branches":[
      {
        "name":"dev",
        "paths":[
          "//depot/dev/APIProject/..."
        ],
        "defaults":{
          "reviewers":{
            "users":{
              "non-admin":{
                "required":true
              }
            },
            "groups":{
              "group1":{
                "required":"1"
              }
            }
          },
          "id":"dev",
          "moderators":[ ],
          "moderators-groups":[ ],
          "workflow":null,
          "retainDefaultReviewers":false,
```

```
        "minimumUpVotes":null
      }
    ],
    "defaults":{
      "reviewers":{
        "users":{
          "swarm-admin":{
            "required":true
          }
        }
      }
    },
    "deleted":false,
    "deploy":{
      "enabled":false,
      "url":""
    },
    "description":"Updated description",
    "emailFlags":[ ],
    "jobview":null,
    "members":[
      "swarm-admin"
    ],
    "minimumUpVotes":null,
    "name":"API Project",
    "owners":[ ],
    "private":false,
    "retainDefaultReviewers":false,
    "subgroups":[ ],
    "tests":{
      "enabled":false,
      "url":""
    },
    "workflow":null
  },
```

```
"readme":"","  
"mode":"edit"  
}
```

Delete a Project

Summary

Delete a Project

DELETE /api/v9/projects/{id}

Description

Mark a Swarm project as deleted. The project ID and name cannot be reused. If a project has owners set, only the owners can perform this action.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Project ID	string	path	Yes

Example response

Successful Response:

```
HTTP/1.1 200 OK  
  
{  
  "id": "testproject"  
}
```

Example usage

Deleting a project:

Super users, administrators, and owners can delete projects. Members can delete projects that have no owners set.

```
curl -u "username:password" -X DELETE "https://my-swarm-  
host/api/v9/projects/testproject3"
```

Assuming that the authenticated user has permission, Swarm responds with the id of the deleted project:

```
HTTP/1.1 200 OK  
  
{  
  "id": "testproject3"  
}
```

Reviews : Swarm Reviews

Important

From Swarm 2022.2, Swarm no longer supports APIs older than v9.

Get reviews for action dashboard

Summary

Get reviews for action dashboard

GET /api/v9/dashboards/action

Description

Gets reviews for the action dashboard for the authenticated user.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).

- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Example response

Successful Commit contains Review and Commit Entities:

HTTP/1.1 200 OK

```
{
  "lastSeen": 120,
  "reviews": [
    {
      "id": 7,
      "author": "swarm_admin",
      "changes": [6],
      "comments": [0,0],
      "commits": [6],
      "commitStatus": [],
      "created": 1485793976,
      "deployDetails": [],
      "deployStatus": null,
      "description": "test\n",
      "groups": ["swarm-project-test"],
      "participants": {"swarm_admin":[]},
      "pending": false,
      "projects": {"test":["test"]},
      "roles": ["moderator|reviewer|required_reviewer|author"],
      "state": "needsReview",
      "stateLabel": "Needs Review",
      "testDetails": [],
      "testStatus": null,
      "type": "default",
      "updated": 1485958875,
      "updateDate": "2017-02-01T06:21:15-08:00"
    }
  ]
}
```



```
],  
  "totalCount": null  
}
```

Example usage

Getting reviews for the action dashboard

To list reviews:

```
curl -u "username:password" "http://my-swarm-  
host/api/v9/dashboards/action"
```

Swarm responds with a list of the latest reviews, a **totalCount** field, and a **lastSeen** value for pagination:

```
HTTP/1.1 200 OK  
  
{  
  "lastSeen": 120,  
  "reviews": [  
    {  
      "id": 7,  
      "author": "swarm_admin",  
      "changes": [6],  
      "comments": [0,0],  
      "commits": [6],  
      "commitStatus": [],  
      "created": 1485793976,  
      "deployDetails": [],  
      "deployStatus": null,  
      "description": "test\n",  
      "groups": ["swarm-project-test"],  
      "participants": {"swarm_admin": []},  
      "pending": false,  
      "projects": {"test": ["test"]},  
      "roles": ["moderator|reviewer|required_reviewer|author"],  
      "state": "needsReview",  
      "stateLabel": "Needs Review",
```

```
    "testDetails": [],
    "testStatus": null,
    "type": "default",
    "updated": 1485958875,
    "updateDate": "2017-02-01T06:21:15-08:00"
  }
],
"totalCount": null
}
```

Get List of Reviews

Summary

Get List of Reviews

GET /api/v9/reviews/

Description

List and optionally filter reviews.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required	Default Value
<code>after</code>	<p>A review ID to seek to.</p> <p>Reviews up to and including the specified <code>id</code> are excluded from the results and do not count towards <code>max</code>. Useful for pagination. Commonly set to the <code>lastSeen</code> property from a previous query.</p> <p>Reviews are returned in the order they were created in, starting with the most recently created review.</p>	integer	query	No	

Parameter	Description	Type	Parameter Type	Required	Default Value
max	Maximum number of reviews to return. This does not guarantee that max reviews are returned. It does guarantee that the number of reviews returned won't exceed max . Server-side filtering may exclude some reviews for permissions reasons.	integer	query	No	1000
fields	An optional comma-separated list (or array) of fields to show. Omitting this parameter or passing an empty value shows all fields.	string	query	No	
author []	One or more authors to limit reviews by. Reviews with any of the specified authors are returned. (v1.2+)	array (of strings)	query	No	

Parameter	Description	Type	Parameter Type	Required	Default Value
<code>change []</code>	One or more change IDs to limit reviews by. Reviews associated with any of the specified changes are returned.	array (of integers)	query	No	
<code>hasReviewers</code>	Boolean option to limit to reviews to those with or without reviewers. Use <code>true</code> or <code>false</code> for JSON-encoded data, <code>1</code> for true and <code>0</code> for false for form-encoded data. The presence of the parameter without a value is evaluated as true.	boolean	query	No	
<code>ids []</code>	One or more review IDs to fetch. Only the specified reviews are returned. This filter cannot be combined with the <code>max</code> parameter.	array (of integers)	query	No	

Parameter	Description	Type	Parameter Type	Required	Default Value
keywords	Keywords to limit reviews by. Only reviews where the description, participants list or project list contain the specified keywords are returned.	string	query	No	
participants[]	One or more participants to limit reviews by. Reviews with any of the specified participants are returned.	array (of strings)	query	No	
project[]	One or more projects to limit reviews by. Reviews affecting any of the specified projects are returned.	array (of strings)	query	No	
state[]	One or more states to limit reviews by. Reviews in any of the specified states are returned.	array (of strings)	query	No	

Parameter	Description	Type	Parameter Type	Required	Default Value
<code>passesTests</code>	Boolean option to limit reviews by tests passing or failing. Use <code>true</code> or <code>false</code> for JSON-encoded data, <code>1</code> for true and <code>0</code> for false for form-encoded data. The presence of the parameter without a value is evaluated as true.	string	query	No	
<code>notUpdatedSince</code>	Option to fetch unchanged reviews. Requires the date to be in the format YYYY-mm-dd, for example 2017-01-01. Reviews to be returned are determined by looking at the last updated date of the review.	string	query	No	

Parameter	Description	Type	Parameter Type	Required	Default Value
hasVoted	Should have the value 'up' or 'down' to filter reviews that have been voted up or down by the current authenticated user.	string	query	No	
myComments	True or false to support filtering reviews that include comments by the current authenticated user.	boolean	query	No	

Examples responses

Successful Response:

HTTP/1.1 200 OK

```
{
  "lastSeen": 12209,
  "reviews": [
    {
      "id": 12206,
      "author": "swarm",
      "changes": [12205],
      "comments": 0,
      "commits": [],
      "commitStatus": [],
      "created": 1402507043,
      "deployDetails": [],
```



```
    "deployStatus": null,  
    "description": "Review Description\n",  
    "participants": {  
      "swarm": []  
    },  
    "pending": true,  
    "projects": [],  
    "state": "needsReview",  
    "stateLabel": "Needs Review",  
    "testDetails": [],  
    "testStatus": null,  
    "type": "default",  
    "updated": 1402518492  
  }  
],  
  "totalCount": 1  
}
```

Note

Swarm returns **null** for **totalCount** if no search filters were provided. **lastSeen** can often be used as an offset for pagination, by using the value in the **after** parameter of subsequent requests.

When no results are found, the **reviews** array is empty:

```
HTTP/1.1 200 OK  
  
{  
  "lastSeen": null,  
  "reviews": [],  
  "totalCount": 0  
}
```

Example usage

Listing reviews

To list reviews:

```
curl -u "username:password" "https://my-swarm-
host/api/v9/reviews?max=2&fields=id,description,author,state"
```

Swarm responds with a list of the latest reviews, a **totalCount** field, and a **lastSeen** value for pagination:

```
HTTP/1.1 200 OK
```

```
{
  "lastSeen": 120,
  "reviews": [
    {
      "id": 123,
      "author": "bruno",
      "description": "Adding .jar that should have been included in
r110\n",
      "state": "needsReview"
    },
    {
      "id": 120,
      "author": "bruno",
      "description": "Fixing a typo.\n",
      "state": "needsReview"
    }
  ],
  "totalCount": null
}
```

The **totalCount** field is populated when keywords are supplied. It indicates how many total matches there are. If keywords are not supplied the **totalCount** field remains **null**, indicating that the list of all reviews is being queried.

Paginating a review listing

To obtain the next page of a reviews list (based on the previous example):

```
curl -u "username:password" "https://my-swarm-
host/api/
v9/reviews?max=2&fields=id,description,author,state&after=120"
```

Swarm responds with the second page of results, if any reviews are present after the last seen review:

```
HTTP/1.1 200 OK

{
  "lastSeen": 100,
  "reviews": [
    {
      "id": 110,
      "author": "bruno",
      "description": "Updating Java files\n",
      "state": "needsReview"
    },
    {
      "id": 100,
      "author": "bruno",
      "description": "Marketing materials for our new cutting-edge
product\n",
      "state": "needsReview"
    }
  ],
  "totalCount": null
}
```

Finding reviews for a change or a list of changes

Given a list of change IDs (5, 6, 7), here is how to check if any of them have reviews attached:

```
curl -u "username:password" "https://my-swarm-
host/api/
v9/reviews?max=2&fields=id,changes,description,author,state&change\
[\]=5&change\[\]=6&change\[\]=7"
```

Swarm responds with a list of reviews that include these changes:

```
HTTP/1.1 200 OK

{
  "lastSeen": 100,
  "reviews": [
    {
```

```

    "id": 110,
    "author": "bruno",
    "changes": [5],
    "description": "Updating Java files\n",
    "state": "needsReview"
  },
  {
    "id": 100,
    "author": "bruno",
    "changes": [6,7],
    "description": "Marketing materials for our new cutting-edge
product\n",
    "state": "needsReview"
  }
],
"totalCount": 2
}

```

If no corresponding reviews are found, Swarm responds with an empty reviews list:

```

HTTP/1.1 200 OK

{
  "lastSeen": null,
  "reviews": [],
  "totalCount": 0
}

```

Finding inactive reviews (by checking the last updated date)

```

curl -u "username:password" "https://my-swarm-host/api/v9
/reviews?max=2&fields=id,changes,description,author,state&notUpdatedSi
nce=2017-01-01"

```

Swarm responds with a list of reviews that have not been updated since the notUpdatedSince date:

```

HTTP/1.1 200 OK

{
  "lastSeen": 100,

```

```
"reviews": [
  {
    "id": 110,
    "author": "bruno",
    "changes": [5],
    "description": "Updating Java files\n",
    "state": "needsReview"
  },
  {
    "id": 100,
    "author": "bruno",
    "changes": [6,7],
    "description": "Marketing materials for our new cutting-edge
product\n",
    "state": "needsReview"
  }
],
"totalCount": 2
}
```

Finding reviews I have voted up

```
curl -u "username:password" "https://my-swarm-
host/api/
v9
/reviews?max=2&fields=id,changes,description,author,state&hasVoted=up"
```

Swarm responds with a list of reviews that include these changes:

```
HTTP/1.1 200 OK

{
  "lastSeen": 100,
  "reviews": [
    {
      "id": 110,
      "author": "bruno",
      "changes": [5],
```

```

    "description": "Updating Java files\n",
    "state": "needsReview"
  },
  {
    "id": 100,
    "author": "bruno",
    "changes": [6,7],
    "description": "Marketing materials for our new cutting-edge
product\n",
    "state": "needsReview"
  }
],
"totalCount": 2
}

```

If no corresponding reviews are found, Swarm responds with an empty reviews list:

```

{
  "lastSeen": null,
  "reviews": [],
  "totalCount": 0
}

```

Finding reviews I have commented on (current authenticated user)

```

curl -u "username:password" "https://my-swarm-host/api/v9
/reviews?max=2&fields=id,changes,description,author,state&myComments=t
rue"

```

Swarm responds with a list of reviews that include these changes:

```

HTTP/1.1 200 OK

{
  "lastSeen": 100,
  "reviews": [
    {
      "id": 110,
      "author": "bruno",
      "changes": [5],

```

```
    "description": "Updating Java files\n",
    "state": "needsReview"
  },
  {
    "id": 100,
    "author": "bruno",
    "changes": [6,7],
    "description": "Marketing materials for our new cutting-edge
product\n",
    "state": "needsReview"
  }
],
"totalCount": 2
}
```

If no corresponding reviews are found, Swarm responds with an empty reviews list:

```
HTTP/1.1 200 OK

{
  "lastSeen": null,
  "reviews": [],
  "totalCount": 0
}
```

Get Review Information

Summary

Get Review Information

GET /api/v9/reviews/{id}

Description

Retrieve information about a review.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Review ID	integer	path	Yes
<code>fields</code>	An optional comma-separated list (or array) of fields to show. Omitting this parameter or passing an empty value shows all fields.	string	query	No

Example usage

Fetching a review

To fetch a review:

```
curl -u "username:password" "https://my-swarm-host/api/v9/reviews/123"
```

Swarm responds with a review entity:

```
HTTP/1.1 200 OK
```

```
{
  "review": {
    "id": 123,
    "author": "bruno",
    "changes": [122,124],
    "commits": [124],
    "commitStatus": [],
    "created": 1399325913,
```



```
"deployDetails": [],
"deployStatus": null,
"description": "Adding .jar that should have been included in
r110\n",
"groups": [],
"participants": {
  "alex_qc": [],
  "bruno": {
    "vote": 1,
    "required": true
  },
  "vera": []
},
"reviewerGroups": {
  "group1" : [],
  "group2" : {
    "required" : true
  },
  "group3" : {
    "required" : true,
    "quorum": "1"
  }
},
"pending": false,
"projects": {
  "swarm": ["main"]
},
"state": "archived",
"stateLabel": "Archived",
"testDetails": {
  "url": "http://jenkins.example.com/job/project_ci/123/"
},
"testStatus": null,
"type": "default",
"updated": 1399325913,
```

```
"versions": [  
  {  
    "difference": 1,  
    "stream": "//jam/main",  
    "streamSpecDifference": 0,  
    "change": 124,  
    "user": "bruno",  
    "time": 1568115902,  
    "pending": true,  
    "addChangeMode": "replace",  
    "testRuns": [  
      8  
    ]  
  },  
  {  
    "difference": 1,  
    "stream": "//jam/main",  
    "streamSpecDifference": 0,  
    "change": 12157,  
    "user": "bruno",  
    "time": 1568115918,  
    "pending": true,  
    "addChangeMode": "replace",  
    "archiveChange": 122,  
    "testRuns": [  
      1,  
      2,  
      3,  
      4,  
      5,  
      6,  
      7  
    ]  
  }  
]
```

```
}  
}
```

Fetching a review that does not exist

If you try to fetch a review that does not exist:

```
curl -u "username:password" "https://my-swarm-  
host/api/v9/reviews/999999"
```

Swarm responds with a 404 response:

```
HTTP/1.1 404 Not Found  
  
{  
  "error": "Not Found"  
}
```

Get transitions for a review

Summary

Get transitions for a review

GET /api/v9/reviews/{id}/transitions

Description

Get the allowed transitions for a review.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
upVoters	A list of users whose vote up will be assumed when determining the transitions. For example if a user has not yet voted but would be the last required vote and asked for possible transitions we would want to include 'approve'	string	query	No

Example response

Successful Response:

```
HTTP/1.1 200 OK
```

```
{
  "isValid": "true",
  "transitions": {
    "needsRevision": "Needs Revision",
    "approved": "Approve",
    "approved:commit": "Approve and Commit",
    "rejected": "Reject",
    "archived": "Archive"
  }
}
```

Example usage

Getting transitions that a review with id 1 can transition to:

```
curl -u "username:password" "https://my-swarm-host/api/v9/reviews/1/transitions?upVoters=bruno"
```

The transitions that are returned depend on the current review state, Swarm setup and the configuration of associated projects and branches

```
HTTP/1.1 200 OK

{
  "isValid": "true"
  "transitions": {
    "needsRevision": "Needs Revision",
    "approved": "Approve",
    "approved:commit": "Approve and Commit",
    "rejected": "Reject",
    "archived": "Archive"
  }
}
```

Create a Review

Summary

Create a Review

POST /api/v9/reviews/

Description

Pass in a changelist ID to create a review. Optionally, you can also provide a description and a list of reviewers.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
change	Change ID to create a review from	integer	form	Yes
description	Description for the new review (defaults to change description)	string	form	No
reviewers	A list of reviewers for the new review	array (of strings)	form	No
requiredReviewers	A list of required reviewers for the new review (v1.1+)	array (of strings)	form	No
reviewerGroups	A list of required reviewers for the new review (v7+)	array	form	No

Example response

Successful Response contains Review Entity:

```
HTTP/1.1 200 OK
```

```
{
  "review": {
    "id": 12205,
    "author": "bruno",
    "changes": [10667],
    "commits": [10667],
    "commitStatus": [],
    "created": 1399325913,
    "deployDetails": [],
```

```
"deployStatus": null,
"description": "Adding .jar that should have been included in
r10145\n",
"participants": {
  "bruno": []
},
"reviewerGroups": {
  "group1" : [],
  "group2" : {
    "required" : true
  },
  "group3" : {
    "required" : true,
    "quorum": "1"
  }
},
"pending": false,
"projects": [],
"state": "archived",
"stateLabel": "Archived",
"testDetails": [],
"testStatus": null,
"type": "default",
"updated": 1399325913
}
```

Example usage

Starting a review

To start a review for a committed change or a non-empty shelved changelist specifying reviewer groups:

```
curl -u "username:password" \
-X POST \
-d "change=122" \
-d "reviewerGroups[0][name]=group1" \
```

```
-d "reviewerGroups[1][name]=group2" \
-d "reviewerGroups[1][required]=true" \
-d "reviewerGroups[2][name]=group3" \
-d "reviewerGroups[2][required]=true" \
-d "reviewerGroups[2][quorum]=1" \
"https://my-swarm-host/api/v9/reviews/"
```

Swarm responds with the new review entity:

```
HTTP/1.1 200 OK
```

```
{
  "review": {
    "id": 123,
    "author": "bruno",
    "changes": [122],
    "commits": [],
    "commitStatus": [],
    "created": 1399325913,
    "deployDetails": [],
    "deployStatus": null,
    "description": "Adding .jar that should have been included in
r110\n",
    "groups": [],
    "participants": {
      "bruno": []
    },
    "reviewerGroups": {
      "group1" : [],
      "group2" : {
        "required" : true
      },
      "group3" : {
        "required" : true,
        "quorum": "1"
      }
    }
  },
}
```



```
"pending": true,
"projects": [],
"state": "needsReview",
"stateLabel": "Needs Review",
"testDetails": [],
"testStatus": null,
"type": "default",
"updated": 1399325913,
"versions": []
}
}
```

Archiving the inactive reviews

Summary

Archiving the inactive reviews (v6+)

POST /api/v9/reviews/archive/

Description

Archiving reviews not updated since the date (v6+).

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>notUpdatedSince</code>	Updated since date. Requires the date to be in the format YYYY-mm-dd, for example 2017-01-01	string	form	Yes
<code>description</code>	A description that is posted as a comment for archiving.	string	form	Yes

Example response

Successful Response:

HTTP/1.1 200 OK

```
{
  "archivedReviews": [
    {
      "id": 836,
      "author": "swarm",
      "changes": [789],
      "commits": [],
      "commitStatus": [],
      "created": 1461164339,
      "deployDetails": [],
      "deployStatus": null,
      "description": "Review description\n",
      "groups": [],
      "participants": {
        "swarm": []
      },
      "pending": false,
```

```
    "projects": [],
    "state": "archived",
    "stateLabel": "Archived",
    "testDetails": [],
    "testStatus": null,
    "type": "default",
    "updated": 1478191607
  }
],
"failedReviews": []
}
```

Example usage

Archiving reviews inactive since 2016/06/30

To archive reviews not updated since 2016/06/30 inclusive:

```
curl -u "username:password" \
  -X POST \
  -d "notUpdatedSince=2016-06-30" \
  "https://my-swarm-host/api/v9/reviews/archive/"
```

Swarm responds with the list of archived reviews and failed reviews if there are any:

```
HTTP/1.1 200 OK

{
  "archivedReviews": [
    {
      "id": 911,
      "author": "swarm",
      "changes": [601],
      "commits": [],
      "commitStatus": [],
      "created": 1461164344,
      "deployDetails": [],
      "deployStatus": null,
      "description": "Touch up references on html pages.\n",

```

```

    "groups": [],
    "participants": {
      "swarm": []
    },
    "pending": false,
    "projects": [],
    "state": "archived",
    "stateLabel": "Archived",
    "testDetails": [],
    "testStatus": null,
    "type": "default",
    "updated": 1478191605
  },
  {
    "id": 908,
    "author": "earl",
    "changes": [605],
    "commits": [],
    "commitStatus": [],
    "created": 1461947794,
    "deployDetails": [],
    "deployStatus": null,
    "description": "Remove (attempted) installation of now deleted
man pages.\n",
    "groups": [],
    "participants": {
      "swarm": []
    },
    "pending": false,
    "projects": [],
    "state": "archived",
    "stateLabel": "Archived",
    "testDetails": [],
    "testStatus": null,
    "type": "default",

```

```
    "updated": 1478191605
  },
],
"failedReviews": [
  {
  }
]
}
```

If no reviews are archived, Swarm responds with an empty reviews list:

```
HTTP/1.1 200 OK

{
  "archivedReviews": [],
  "failedReviews": []
}
```

Add Change to Review

Summary

Add Change to Review

POST /api/v9/reviews/{id}/changes/

Description

Links the given change to the review and schedules an update.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required	Default Value
id	Review ID	integer	path	Yes	
change	Change ID	integer	form	Yes	
mode	The mode of operation, currently 'replace' or 'append'	string	form	No	replace

Example response

Successful Response contains Review Entity:

```
HTTP/1.1 200 OK
```

```
{
  "review": {
    "id": 12206,
    "author": "bruno",
    "changes": [10667, 12000],
    "commits": [10667, 12000],
    "commitStatus": [],
    "created": 1399325913,
    "deployDetails": [],
    "deployStatus": null,
    "description": "Adding .jar that should have been included in
r10145\n",
    "participants": {
      "bruno": []
    },
    "pending": false,
    "projects": [],
    "state": "archived",
    "stateLabel": "Archived",
```

```
    "testDetails": [],
    "testStatus": null,
    "type": "default",
    "updated": 1399325913
  }
}
```

Example usage

Adding a change to a review

You may want to update a review from a shelved or committed change that is different from the initiating change. This is done by adding a change to the review.

To replace a changelist:

```
curl -u "username:password" \
  -X POST \
  -d "change=124" \
  "https://my-swarm-host/api/v9/reviews/123/changes/"
```

To append a changelist:

```
curl -u "username:password" -X POST -d "change=124" -d "mode=append"
https://my-swarm-host/api/v9/reviews/123/changes/"
```

Swarm responds with the updated review entity:

```
HTTP/1.1 200 OK

{
  "review": {
    "id": 123,
    "author": "bruno",
    "changes": [122, 124],
    "commits": [],
    "commitStatus": [],
    "created": 1399325913,
    "deployDetails": [],
    "deployStatus": null,
    "description": "Adding .jar that should have been included in
r110\n",
```

```
{
  "groups": [],
  "participants": {
    "bruno": []
  },
  "pending": true,
  "projects": [],
  "state": "needsReview",
  "stateLabel": "Needs Review",
  "testDetails": [],
  "testStatus": null,
  "type": "default",
  "updated": 1399325913,
  "versions": [
    {
      "difference": 1,
      "stream": null,
      "change": 124,
      "user": "bruno",
      "time": 1399330003,
      "pending": true,
      "archiveChange": 124
    }
  ]
}
```

Clean up a review

Summary

Clean up a review (v6+)

POST /api/v9/reviews/{id}/cleanup

Description

Clean up a review for the given id.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
reopen	Expected to be a boolean (defaulting to false). If true then an attempt will be made to reopen files into a default changelist	boolean	form	No

Example response**Successful Response:**

```
HTTP/1.1 200 OK

{
  "complete": [
    {
      "1": ["2"]
    }
  ],
  "incomplete": []
}
```

Example usage**Cleaning up a review with id 1**

Cleanup review number 1, reopening any files into the default changelist.

```
curl -u "username:password" \
  -X POST \
  -d "reopen=true" \
  "https://my-swarm-host/api/v9/reviews/1/cleanup"
```

Swarm responds with the review and the changelists cleaned. Depending on the completion they will be either detailed in 'complete' or 'incomplete'. Incomplete changelists will have messages indicating why it was not possible to complete:

```
HTTP/1.1 200 OK

{
  "complete": [
    {
      "1": ["2"]
    }
  ],
  "incomplete": []
}
```

Obliterate a review

Summary

Obliterate a review

POST /api/v9/reviews/{id}/obliterate

Description

Obliterate a review for the given id.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Example response

Successful Response:

```
HTTP/1.1 200 OK

{
  "isValid": true,
  "message": "The review with id [1] has been obliterated.",
  "code": 200
}
```

Example usage

Obliterating a review with id 1

Obliterating review number 1.

```
curl -u "username:password" \
  -X POST \
  "https://my-swarm-host/api/v9/reviews/1/obliterate"
```

Swarm responds with a message informing you that is has successfully Obliterated the review:

```
HTTP/1.1 200 OK

{
  "isValid": true,
  "message": "The review with id [1] has been obliterated.",
  "code": 200
}
```

Set vote for the authenticated user to up, down, or cleared

Summary

Set the vote for the authenticated user to be up, down or cleared

POST /api/v9/reviews/{id}/vote/

Description

Set the vote for the authenticated user to be up, down or clear.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
version	Expected to be a valid review version to vote on if supplied, ignored if the version does not exist and the vote will apply to the latest version	string	form	No
vote	Expected to be a valid vote. Valid votes are 'up', 'down' and 'clear'	string	form	Yes

Example response

Successful Response:

```
HTTP/1.1 200 OK
```

```
{
  "isValid": "true",
  "messages": ["User username set vote to up on review 1"]
}
```

Example usage

Voting up for review with id 1

```
curl -u "username:password" \
  -X POST \
  -d "vote[value]=up" -d "vote[version]=1" \
  "https://my-swarm-host/api/v9/reviews/1/vote/"
```

Swarm responds with

```
HTTP/1.1 200 OK

{
  "isValid": "true",
  "messages": ["User username set vote to up on review 1"]
}
```

Transition the Review State

Summary

Transition the Review State (v2+)

PATCH /api/v9/reviews/{id}/state/

Description

Transition the review to a new state. When transitioning to approved, you can optionally commit the review. (v2+).

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Review ID	integer	path	Yes
<code>state</code>	Review State. Valid options: needsReview, needsRevision, approved, archived, rejected	string	form	Yes
<code>description</code>	An optional description that is posted as a comment for non-commit transitions. Commits that do not include a description default to using the Review description in the resulting change description.	string	form	No
<code>commit</code>	Set this flag to true and provide a state of approved in order to trigger the Approve and Commit action in Swarm.	boolean	form	No
<code>wait</code>	Instruct Swarm to wait for a commit to finish before returning.	boolean	form	No

Parameter	Description	Type	Parameter Type	Required
<code>jobs[]</code>	When performing an 'Approve and Commit', one or more jobs can be attached to the review as part of the commit process.	stringArray	form	No
<code>fixStatus</code>	Provide a fix status for the attached job (s) when performing an 'Approve and Commit'. Possible status values vary by job specification, but often include: open, suspended, closed, review, fixed.	string	form	No

Example responses

Successful Response contains Review Entity:

```
HTTP/1.1 200 OK
```

```
{
  "review": {
    "id": 12207,
    "author": "bruno",
    "changes": [10667, 12000],
    "commits": [],
    "commitStatus": [],
    "created": 1399325913,
    "deployDetails": [],
    "deployStatus": null,
```

```

    "description": "Adding .jar that should have been included in
r10145\n",
    "participants": {
      "bruno": []
    },
    "pending": false,
    "projects": [],
    "state": "needsRevision",
    "stateLabel": "Needs Revision",
    "testDetails": [],
    "testStatus": null,
    "type": "default",
    "updated": 1399325913
  },
  "transitions": {
    "needsReview": "Needs Review",
    "approved": "Approve",
    "rejected": "Reject",
    "archived": "Archive"
  }
}

```

Successful Commit contains Review and Commit Entities:

```

HTTP/1.1 200 OK

{
  "review": {
    "id": 12208,
    "author": "bruno",
    "changes": [10667, 12000, 12006],
    "commits": [12006],
    "commitStatus": {
      "start": 1399326910,
      "change": 12006,
      "status": "Committed",

```



```
    "committer": "bruno",
    "end": 1399326911
  },
  "created": 1399325900,
  "deployDetails": [],
  "deployStatus": null,
  "description": "Adding .jar that should have been included in
r10145\n",
  "participants": {
    "bruno": []
  },
  "pending": false,
  "projects": [],
  "state": "needsRevision",
  "stateLabel": "Needs Revision",
  "testDetails": [],
  "testStatus": null,
  "type": "default",
  "updated": 1399325905
},
"transitions": {
  "needsReview": "Needs Review",
  "needsRevision": "Needs Revision",
  "rejected": "Reject",
  "archived": "Archive"
},
"commit": 12006
}
```

Example usage

Committing a review

To commit a review:

```
curl -u "username:password" \
-X PATCH \
```

```
-d "state=approved" -d "commit=1" \
"https://my-swarm-host/api/v9/reviews/123/state/"
```

Swarm responds with the updated review entity, as well as a list of possible transitions for the review:

HTTP/1.1 200 OK

```
{
  "review": {
    "id": 123,
    "author": "bruno",
    "changes": [122, 124],
    "commits": [124],
    "commitStatus": {
      "start": 1399326910,
      "change": 124,
      "status": "Committed",
      "committer": "bruno",
      "end": 1399326911
    },
    "created": 1399325913,
    "deployDetails": [],
    "deployStatus": null,
    "description": "Adding .jar that should have been included in
r110\n",
    "groups": [],
    "participants": {
      "bruno": []
    },
    "pending": false,
    "projects": [],
    "state": "approved",
    "stateLabel": "Approved",
    "testDetails": [],
    "testStatus": null,
    "type": "default",
    "updated": 1399325913,
```

```
    "versions": [],
  },
  "transitions": {
    "needsReview": "Needs Review",
    "approved": "Approve",
    "rejected": "Reject",
    "archived": "Archive"
  }
}
```

Update Review Description

Summary

Update Review Description

PATCH /api/v9/reviews/{review_id}

Description

Update the description field of a review.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>review_id</code>	Review ID	integer	path	Yes

Parameter	Description	Type	Parameter Type	Required
author	The new author for the specified review. (At least one of Author or Description are required.)	string	form	No
description	The new description for the specified review. (At least one of Description or Author are required.)	string	form	No
_method	Method Override. If your client cannot submit HTTP PATCH, use an HTTP POST with the parameter ?_method=PATCH to override.	string	query	No

Example responses

Successful Response:

HTTP/1.1 200 OK

```
{
  "review": {
    "id": 12306,
    "author": "swarm",
    "changes": [12205],
    "comments": 0,
    "commits": [],
    "commitStatus": [],
    "created": 1402507043,
    "deployDetails": [],
    "deployStatus": null,
    "description": "Updated Review Description\n",
    "participants": {
```

```
    "swarm": []
  },
  "pending": true,
  "projects": [],
  "state": "needsReview",
  "stateLabel": "Needs Review",
  "testDetails": [],
  "testStatus": null,
  "type": "default",
  "updated": 1402518492
},
"transitions": {
  "needsRevision": "Needs Revision",
  "approved": "Approve",
  "rejected": "Reject",
  "archived": "Archive"
},
"canEditAuthor": true
}
```

Note

Swarm returns **null** for **totalCount** if no search filters were provided. **lastSeen** can often be used as an offset for pagination, by using the value in the **after** parameter of subsequent requests.

When no results are found, the **reviews** array is empty:

```
HTTP/1.1 200 OK

{
  "lastSeen": null,
  "reviews": [],
  "totalCount": 0
}
```

Servers : Swarm Servers API

Important

From Swarm 2022.2, Swarm no longer supports APIs older than v9.

Get a list of servers

Summary

Gets a list of servers

GET /api/v9/servers/

Description

Gets a list of servers

Example response

Successful Response:

```
HTTP/1.1 200 OK

{
  "servers": {
    "Main": {
      "port": "ssl:10.33.44.55:1666"
    },
    "Artifacts": {
      "port": "10.55.66.77:1666"
    }
  }
}
```

Example usage

Get a list of servers

Get a list of servers that Swarm is aware of

```
curl -u "username:password" "https://myswarm.url/api/v9/servers/"
```

If multiple p4d are configured it will give an output like

```
HTTP/1.1 200 OK

{
  "servers": {
    "Main": {
      "port": "ssl:10.33.44.55:1666"
    },
    "Artifacts": {
      "port": "10.55.66.77:1666"
    }
  }
}
```

If a single p4d is configured it will give an output like

```
HTTP/1.1 200 OK

{
  "servers": {
    "p4": {
      "port": "ssl:10.33.44.55:1666"
    }
  }
}
```

Users : Swarm Users

Important

From Swarm 2022.2, Swarm no longer supports APIs older than v9.

Get List of Users

Summary

Get List of Users

GET /api/v9/users/

Description

Returns a list of users in Swarm.

Parameters

Parameter	Description	Type	Parameter Type	Required
fields	An optional comma-separated list (or array) of fields to show for each user. Omitting this parameter or passing an empty value shows all fields. Be aware, the fields are case sensitive for users. You can use one of the following: User , Type , Email , Updated , Access , FullName , JobView , Password , AuthMethod , Reviews .	string	query	No
users	An optional comma-separated list (or array) of users to display. Omitting this parameter or passing an empty value shows all users.	string	query	No

Parameter	Description	Type	Parameter Type	Required
group	An optional to get users from a group. Cannot be used with users parameter	string	query	No
ignoreExcludeList	Determines if the list of users has the user_exclude_list filter applied or not. Add the parameter to ignore the user_exclude_list filter.	boolean	query	No

Example response

Successful Response:

HTTP/1.1 200 OK

```
{
  "User": "bruno",
  "Type": "standard",
  "Email": "bruno@perforce.com",
  "Updated": "2005/10/11 21:05:15",
  "Access": "2018/04/12 14:55:10",
  "FullName": "Bruno First",
  "JobView": null,
  "Password": null,
  "AuthMethod": "perforce",
  "Reviews": []
}
```

Example usage

Listing users

To list all users:

```
curl -u "username:password" "https://my-swarm-host/api/v9/users?fields=User,FullName&users=super,bruno"
```

Pagination is not currently supported by this endpoint. Swarm responds with a list of all users:

```
HTTP/1.1 200 OK
```

```
{
  {
    "User": "bruno",
    "FullName": "Bruno First"
  },
  {
    "User": "super",
    "FullName": "Super Second"
  }
}
```

Unfollow all Users and Projects

Summary

Unfollow all Users and Projects

POST /api/v9/users/{user}/unfollowall

Description

Admin and super users are permitted to execute unfollow all against any target user. Other users are only permitted to execute the call if they themselves are the target user

Example response

Successful Response:

```
HTTP/1.1 200 OK
```

```
{
  "isValid": true,
  "messages": "User {user} is no longer following any Projects or
Users."
}
```

Workflows : Controller for querying, creating and updating workflows

Important

From Swarm 2022.2, Swarm no longer supports APIs older than v9.

Get workflows

Summary

Gets workflows

GET /api/v9/workflows/

Description

Get all workflows

Parameters

Parameter	Description	Type	Parameter Type	Required
fields	An optional comma-separated list (or array) of fields to show for each workflow. Omitting this parameter or passing an empty value shows all fields.	string	query	No

Parameter	Description	Type	Parameter Type	Required
noCache	If provided and has a value of 'true' a query will always be performed and the cache of workflows is ignored. Otherwise the cache will be used if it exists.	boolean	query	No

Usage example

Get a list of workflows

```
curl -u "username:password" "https://my-swarm-host/api/v9/workflows"
```

JSON Response:

```
HTTP/1.1 200 OK
```

```
{
  "workflows": [
    {
      "on_submit": {
        "without_review": {
          "rule": "no_checking"
        },
        "with_review": {
          "rule": "no_checking"
        }
      },
      "name": "myWorkflow",
      "description": "A description",
      "shared": "true",
      "owners": [
        "user1",
        "user2"
      ],
      "end_rules": {
```

```
        "update":{
            "rule":"no_revision"
        },
        "auto_approve":{
            "rule": "never"
        },
        "counted_votes":{
            "rule": "anyone"
        },
        "group_exclusions":{
            "rule": [ ]
        },
        "user_exclusions":{
            "rule": [ ]
        },
        "id": "1"
    },
    {
        "on_submit":{
            "without_review":{
                "rule": "no_checking"
            },
            "with_review":{
                "rule": "no_checking"
            }
        },
        "name": "myWorkflow 2",
        "description": "Another description",
        "shared": "true",
        "owners": [
            "user3",
            "user4"
        ],
        "end_rules":{
```

```
        "update":{
            "rule":"no_revision"
        },
        "auto_approve":{
            "rule": "votes"
        },
        "counted_votes": {
            "rule": "members"
        },
        "group_exclusions":{
            "rule": [ ]
        },
        "user_exclusions":{
            "rule": [ ]
        },
        "id": "2"
    },
]
}
```

Get a workflow by id

Summary

Gets a workflow by id

GET /api/v9/workflows/{id}

Description

Gets a workflow by id

Parameters

Parameter	Description	Type	Parameter Type	Required
fields	An optional comma-separated list (or array) of fields to show for each workflow. Omitting this parameter or passing an empty value shows all fields.	string	query	No

Example usage

Get a workflows by id

```
curl -u "username:password" "https://my-swarm-host/api/v9/workflows/1"
```

JSON Response:

```
HTTP/1.1 200 OK
```

```
{
  "workflows": [
    {
      "on_submit": {
        "without_review": {
          "rule": "no_checking"
        },
        "with_review": {
          "rule": "no_checking"
        }
      },
      "name": "myWorkflow",
      "description": "A description",
      "shared": "true",
      "owners": [
        "user1",
        "user2"
      ],
      "end_rules": {
        "update": {
```

```
        "rule": "no_revision"
      },
    },
    "auto_approve": {
      "rule": "never"
    },
    "counted_votes": {
      "rule": "anyone"
    },
    "group_exclusions": [ ],
    "rule": [ ]
  },
  "user_exclusions": [ ],
  "rule": [ ]
},
"id": "1"
}
```

Create a workflow

Summary

Create a workflow

POST /api/v9/workflows/

Description

Create a new workflow.

Parameters

Parameter	Description	Type	Parameter Type	Required
name	The workflow name. Will be compared against other workflows and rejected if not unique	string	form	Yes
description	Description for the new workflow	string	form	No
shared	Whether this workflow is shared for other users that do not own it. Defaults to not shared	boolean	form	No
owners	A list owners for the workflow. Can be users or group names (prefixed with swarm-group-). Users and group names must exist or the workflow will be rejected	array (of strings)	form	No
on_submit	Data for rules when changes are submitted. Valid values for with_review are no_checking, approved, strict. Valid values for without review are no_checking, auto_create, reject	array	form	No
end_rules	Data for rules when changes are submitted. Valid values are no_checking, no_revision.	array	form	No

Parameter	Description	Type	Parameter Type	Required
auto_approve	Data for rules when changes are submitted. Valid values are votes, never.	array	form	No
counted_votes	Data for rules when counting votes up. Valid values are anyone, members.	array	form	No

Example usage

Create a workflow

```
curl -u "username:password" \
  -X POST \
  -d "on_submit[with_review][rule]=no_checking" \
  -d "on_submit[without_review][rule]=no_checking" \
  -d "name=myWorkflow" \
  -d "description=A description" \
  -d "shared=true" \
  -d "owners[]=Francois_Piccard" \
  -d "owners[]=Anna_Schmidt" \
  -d "end_rules[update][rule]=no_checking" \
  -d "auto_approve:[rule]never" \
  -d "counted_votes:[rule]members" \
  "https://my-swarm-host/api/v9/workflows"
```

JSON Response:

```
HTTP/1.1 200 OK

{
  "workflow":{
    "on_submit":{
      "with_review":{
        "rule":"no_checking"
```

```

        },
        "without_review":{
            "rule":"no_checking"
        }
    },
    "name":"myWorkflow",
    "description":"A description",
    "shared":true,
    "owners":[
        "Anna_Schmidt",
        "Francois_Piccard"
    ],
    "end_rules":{
        "update":{
            "rule":"no_checking"
        }
    },
    "auto_approve":{
        "rule":"never"
    },
    "counted_votes":{
        "rule":"members"
    },
    "group_exclusions":[ ],
        "rule": [ ]
    },
    "user_exclusions":[ ],
        "rule": [ ]
    },
    "id":1
}
}

```

Patch a workflow

Summary

Patch a workflow

PATCH /api/v9/workflows/{id}

Description

Patch a workflow.

Parameters

Parameter	Description	Type	Parameter Type	Required
id	The id of the workflow being patched	string	form	Yes
name	The workflow name. Will be compared against other workflows and rejected if not unique	string	form	No
description	Description for the new workflow	string	form	No
shared	Whether this workflow is shared for other users that do not own it. Defaults to not shared	boolean	form	No
owners	A list owners for the workflow. Can be users or group names (prefixed with swarm-group-). Users and group names must exist or the workflow will be rejected	array (of strings)	form	No

Parameter	Description	Type	Parameter Type	Required
on_submit	Data for rules when changes are submitted. Valid values for with_review are no_checking, approved, strict. Valid values for without review are no_checking, auto_create, reject	array	form	No
end_rules	Data for rules when changes are submitted. Valid values are no_checking, no_revision.	array	form	No
auto_approve	Data for rules when changes are submitted. Valid values are votes, never.	array	form	No
counted_votes	Data for rules when counting votes up. Valid values are anyone, members.	array	form	No

Example usage

Patch workflow name and description

```
curl -u "username:password" \
  -X PATCH \
  -d "name=myWorkflow" \
  -d "description=A description" \
  "https://my-swarm-host/api/v9/workflows/1"
```

JSON Response:

```
HTTP/1.1 200 OK

{
```

```
"workflow":{
  "on_submit":{
    "with_review":{
      "rule":"no_checking"
    },
    "without_review":{
      "rule":"no_checking"
    }
  },
  "name":"myWorkflow",
  "description":"A description",
  "shared":false,
  "owners":[
    "bruno"
  ],
  "end_rules":{
    "update":{
      "rule":"no_revision"
    }
  },
  "auto_approve":{
    "rule":"votes"
  },
  "counted_votes":{
    "rule":"members"
  },
  "group_exclusions":[ ],
  "rule": [ ]
},
"user_exclusions":[ ],
"rule": [ ]
},
"id":1
}
}
```

Patch on_submit with_review

```
curl -u "username:password" \  
  -X PATCH \  
  -d "on_submit[with_review][rule]=no_checking" \  
  "https://my-swarm-host/api/v9/workflows/1"
```

JSON Response:

```
HTTP/1.1 200 OK  
  
{  
  "workflow":{  
    "on_submit":{  
      "with_review":{  
        "rule":"no_checking"  
      },  
      "without_review":{  
        "rule":"no_checking"  
      }  
    },  
    "name":"test",  
    "description":"test line 2",  
    "shared":false,  
    "owners":[  
      "bruno"  
    ],  
    "end_rules":{  
      "update":{  
        "rule":"no_revision"  
      }  
    },  
    "auto_approve":{  
      "rule":"votes"  
    },  
    "counted_votes":{  
      "rule":"members"  
    },  
  },  
}
```

```

    "group_exclusions":[ ],
      "rule": [ ]
    },
    "user_exclusions":[ ],
      "rule": [ ]
    },
    "id":1
  }
}

```

Patch on_submit without_review

```

curl -u "username:password" \
  -X PATCH \
  -d "on_submit[without_review][rule]=no_checking" \
  "https://my-swarm-host/api/v9/workflows/1"

```

JSON Response:

```

HTTP/1.1 200 OK

{
  "workflow":{
    "on_submit":{
      "with_review":{
        "rule":"no_checking"
      },
      "without_review":{
        "rule":"no_checking"
      }
    },
    "name":"test",
    "description":"test line 2",
    "shared":false,
    "owners":[
      "bruno"
    ],
    "end_rules":{

```



```
        "update":{
            "rule":"no_revision"
        },
        "auto_approve":{
            "rule":"votes"
        },
        "counted_votes":{
            "rule":"members"
        },
        "group_exclusions":[ ],
        "rule": [ ]
    },
    "user_exclusions":[ ],
    "rule": [ ]
},
"id":1
}
}
```

Patch Auto_approve

```
curl -u "username:password" \
-X PATCH \
-d "auto_approve[rule]=never" \
"https://my-swarm-host/api/v9/workflows/1"
```

JSON Response:

```
HTTP/1.1 200 OK

{
    "workflow":{
        "on_submit":{
            "with_review":{
                "rule":"no_checking"
            },
            "without_review":{
```

```

        "rule":"no_checking"
      }
    },
    "name":"test",
    "description":"test line 2",
    "shared":false,
    "owners":[
      "bruno"
    ],
    "end_rules":{
      "update":{
        "rule":"no_revision"
      }
    },
    "auto_approve":{
      "rule":"never"
    },
    "counted_votes":{
      "rule":"members"
    },
    "group_exclusions":[ ],
    "rule": [ ]
  },
  "user_exclusions":[ ],
  "rule": [ ]
},
"id":1
}
}

```

Patch counted_votes

```

curl -u "username:password" \
  -X PATCH \
  -d "counted_votes[rule]=anyone" \
  "https://my-swarm-host/api/v9/workflows/1"

```

JSON Response:

HTTP/1.1 200 OK

```
{
  "workflow":{
    "on_submit":{
      "with_review":{
        "rule":"no_checking"
      },
      "without_review":{
        "rule":"no_checking"
      }
    },
    "name":"test",
    "description":"test line 2",
    "shared":false,
    "owners":[
      "bruno"
    ],
    "end_rules":{
      "update":{
        "rule":"no_revision"
      }
    },
    "auto_approve":{
      "rule":"never"
    },
    "counted_votes":{
      "rule":"anyone"
    },
    "group_exclusions":[ ],
    "rule": [ ]
  },
  "user_exclusions":[ ],
  "rule": [ ]
},
```

```
    "id":1
  }
}
```

Patch end_rules

```
curl -u "username:password" \
  -X PATCH \
  -d "end_rules[update][rule]=no_checking" \
  "https://my-swarm-host/api/v9/workflows/1"
```

JSON Response:

```
HTTP/1.1 200 OK

{
  "workflow":{
    "on_submit":{
      "with_review":{
        "rule":"no_checking"
      },
      "without_review":{
        "rule":"no_checking"
      }
    },
    "name":"test",
    "description":"test line 2",
    "shared":false,
    "owners":[
      "bruno"
    ],
    "end_rules":{
      "update":{
        "rule":"no_checking"
      }
    },
    "auto_approve":{
      "rule":"never"
    }
  }
}
```

```
    },
    "counted_votes":{
      "rule":"members"
    },
    "group_exclusions":[ ],
      "rule": [ ]
    },
    "user_exclusions":[ ],
      "rule": [ ]
    },
    "id":1
  }
}
```

Patch shared

```
curl -u "username:password" \
-X PATCH \
-d "shared=true" \
"https://my-swarm-host/api/v9/workflows/1"
```

JSON Response:

```
HTTP/1.1 200 OK

{
  "workflow":{
    "on_submit":{
      "with_review":{
        "rule":"no_checking"
      },
      "without_review":{
        "rule":"no_checking"
      }
    },
    "name":"test",
    "description":"test line 2",
    "shared":true,
```

```

    "owners":[
        "bruno"
    ],
    "end_rules":{
        "update":{
            "rule":"no_checking"
        }
    },
    "auto_approve":{
        "rule":"never"
    },
    "counted_votes":{
        "rule":"anyone"
    },
    "group_exclusions":[ ],
        "rule": [ ]
    },
    "user_exclusions":[ ],
        "rule": [ ]
    },
    "id":1
}
}

```

Patch Owners

```

curl -u "username:password" \
  -X PATCH \
  -d "owners[]=Francois_Piccard" \
  -d "owners[]=Anna_Schmidt" \
  -d "owners[]=bruno" \
  "https://my-swarm-host/api/v9/workflows/1"

```

JSON Response:

```

HTTP/1.1 200 OK

{

```

```
"workflow":{
  "on_submit":{
    "with_review":{
      "rule":"no_checking"
    },
    "without_review":{
      "rule":"no_checking"
    }
  },
  "name":"test",
  "description":"test line 2",
  "shared":true,
  "owners":[
    "Anna_Schmidt",
    "Francois_Piccard",
    "bruno"
  ],
  "end_rules":{
    "update":{
      "rule":"no_checking"
    }
  },
  "auto_approve":{
    "rule":"never"
  },
  "counted_votes":{
    "rule":"anyone"
  },
  "group_exclusions":[ ],
    "rule": [ ]
  },
  "user_exclusions":[ ],
    "rule": [ ]
  },
  "id":1
```

```
}  
}
```

Update a workflow

Summary

Update a workflow

PUT /api/v9/workflows/{id}

Description

Update a workflow. All values should be provided in the request. If not provided any missing values are reverted to default.

Parameters

Parameter	Description	Type	Parameter Type	Required
id	The id of the workflow being patched	string	form	Yes
name	The workflow name. Will be compared against other workflows and rejected if not unique	string	form	Yes
owners	A list owners for the workflow. Can be users or group names (prefixed with swarm-group-). Users and group names must exist or the workflow will be rejected	array (of strings)	form	No
description	Description for the new workflow	string	form	No

Parameter	Description	Type	Parameter Type	Required
shared	Whether this workflow is shared for other users that do not own it. Defaults to not shared	boolean	form	No
on_submit	Data for rules when changes are submitted. Valid values for with_review are no_checking, approved, strict. Valid values for without review are no_checking, auto_create, reject	array	form	No
end_rules	Data for rules when changes are submitted. Valid values are no_checking, no_revision.	array	form	No
auto_approve	Data for rules when changes are submitted. Valid values are votes, never.	array	form	No
counted_votes	Data for rules when counting votes up. Valid values are anyone, members.	array	form	No

Example usage

Update a workflow

```
curl -u "username:password" \
  -X PUT \
  -d "on_submit[with_review][rule]=no_checking" \
  -d "on_submit[without_review][rule]=no_checking" \
  -d "name=myWorkflow" \
```

```
-d "description=A description" \  
-d "shared=true" \  
-d "owners[]=Francois_Piccard" \  
-d "owners[]=Anna_Schmidt" \  
"https://my-swarm-host/api/v9/workflows/1"
```

JSON Response:

HTTP/1.1 200 OK

```
{  
  "workflow":{  
    "on_submit":{  
      "with_review":{  
        "rule":"no_checking"  
      },  
      "without_review":{  
        "rule":"no_checking"  
      }  
    },  
    "name":"myWorkflow",  
    "description":"A description",  
    "shared":true,  
    "owners":[  
      "Anna_Schmidt",  
      "Francois_Piccard"  
    ],  
    "end_rules":{  
      "update":{  
        "rule":"no_checking"  
      }  
    },  
    "auto_approve":{  
      "rule":"never"  
    },  
    "counted_votes":{  
      "rule":"anyone"  
    }  
  }  
}
```

```
    },
    "group_exclusions": [ ],
    "rule": [ ]
  },
  "user_exclusions": [ ],
  "rule": [ ]
},
"id":1
}
```

Delete a workflow

Summary

Delete a workflow

DELETE /api/v9/workflows/{id}

Description

Delete a workflow for the provided id. This call must be authenticated and the user must have permission to edit the workflow. If the workflow is in use it cannot be deleted and an error message will be returned.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<u>id</u>	The id of the workflow being deleted	string	form	Yes

Example usage

Delete a workflow not in use

```
curl -u "username:password" -X DELETE "https://my-swarm-host/api/v9/workflows/1"
```

JSON Response

```
HTTP/1.1 200 OK

{
  "isValid": true,
  "messages": [
    "Workflow [1] was deleted"
  ]
}
```

Delete a workflow that is in use on public projects

```
curl -u "username:password" -X DELETE "https://my-swarm-host/api/v9/workflows/1"
```

JSON Response

```
HTTP/1.1 200 OK

{
  "isValid": false,
  "messages": [
    "Cannot delete workflow [1], it is in use on project [project1, project2]"
  ]
}
```

Delete a workflow that is in use on a public and some private projects

```
curl -u "username:password" -X DELETE "https://my-swarm-host/api/v9/workflows/1"
```

JSON Response

```
HTTP/1.1 200 OK

{
  "isValid": false,
  "messages": [
    "Cannot delete workflow [1], it is in use on project [project1] and others"
  ]
}
```

Delete a workflow that is in use on only private projects

```
curl -u "username:password" -X DELETE "https://my-swarm-host/api/v9/workflows/1"
```

JSON Response

```
HTTP/1.1 200 OK

{
  "isValid": false,
  "messages": [
    "Cannot delete workflow [1], it is in use"
  ]
}
```

API (v9) examples

This section contains some API (v9) examples.

Extended API example

This section contains an extended API example, involving multiple API calls to answer a more complicated kind of question than any single API endpoint can provide: which reviews does a specific *userid* need to attend to?

The code

```
<?php
/**
```

```
* vim:set ai si et ts=4 sw=4 syntax=php:
*
* reviews.php
*
* Queries the Swarm API and reports which reviews a specified user
* needs to attend to.
*
* Required attention is determined by the following criteria:
* - the user is a participant in a review
* - and the user has not voted on the review
* - and the user has not commented on the review
* - or the user's comment on the review is a
*   task that has been addressed and needs verification
*/

if (ini_set('track_errors', 1) === false) {
    echo "Warning: unable to track errors.\n";
}

# process command-line arguments
$options = getopt(
    'hs:r:v',
    array('help', 'swarm:', 'reviewer', 'verbose')
);

$swarm = '';
if (isset($options['s'])) {
    $swarm = $options['s'];
}
if (isset($options['swarm'])) {
    $swarm = $options['swarm'];
}
if (!$swarm) {
    usage('Swarm API URL not provided.');
```

```
$reviewer = '';
if (isset($options['r'])) {
    $reviewer = $options['r'];
}
if (isset($options['reviewer'])) {
    $reviewer = $options['reviewer'];
}
if (!$reviewer) {
    usage('Swarm reviewer not provided.');
```



```
$verbose = false;
if (isset($options['v']) || isset($options['verbose'])) {
    $verbose = true;
}

if (isset($options['h']) || isset($options['help'])) {
    usage();
}

function usage($message = null)
{
    if ($message) {
        echo "$message\n\n";
    }

    $script = basename(__FILE__);
    echo <<<EOU
$script: -s <Swarm URL> -u <API userid> -p <API user's password> \
    -r <reviewer userid to report on> -h

-s|--swarm      Swarm's URL (e.g. https://user@password:myswarm.url/)
-r|--reviewer   The reviewer to report on.
-h|--help       This help text.
```

```
-v|--verbose    Verbose output.
```

This script queries the Swarm API and reports on reviews that the specified user needs to attend to.

Note: If your Helix Core Server (p4d) has security level 3 set, you cannot use a password to authenticate; you must acquire a host-unlocked

ticket from p4d, and use the ticket in place of a password when communicating with the Swarm API connected to p4d.

```
EOU;
```

```
    exit;
```

```
}
```

```
function msg($message)
```

```
{
```

```
    global $verbose;
```

```
    if ($verbose) {
```

```
        echo $message;
```

```
    }
```

```
}
```

```
function call_api($url, $params)
```

```
{
```

```
    global $php_errormsg;
```

```
    $query    = http_build_query($params);
```

```
    $request  = $url . '?' . $query;
```

```
    $response = @file_get_contents($request);
```

```
    if ($php_errormsg) {
```

```
        echo "Unable to call api: $php_errormsg\n";
```

```
        exit;
```

```
    }
```



```
$json = @json_decode($response, true);
if ($php_errormsg) {
    echo "Unable to decode api response: $php_errormsg\n";
    exit;
}

return $json;
}

# remove trailing / from URL, if it exists
$swarm = rtrim(trim($swarm), '/');

# fetch the list of reviews
$reviews = call_api(
    "$swarm/api/v9/reviews",
    array(
        'hasReviewers' => 1, # only reviews with participants
        'participants' => array($reviewer), # only review for this
reviewer
        'max'          => 9, # get plenty of reviews, if available
        'fields'       => array('id', 'description', 'commits'), # get
these fields
    )
);

$report = array();
foreach ($reviews['reviews'] as $review) {
    if (is_null($review)) {
        continue;
    }

    $flag = false;
    msg('Review: ' . $review['id'] . ' ');
```

```

    # if the review is already committed, it likely does not need
attention
    if (array_key_exists('commits', $review)
        && count($review['commits']))
    ) {
        msg("is committed, skipping...\n");
        continue;
    }

    # if the review has a vote from the reviewer, they are already
aware
    if (array_key_exists('participants', $review)
        && array_key_exists('vote', $review['participants']
[$reviewer]))
    ) {
        msg("has vote from reviewer, skipping...\n");
        continue;
    }

    # if there are no open comments on the review, the reviewer's
# attention is required
    if (array_key_exists('comments', $review)
        && $review['comments'][0] == 0
    ) {
        msg("has no open comments, skipping...\n");
        continue;
    }

    # fetch the comments for this review
    $comments = call_api(
        "$swarm/api/v9/comments",
        array(
            'topic' => 'reviews/' . $review['id'], # comments for this
review
            'max'    => 9, # get plenty of comments, if available

```

```
    )
);

foreach ($comments['comments'] as $comment) {
    msg("\n Comment: " . $comment['id'] . ' ');

    // skip over comments from other reviewers
    if (array_key_exists('user', $comment) && $reviewer !=
$comment['user']) {
        msg("is by another user, carry on...\n");
        continue;
    }

    # skip archived comments
    if (array_key_exists('flags', $comment)
        && count($comment['flags']) > 0
        && $comment['flags'][0] == 'closed'
    ) {
        msg("is archived, carry on...\n");
        continue;
    }

    # skip marked tasks
    if (array_key_exists('taskState', $comment)
        && ($comment['taskState'] == 'comment'
            || $comment['taskState'] == 'verified'
            || $comment['taskState'] == 'open'
        )
    ) {
        msg("reviewer's comment needs attention, carry on...\n");
        continue;
    }

    // anything else means that the reviewer's comment needs
attention
```

```

        // by the reviewer
        $flag = true;
        msg("needs attention!\n");
        break;
    }

    // evaluation is complete. Does this review need attention?
    if ($flag) {
        $report[] = $review;
    }
}

if (count($report)) {
    echo "User '$reviewer' needs to attend to these reviews:\n";
    foreach ($report as $review) {
        $description = trim($review['description']);
        if (strlen($description) > 60) {
            $description = substr($description, 0, 60) . ' ...';
        }
        echo $review['id'] . ": $description\n";
    }
} else {
    echo "User '$reviewer' has no reviews to attend to.\n";
}

```

Executing the example

The example is written in PHP. To use it copy and paste it into a file called **reviews.php**. Then, execute it like this:

```
$ php reviews.php -s https://myswarm.host:port/ -r bob
```

Replace **http://myswarm.host/** with the URL to your Swarm installation. Replace **bob** with the userid you'd like to report on.

To authenticate, insert **username:password@** before the hostname. If your Helix Core Server security counter is set to **3** or higher, you need to acquire a ticket and use the ticket in place of the password (see [Authentication](#) for details). If your Swarm is installed on a **custom port**, or is installed in a **sub-folder**, include those elements in the URL as well. For example:

```
$ php reviews.php -s
https://me:F0FC33068BA244B1BBD8196CC9166F34@my.host:8080/swarm/ -r bob
```

If you do not specify the URL correctly, you might see an error like:

```
Unable to call api: file_get_contents
(http://...@my.host:8080/swarm/api/v9
/reviews?hasReviewers=1&participants%5B0%5D=bob&max=9&fields%5B0%5D
=id&fields%5B1%5D=description&fields%5B2%5D=commits): failed to
open stream: HTTP request failed! HTTP/1.1 404 Not Found
```

If there are no errors, and the specified userid does have reviews to attend to, the output might look like:

```
1234: Added grapple-grommit support to woozlewobble class. @bob sh
...
```

1234 is the id of a review that **bob** should attend to, followed by the first 60 characters of the review's description.

Pending review cleanup API example

This section contains an example script that cleans up pending changelists which are no longer needed. See the [review cleanup](#) options for how this can be done automatically when a review is committed.

For pending changelists which were present before this option was available, or for reviews which have been contributed to by multiple authors and so require super user access to tidy up, there is an API which allows the super user to bulk remove such changelists.

The script demonstrates how this API could be used. It isn't meant as a complete solution, just a guide to demonstrate what is possible.

The code

```
<?php
/**
 * Perforce Swarm
 *
 * @copyright 2017 Perforce Software. All rights reserved.
 * @license   Please see LICENSE.txt in top-level folder of this
distribution.
 * @version   <release>/<patch>
 */
/**
 * This example script can be used to clean up (delete) Perforce
Server pending changelists automatically when run
```

```

* as a super user. It is able to query for reviews based on
parameters to establish which changelists are eligible
* for clean up. In this way, it can be tailored to run against
reviews of a user's choice.

* Requirements for this script:
*   - MUST be a super user
*   - MUST populate the parameters below
*   - MUST be using Swarm 2017.1 or later
*
* Usage of script
*   php superUserReviewCleanUp.php max=10 notUpdatedSince=2017-04-01
state=approved
*   php superUserReviewCleanUp.php max=10 author=bruno state=approved
*
* Each of the parameters that you can use are documented at the
following URL:
*
https://www.perforce.com/perforce/doc.current/manuals/swarm/Content/Swarm/swarm-apidoc\_endpoint\_reviews.html
*
*
* This returns a JSON object that contains four main objects.
*
* {
*   "error": "",
*   "help": "",
*   "results": "[]",
*   "search_criteria": {
*     "fields": "id",
*     "max": "10",
*     "notUpdatedSince": "2017-04-01"
*   }
* }
*

```

```
* Referencing the example above:
*
* The error section will contain any errors that have been
encountered trying to execute the script.
*
* The help section will populated if you have run the command php
superUserReviewCleanUp.php help
*
* The search_criteria section indicates which parameters have been
used to fetch the reviews list.
*
* The results section returns a JSON object of each of the reviews it
has processed, which may
* include actions that were incomplete.
*
* Below is an example of results being processed:
* "814": {
*   "complete": [],
*   "incomplete": {
*     "814": {
*       "813": [
*         "0": "Command failed: No shelved files in changelist
to delete.",
*       ]
*     }
*   }
* },
* "818": {
*   "complete": [],
*   "incomplete": []
* }
* "820": {
*   "complete": [],
*   "incomplete": {
*     "820": {
```

```

*           "821": [
*               "0": "Command failed: No shelved files in changelist
to delete.",
*               "1": "Command failed: Usage: fix [ -d ] [ -s status ]
-c changelist# jobName ...\nMissing/wrong
*                   number of arguments."
*           ]
*       }
*   }
* },
*
* Some reviews may have no actions or incomplete actions. Incomplete
actions indicate that additional work is required
* and the review could not be entirely cleaned up. In the example
above, the first message indicates a changelist was
* not found. This could be because the end user has already deleted
it.
*
* An error with the fix command can indicate that the pending
changelist doesn't have any jobs linked or that the jobs
* have already been removed.
*
* NOTES:
* To make the output print nicely, you can use the python command
like this:
*
* php superUserReviewCleanUp.php max=10 notUpdatedSince=2018-04-01
state=approved | python -m json.tool
*
*/

/* ***** */
/* These values MUST be set before running the script. */
/* ***** */

```



```
// URL to access swarm. Must not include trailing slash.
$swarmURL = 'http://my.swarm.com';
// Username of super user
$username = '';
// Ticket for user, can be created by running "p4 -u $username login -
pa"
$ticket = '';

/* ***** */

// If the super user wants to reopen the files of the end user.
$reopen = true;
// Prebuild the the return message helper array.
$help = array( "help" => "", "error" => "", "results" => "", "search_
criteria" => "");
// @codingStandardsIgnoreEnd
/**
 * function that make the GET or POST requests
 *
 * @param $url          Url in which we want to make our request to
 * @param bool $fetch   Set to true for a GET request, otherwise will
do a POST.
 * @param array $args   These are the parameter that we pass the the
GET or POST request to filter the reviews
 * @return JSON        We return a JSON object back to be worked on.
 */
function request($url, $fetch = false, $args = array())
{
    // Fetch the settings to allow this function to access them.
    global $username, $ticket, $reopen, $help;
    $curl = '';

    // If GET request fetch should be true and args shouldn't be empty
```

```

    if ($fetch === true) {
        // If is args is empty just give the url, otherwise build a
http query with args elements
        $curl = empty($args) ? curl_init($url) : curl_init
($url."?".http_build_query($args));
    } else {
        // Assume fetch is false and build a POST request.
        $curl = curl_init($url."/". $args['id'].' /cleanup');
        curl_setopt($curl, CURLOPT_POSTFIELDS, http_build_query(array
('reopened'=>$reopen)));
        curl_setopt($curl, CURLOPT_POST, count($reopen));
    }
    curl_setopt($curl, CURLOPT_USERPWD, "$username:$ticket");
    curl_setopt($curl, CURLOPT_RETURNTRANSFER, true);
    curl_setopt($curl, CURLOPT_HTTPAUTH, CURLAUTH_BASIC);

    $result = curl_exec($curl);

    // Catch the error code in case Ticket expired or url doesn't
return.
    $statusCode = curl_getinfo($curl, CURLINFO_HTTP_CODE);
    if ($statusCode != 200) {
        $help['error'] = array("HTTP code" => $statusCode);
    }

    curl_close($curl);
    return json_decode($result, true);
}

/**
 * Fetch a list of Reviews based on parameters
 *
 * @param $elements Each of the args passed from command line are
treated as elements
 * @return JSON      We return a JSON object back to be worked on.

```

```
*/
function fetchReviews($elements)
{
    global $swarmURL, $help;
    $parameters = array();

    if ($elements !== null) {
        // Loop through each of the args we want to fetch reviews based
on
        foreach ($elements as $key => $value) {
            // Break each of the element into key and value to allow
use to build a array to pass to url
            $brokenDown = explode("=", $value);
            if ($key !== 0 && sizeof($brokenDown) === 2) {
                $parameters[$brokenDown[0]] = $brokenDown[1];
            }
        }
        // We only require the field id to limit the amount of data.
        $parameters['fields'] = 'id';
        // Helpful for debugging which parameters we have passed in
the queue.
        $help['search_criteria'] = $parameters;
    }
    // Now make the request to the Swarm server with your field
options.
    $result = request(
        "$swarmURL/api/v9/reviews",
        true,
        $parameters
    );

    // Return the JSON object back to be worked on.
    return $result;
}
```

```

/**
 * Loop though each of the Reviews passed in and run clean up for
 * them.
 *
 * @param $reviews  JSON object of all the reviews we want to run
 * cleanup on
 * $reviews => array(
 *     'reviews' => array(
 *         0 => array (
 *             'id' => 134,
 *         ),
 *         1 => array (
 *             'id' => 136,
 *         ),
 *         2 => array (
 *             'id' => 143,
 *         ),
 *         3 => array (
 *             'id' => 158,
 *         )
 *     )
 * )
 *
 * @return $array  return the array of work that has been carried
 * out.
 */
function runCleanUp($reviews)
{
    global $swarmURL;
    $results = array();
    // Check if there is an reviews element of the array
    if (isset($reviews['reviews'])) {
        foreach ($reviews['reviews'] as $review) {
            // Now make the request to the Swarm for each review.
            $results[$review['id']] = request(

```

```
        "$swarmURL/api/v9/reviews",
        false,
        $review
    );
}
}
return $results;
}

/**
 * The help function in case a user doesn't set the basic settings.
 *
 * @param $help    Pass the help array from main script to append the
helpful message.
 * @return $array  Return the array that will be presented to the
users.
 */
function helpMessage($help)
{
    $help["help"]      = array( "1" => "", "2" => "", "3" => "" );
    $help["help"]["1"] = "Please ensure you have set the Username,
Ticket and Swarm URL before using this script";
    $help["help"]["2"] = "Running the script can be done by using any
of the standard Swarm API fields for reviews";
    $help["help"]["3"] = "Visit
https://www.perforce.com/perforce/doc.current/manuals/swarm/index.htm
l";
    return $help;
}

// check the first argument is not help.
$helpSet = isset($argv[1]) && $argv[1] == "help" ? "help" : null;

// Check if the user has given help as a command to this script.
if (isset($argv[1]) && $helpSet == "help") {
```

```

    // Set the $help array with the help message.
    $help = helpMessage($help);
}

// Check if the basic user ticket and swarmurl are set.
if (!empty($username) && !empty($ticket) && !empty($swarmURL) &&
$helpSet != "help") {
    try {
        $help["results"] = runCleanup(fetchReviews($argv));
    } catch (Exception $e) {
        $help = helpMessage($help);
    }
} else {
    $message = "Please ensure you have set the below before using
this script";
    $dataArray = array("message" => $message, "parameter" => "");

    $missingParameter = array();

    // Now check if the basic settings are empty and show the end
user.
    empty($username) ? $missingParameter[] = "Username:''";
    empty($ticket) ? $missingParameter[] = "Ticket:''";
    empty($swarmURL) ? $missingParameter[] = "SwarmURL:''";

    $dataArray['parameter'] = $missingParameter;

    $help["error"] = $dataArray ;
}

// Output the end result of the what the script does.
echo json_encode($help, JSON_FORCE_OBJECT);
// @codingStandardsIgnoreEnd

```

Executing the script

The example is written in PHP, and demonstrates how to make use of the APIs which remove unneeded pending changelists. It **must** be run as a super user.

For a full set of instructions on how to use the example script, see the comments in the script itself.

Abbreviated instructions:

1. Set the value of the `$swarmURL`, `$username` and `$ticket` variables.
2. Run the script by using a command similar to the following:

```
$ php pendingReviewCleanUp.php max=10 author=bruno state=approved
```

API endpoints (v10)

This section includes coverage for the endpoints provided by the API (v10).

Activity: Swarm activity

Get a list of all activities

Summary

Get a list of all activities.

GET /api/v10/activity

Description

Get a list of all activities the user has permission to view.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
followed	<p>Set the followed parameter to true to limit results to activities followed by the user making the request.</p> <p>If the followed parameter is set to false or is not included in the request, results are not filtered by followed activities.</p>	Boolean	query	No
after	<p>An activity ID to seek to. Activity entries up to and including the specified ID are excluded from the results and do not count towards max. Useful for pagination. Commonly set to the lastSeen property from a previous query.</p>	integer	query	No
max	<p>Maximum number of activity entries to return. This does not guarantee that max entries are returned. It does guarantee that the number of entries returned won't exceed max. Server-side filtering may exclude some activity entries for permissions reasons.</p> <p>If the max parameter is not included in the request, only the most recent 100 activities are returned.</p>	integer	query	No

Parameter	Description	Type	Parameter Type	Required
fields	An optional comma-separated list (or array) of fields to show. Omitting this parameter or passing an empty value shows all fields.	string	query	No

Example usage

Get a list of all activities

```
curl -u "username:ticket" "https://my-swarm-host/api/v10/activity"
```

Swarm responds with:

Tip

By default, this is limited to the 100 most recent activities. This can be changed by adding a **max** parameter value to the request.

```
{
  "error": null,
  "messages": null,
  "data": {
    "activity": [
      {
        "id": 1646,
        "type": "review",
        "link": [
          "review",
          {
            "review": "1236",
            "version": 4
          }
        ],
        "user": null
      },
      {
        "action": "Automated tests reported",
        "target": "review 1236",
```

```

    "preposition": "failed tests for",
    "description": "Make a review.",
    "details": [ ],
    "topic": "reviews/1236",
    "depotFile": null,
    "time": 1620228301,
    "behalfOf": null,
    "projects": [ ],
    "followers": [
        "bruno"
        "bob"
    ],
    "streams": {
        "review-1236",
        "user-",
        "personal-",
        "personal-macy.winter",
        "personal-super"
    },
    "change": 1236
},
{
    Other ids formatted as above
}
],
"totalCount": 100,
"lastSeen": 1646
}
}

```

Activity pagination

To get the second page of activity entries limited to a max of 1 and showing specified fields only (based on the previous example):

```
curl -u "username:ticket"
"https://myswarm.url/api/
v10
/activity?max=1&after=1646&fields=id,action,time,description,type"
```

Swarm again responds with:

```
{
  "error": null,
  "messages": null,
  "data": {
    "activity": [
      {
        "id": 1645,
        "type": "review",
        "action": "Automated tests reported",
        "description": "Make a review.",
        "time": 1626945417
      }
    ],
    "totalCount": 1,
    "lastSeen": 1645
  }
}
```

Get a list of all activities the requester follows

```
curl -u "username:ticket" "https://my-swarm-
host/api/v10/activity?followed=true"
```

Swarm responds with:

```
{
  "error": null,
  "messages": null,
  "data": {
    "activity": [
      {
        "id": 1646,
        "type": "review",
        "link": [
```

```

        "review",
        {
            "review": "1236",
            "version": 4
        }
    ],
    "user": null
    "action": "Automated tests reported",
    "target": "review 1236",
    "preposition": "failed tests for",
    "description": "Make a review.",
    "details": [ ],
    "topic": "reviews/1236",
    "depotFile": null,
    "time": 1620228301,
    "behalfOf": null,
    "projects": [ ],
    "followers": [
        "bruno"
        "bob"
    ],
    "streams": {
        "review-1236",
        "user-",
        "personal-",
        "personal-macy.winter",
        "personal-super"
    },
    "change": 1236
    },
    {
        Other ids formatted as above
    }
    ],
    "totalCount": 100,

```

```
"lastSeen": 1646
}
```

If a request fails

<error code>:

- **400** one of the following:
 - the **max** parameter value must be an integer greater than 0
 - the **after** parameter must be an integer greater than 0
- **404** record not found
- **500** internal error, request failed

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Get a list of all activities of a specified type

Summary

Get a list of all activities of a specified type.

GET /api/v10/activity/{type}

Description

Get a list of all activities of a specified type that the user has permission to view.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>followed</code>	Set the <code>followed</code> parameter to true to limit results to activities followed by the user making the request. If the <code>followed</code> parameter is set to <code>false</code> or is not included in the request, results are not filtered by followed activities.	Boolean	query	No
<code>after</code>	An activity ID to seek to. Activity entries up to and including the specified ID are excluded from the results and do not count towards <code>max</code> . Useful for pagination. Commonly set to the <code>lastSeen</code> property from a previous query.	integer	query	No

Parameter	Description	Type	Parameter Type	Required
max	Maximum number of activity entries to return. This does not guarantee that max entries are returned. It does guarantee that the number of entries returned won't exceed max . Server-side filtering may exclude some activity entries for permissions reasons. If the max parameter is not included in the request, only the most recent 100 activities are returned.	integer	query	No
fields	An optional comma-separated list (or array) of fields to show. Omitting this parameter or passing an empty value shows all fields.	string	query	No

Example usage

Get a list of all review activities

```
curl -u "username:ticket" "https://my-swarm-host/api/v10/activity/review"
```

Swarm responds with:

Tip

By default, this is limited to the 100 most recent activities. This can be changed by adding a **max** parameter value to the request.

```
{
  "error": null,
  "messages": null,
  "data": {
    "activity": [
```

```
{
  "id": 1646,
  "type": "review",
  "link": [
    "review",
    {
      "review": "1236",
      "version": 4
    }
  ],
  "user": null
  "action": "Automated tests reported",
  "target": "review 1236",
  "preposition": "failed tests for",
  "description": "Make a review.",
  "details": [ ],
  "topic": "reviews/1236",
  "depotFile": null,
  "time": 1620228301,
  "behalfOf": null,
  "projects": [ ],
  "followers": [
    "bruno"
    "bob"
  ],
  "streams": {
    "review-1236",
    "user-",
    "personal-",
    "personal-macy.winter",
    "personal-super"
  },
  "change": 1236
},
{
```



```
        Other ids formatted as above
    },
    ],
    "totalCount": 100,
    "lastSeen": 1646
  }
}
```

Activity pagination

To get the second page of review activity entries limited to a max of 1 and showing specified fields only (based on the previous example):

```
curl -u "username:ticket" "https://myswarm.url/api/v10
/activity/review?max=1&after=1646&fields=id,action,date,description
,type"
```

Swarm responds with:

```
{
  "error": null,
  "messages": null,
  "data": {
    "activity": [
      {
        "id": 1645,
        "type": "review",
        "action": "Automated tests reported",
        "description": "Make a review.",
        "time": 1626945417
      }
    ],
    "totalCount": 1,
    "lastSeen": 1645
  }
}
```

Get a list of all review activities the requester follows

```
curl -u "username:ticket" "https://my-swarm-
host/api/v10/activity/review?followed=true"
```

Swarm responds with:

```
{
  "error": null,
  "messages": null,
  "data": {
    "activity": [
      {
        "id": 1646,
        "type": "review",
        "link": [
          "review",
          {
            "review": "1236",
            "version": 4
          }
        ],
        "user": null
        "action": "Automated tests reported",
        "target": "review 1236",
        "preposition": "failed tests for",
        "description": "Make a review.",
        "details": [ ],
        "topic": "reviews/1236",
        "depotFile": null,
        "time": 1620228301,
        "behalfOf": null,
        "projects": [ ],
        "followers": [
          "bruno"
          "bob"
        ],
        "streams": {
          "review-1236",
          "user-",
          "personal-",
          "personal-macy.winter",

```

```
        "personal-super"
      },
      "change": 1236
    },
    {
      Other ids formatted as above
    }
  ],
  "totalCount": 100,
  "lastSeen": 1646
}
```

If a request fails

<error code>:

- **400** one of the following:
 - the **max** parameter value must be an integer greater than 0
 - the **after** parameter must be an integer greater than 0
- **404** record not found
- **500** internal error, request failed

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Get a list of all activities for a specified review

Summary

Get a list of all activities for a specified review.

GET /api/v10/reviews/{id}/activity

Description

Get a list of all activities for a specified review that the user has permission to view.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>followed</code>	Set the <code>followed</code> parameter to true to limit results to activities followed by the user making the request. If the <code>followed</code> parameter is set to <code>false</code> or is not included in the request, results are not filtered by followed activities.	Boolean	query	No
<code>after</code>	An activity ID to seek to. Activity entries up to and including the specified ID are excluded from the results and do not count towards <code>max</code> . Useful for pagination. Commonly set to the <code>lastSeen</code> property from a previous query.	integer	query	No

Parameter	Description	Type	Parameter Type	Required
max	Maximum number of activity entries to return. This does not guarantee that max entries are returned. It does guarantee that the number of entries returned won't exceed max . Server-side filtering may exclude some activity entries for permissions reasons. If the max parameter is not included in the request, only the most recent 100 activities are returned.	integer	query	No
fields	An optional comma-separated list (or array) of fields to show. Omitting this parameter or passing an empty value shows all fields.	string	query	No

Example usage

Get a list of all activities for review 12345

```
curl -u "username:ticket" "https://my-swarm-host/api/v10/reviews/12345/activity"
```

Swarm responds with:

Tip

By default, this is limited to the 100 most recent activities. This can be changed by adding a **max** parameter value to the request.

```
{
  "error": null,
  "messages": [],
  "data": {
    "activity": [
```

```
{
  "id": 790,
  "type": "review",
  "link": [
    "review",
    {
      "review": "12345",
      "version": 1
    }
  ],
  "user": "bruno",
  "action": "changed author of",
  "target": "review 12345 (revision 1)",
  "preposition": "for",
  "description": "Has comment that is outside of the Diff
chunks.",
  "details": {
    "author": {
      "oldAuthor": "dai",
      "newAuthor": "earl"
    }
  },
  "topic": "reviews/12345",
  "depotFile": null,
  "time": 1640102250,
  "behalfOf": null,
  "projects": {
    "jam": [
      "main"
    ]
  },
  "followers": [],
  "streams": {
    "0": "review-12345",
    "1": "user-bruno",
```

```
        "2": "personal-bruno",
        "3": "project-jam",
        "4": "group-Administrators",
        "6": "personal-dai",
        "7": "personal-earl",
        "8": "personal-ona",
        "9": "personal-raj",
        "10": "personal-swarm"
    },
    "change": 12344
  },
  {
    Other ids formatted as above
  }
],
"totalCount": 20,
"lastSeen": 790
}
}
```

Activity pagination

To get the second page of review activity entries limited to a max of 1 and showing specified fields only (based on the previous example):

```
curl -u "username:ticket" "https://myswarm.url/api/v10
/reviews/12345/activity?max=1&after=790&fields=id,action,date,description,type"
```

Swarm responds with:

```
{
  "error": null,
  "messages": null,
  "data": {
    "activity": [
      {
        "id": 789,
        "type": "review",
        "action": "Automated tests reported",
```

```

        "description": "Make a review.",
        "time": 1626945417
      }
    ],
    "totalCount": 1,
    "lastSeen": 789
  }
}

```

If a request fails

<error code>:

- **400** one of the following:
 - the **max** parameter value must be an integer greater than 0
 - the **after** parameter must be an integer greater than 0
- **404** review not found
- **500** internal error, request failed

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```

Attachments: Swarm comment attachments

Get attachments details

Summary

Get attachment details

GET /api/v10/attachments?ids[]={id}

Description

Get details of attachments the user has permission to view.

Parameters

Parameter	Description	Type	Parameter Type	Required
id	Specify the attachments you want details returned for by specifying attachment ids on the request.	Array of ids (integers)	query	Yes

Example usage

Get details for attachments 1, 2, and 500

```
curl -u "username:ticket" "https://myswarm-url/api/v10/attachments?ids[]=1&ids[]=2&ids[]=500"
```

Swarm responds with details for attachment ids 1, 2, and 500:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "attachments": [
      {
        "id": 1,
        "name": "motorbike.png",
        "type": "image/png",
        "size": 19844,
        "depotFile": "//.swarm/attachments/0000000001-motorbike.png",
        "references": {
          "comment": [
            502
          ]
        }
      }
    ]
  }
}
```

```

    },
    {
      "id": 2,
      "name": "house.png",
      "type": "image/png",
      "size": 16109,
      "depotFile": "//.swarm/attachments/0000000002-house.png",
      "references": {
        "comment": [
          503
        ]
      }
    },
    {
      "id": 500,
      "name": "car.png",
      "type": "image/png",
      "size": 26123,
      "depotFile": "//.swarm/attachments/0000000500-car.png",
      "references": {
        "comment": [
          890
        ]
      }
    }
  ]
}

```

If a request fails

<error code>:

404 Attachment does not exist

HTTP/1.1 <response error code>

```
{
```

```
"error": <error code>,  
"messages": [{  
  "code" : "<code string>",  
  "text" : "<error message>"  
}],  
"data" : null  
}
```

Get attachment thumbnail

Summary

Get the thumbnail of an attachment.

GET /api/v10/attachments/{id}/content/{filename}/thumb

Description

Get the thumbnail of an attachment the user has permission to view.

Parameters

Parameter	Description	Type	Parameter Type	Required
id	Attachment id	integer	query	Yes
filename	This is the attachment filename and extension encoded to URL safe Base64 . Various tools are available online to encode the filename to URL safe Base64 , for example: https://www.base64encode.org	string	path	Yes

Example usage

Get the thumbnail for attachment ID 118 (**house.png** which is **aG91c2UucG5n** when encoded in **URL safe Base64**):

```
curl -u "username:ticket" "https://myswarm-url/api/v10/attachments/118/content/aG91c2UucG5n/thumb"
```

Swarm responds with the thumbnail for the attachment.

If a request fails

<error code>:

404 Attachment does not exist

```
HTTP/1.1 <response error code>

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Open an attachment

Summary

Open an attachment.

GET /api/v10/attachments/{id}/content/{filename}

Description

Open the attachment the user has permission view.

Parameter	Description	Type	Parameter Type	Required
id	Attachment id	integer	query	Yes
filename	This is the attachment filename and extension encoded to URL safe Base64 . Various tools are available online to encode the filename to URL safe Base64 , for example: https://www.base64encode.org	string	path	Yes

Example usage

Open attachment ID 118 (`house.png` which is `aG91c2UucG5n` when encoded in **URL safe Base64**):

```
curl -u "username:ticket" "https://myswarm-  
url/api/v10/attachments/118/content/aG91c2UucG5n"
```

Swarm opens the attachment.

If a request fails

<error code>:

404 Attachment does not exist

HTTP/1.1 <response error code>

```
{  
  "error": <error code>,  
  "messages": [{  
    "code" : "<code string>",  
    "text" : "<error message>"  
  }],  
  "data" : null  
}
```

Download an attachment

Summary

Download an attachment.

GET `/api/v10/attachments/{id}/content/{filename}/download`

Description

Download the attachment the user has permission to view.

Parameters

Parameter	Description	Type	Parameter Type	Required
id	Attachment id	integer	query	Yes
filename	This is the attachment filename and extension encoded to URL safe Base64 . Various tools are available online to encode the filename to URL safe Base64 , for example: https://www.base64encode.org	string	path	Yes

Example usage

Download attachment ID 118 (`house.png` which is `aG91c2UucG5n` when encoded in **URL safe Base64**):

```
curl -u "username:ticket" "https://myswarm-url/api/v10/attachments/118/content/aG91c2UucG5n/download"
```

Swarm responds by downloading the attachment.

If a request fails

<error code>:

404 Attachment does not exist

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Create an attachment

Summary

Create an attachment.

POST /api/v10/attachments

Description

Create a standalone attachment.

Once the attachment is created, Swarm dynamically assigns a unique id to the attachment. To add an attachment to a comment, pass the unique attachment id as a parameter using the ["Add a comment with an attachment to a review"](#) on [page 1093](#) API endpoint.

Parameters

Parameter	Description	Type	Parameter Type	Required
file	This is the complete path with the attachment file name.	string	path	Yes

Example usage

Create an attachment for the `house.png` file:

```
curl -X POST -H "Content-Type: multipart/form-data" -u "username:ticket" "https://myswarm-url/api/v10/attachments/" -F "file=@/home/perforce/client1/house.jpg"
```

Swarm responds with the details of the attachment:

```
HTTP/1.1 200 OK
```

```
{
  "error": null,
  "messages": [],
  "data": {
    "attachments": [
      {
        "id": 12,
        "name": "house.png",
        "type": "image/png",
```

```

    "size": 106689,
    "blur": "L1SPX_I,Mw~W~qofa{WV56yDcDkq",
    "depotFile": "//depot/Jam/MAIN/src/attachments/0000000118-
house.png",
    "references": []
  }
]
}
}

```

If a request fails

<error code>:

404 Attachment does not exist

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```

Changes: Swarm changes

Get a list of changes

Summary

Get a list of changes.

GET /api/v10/changes

Description

Get a list of changes the user is able to view.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>pending</code>	<p>Determines if all changelists, only pending changelists, or only submitted changelists are returned by the request.</p> <ul style="list-style-type: none"> ▪ Not set: All changelists (submitted and pending) are returned. ▪ <code>true</code>: Only pending changelists are returned. ▪ <code>false</code>: Only submitted changelists are returned. 	boolean	query	No
<code>user</code>	The <code>user</code> parameter enables you to return just the changelists created by a specific user.	string	query	No
<code>rootPath</code>	<p>The <code>rootPath</code> parameter enables you to return all the changelists in a specific root path. The <code>rootPath</code> parameter is encoded to URL safe Base64.</p> <p>Various tools are available online to encode the id to URL safe Base64, for example: https://www.base64encode.org</p>	string	query	No

Example usage

Get a list of changes

```
curl -u "username:ticket" "https://myswarm-url/api/v10/changes"
```

Swarm responds with the changes starting with the newest change:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "changes": [
      {
        "id": 12853,
        "date": "2022/04/14 08:59:59",
        "client": "swarm-test",
        "user": "bruno",
        "status": "submitted",
        "type": "public",
        "importedBy": null,
        "identity": null,
        "description": "Fix display bug. ",
        "jobStatus": null,
        "jobs": [],
        "stream": null,
        "files": [
          "//depot/main/myfile.txt#6"
        ]
      },
      {
        ...
      }
    ]
  }
}
```

```
}  
}
```

Get a list of the pending changes

```
curl -u "username:ticket" "https://myswarm-  
url/api/v10/changes?pending=true"
```

Swarm responds with a list of pending changes starting with the newest pending change:

```
HTTP/1.1 200 OK  
  
{  
  "error": null,  
  "messages": [],  
  "data" : {  
    "changes": [  
      {  
        "id": 12382,  
        "date": "2022/04/10 12:59:59",  
        "client": "swarm-test",  
        "user": "paula.boyle",  
        "status": "pending",  
        "type": "public",  
        "importedBy": null,  
        "identity": null,  
        "description": "Update license details. ",  
        "jobStatus": null,  
        "jobs": [],  
        "stream": null,  
        "files": [  
          "//depot/main/license.txt#2"  
        ]  
      },  
      {  
        ...  
        <other change ids formatted as above>  
        ...  
      }  
    ]  
  }  
}
```

```

    }
  ]
}
}

```

Get a list of changes created by Bruno

```
curl -u "username:ticket" "https://myswarm-
url/api/v10/changes?user=bruno"
```

Swarm responds with the changes created by Bruno starting with the newest change:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "changes": [
      {
        "id": 12853,
        "date": "2022/04/14 08:59:59",
        "client": "swarm-test",
        "user": "bruno",
        "status": "submitted",
        "type": "public",
        "importedBy": null,
        "identity": null,
        "description": "Fix display bug. ",
        "jobStatus": null,
        "jobs": [],
        "stream": null,
        "files": [
          "//depot/main/myfile.txt#6"
        ]
      },
      {
        ...
      }
    ]
  }
}
```

```
    <other change ids formatted as above>
    ...
  }
]
}
}
```

Get a list of changes in `//depot/Jam/MAIN/...`

`//depot/Jam/MAIN/...` encoded to **URL safe Base64** is **Ly9kZXBvdC9KYW0vTUFJT i8uLi4=**.

```
curl -u "username:ticket" "https://myswarm-
url/api/v10/changes?rootPath=Ly9kZXBvdC9KYW0vTUFJT i8uLi4="
```

Swarm responds with changelists in `//depot/Jam/MAIN/...` starting with the newest change:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "changes": [
      {
        "id": 12853,
        "date": "2022/04/14 08:59:59",
        "client": "swarm-test",
        "user": "paula.boyle",
        "status": "pending",
        "type": "public",
        "importedBy": null,
        "identity": null,
        "description": "Update the XML and Word files. ",
        "jobStatus": null,
        "jobs": [],
        "stream": null,
        "files": [
          "//depot/Jam/MAIN/src/anotherfile.xml#3",
          "//depot/Jam/MAIN/src/myfile.xml#15",
```

```

        "//depot/Jam/MAIN/src/yetanotherfile.docx#4"
    ]
},
{
    ...
    <other change ids formatted as above>
    ...
}
]
}
}

```

If a request fails

<error code>:

403 Insufficient permissions to access the changes

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```

Get a list of files on a change

Summary

Get a list of files on a change.

GET /api/v10/changes/{id}/files

Description

Get a list of files on a change. You must be authenticated to view the files on the change.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Change ID	integer	path	Yes

Example usage

Get a list of files on change 12106:

```
curl -u "username:ticket" "https://myswarm-url/api/v10/changes/12106/files"
```

Swarm responds with a list of files on change 12106:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "root": "//depot/triggers",
    "files": [
      {
        "depotFile": "//depot/triggers/swarm-trigger.conf",
        "action": "add",
        "type": "text",
        "rev": "1",
        "fileSize": "129",
        "digest": "635CC5C20C0C2F3534D94207C7F7FB25"
      },

```

```

    {
      "depotFile": "//depot/triggers/swarm-trigger.pl",
      "action": "add",
      "type": "text+x",
      "rev": "1",
      "fileSize": "51239",
      "digest": "0C0E14975CD4E6195E95805DEC98DA03"
    }
  ],
  "limited": false
}

```

If a request fails

<error code>:

- **403** Insufficient permissions to access the change
- **404** Change does not exist

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```

Get job information for jobs associated with a change

Summary

Get job information for jobs associated with a change

GET /api/v10/changes/{id}/jobs

Description

Get job information for jobs associated with a change. You must be authenticated to view the jobs associated with the change.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Change ID	integer	path	Yes

Example usage

Get job information for change 12344

```
curl -u "username:<ticket>" "https://myswarm-url/api/v10/changes/12344/jobs"
```

Swarm responds with:

```
HTTP/1.1 200 OK
```

```
{
  "error": null,
  "messages": [],
  "data" : {
    "jobs": [
      {
        "job": "job000020",
        "link": "/jobs/job000020",
        "fixStatus": "open",
        "description": "Fix final bugs for beta release.\n",
```

```

      "description-markdown": "<span class=\"first-line\">Fix final
bugs for beta release.</span>"
    }
  ]
}
}

```

If a request fails

<error code>:

- **403** Insufficient permissions to access the change
- **404** Change does not exist

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```

Add a job to a change

Summary

Add a job to a change.

PUT /api/v10/changes/{id}/jobs/{job_id}

Description

Add a job to a change. You must be authenticated to add a job to a change.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the

project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).

- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Change ID	integer	path	Yes
<code>job_id</code>	Job ID	integer	path	Yes

Example usage

Add a job to a change

Add job000020 to change12344

```
curl -u "username:<ticket>" PUT "https://myswarm-url/api/v10/changes/12344/jobs/job000020"
```

Swarm responds with all of the jobs associated with the change:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "jobs": [
      {
        "job": "job000020",
        "link": "/jobs/job000020",
        "fixStatus": "open",
        "description": "Fix final bugs for beta release.\n",
        "description-markdown": "<span class=\"first-line\">Fix final
bugs for beta release.</span>"
      }
    ]
  }
}
```

```

    ]
  }
}

```

Add a job to a change that is already associated with another job

Add job001001 to change12344

```
curl -u "username:<ticket>" PUT "https://myswarm-
url/api/v10/changes/12344/jobs/job001001"
```

Swarm responds with all of the jobs associated with the change:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "jobs": [
      {
        "job": "job000020",
        "link": "/jobs/job000020",
        "fixStatus": "open",
        "description": "Fix final bugs for beta release.\n",
        "description-markdown": "<span class=\"first-line\">Fix final
bugs for beta release.</span>"
      },
      {
        "job": "job001001",
        "link": "/jobs/job001001",
        "fixStatus": "fixed",
        "description": "Bug fixes by Terry\n",
        "descriptionMarkdown": "<span class=\"first-line\">Bug fixes
by Terry</span>"
      }
    ]
  }
}

```

If a request fails

<error code>:

- **403** Insufficient permissions to access the change
- **404** Change does not exist
- **500** Job does not exist

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Remove a job from a change

Summary

Remove a job from a change.

DELETE /api/v10/changes/{id}/jobs/{job_id}

Description

Remove a job from a change. You must be authenticated to remove a job from a change.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Change ID	integer	path	Yes
<code>job_id</code>	Job ID	integer	path	Yes

Example usage

Remove a job from a change

Remove job000020 from change12344:

```
curl -u "username:<ticket>" PUT "https://myswarm-url/api/v10/changes/12344/jobs/job000020"
```

Swarm responds with all of the jobs associated with the change:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "jobs": []
  }
}
```

Remove a job from a change that is associated with more than one job

Remove job000020 from change12344:

```
curl -u "username:<ticket>" PUT "https://myswarm-url/api/v10/changes/12344/jobs/job000020"
```

Swarm responds with all of the jobs associated with the change:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
```

```
"jobs": [
  {
    "job": "job001001",
    "link": "/jobs/job001001",
    "fixStatus": "fixed",
    "description": "Bug fixes by Terry\n",
    "description-markdown": "<span class=\"first-line\">Bug fixes
by Terry</span>"
  }
]
```

If a request fails

<error code>:

- **403** Insufficient permissions to access the change
- **404** Change does not exist
- **500** Job does not exist

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Comments: Swarm comments

Get comments for a review, change, or job

Summary

Get a list of comments for a review, change, or job.

GET /api/v10/comments/<topic_subject>/<topic_id>

Description

Get a list of comments for a review, change, or job. Only comments for the specified **topic_subject** and **topic_id** are returned. Valid **topic_subject** values are **reviews**, **changes**, and **jobs**.

For example:

- GET /api/v10/comments/reviews/1234
- GET /api/v10/comments/changes/1234
- GET /api/v10/comments/jobs/job001234

Parameters

Parameter	Description	Type	Parameter Type	Required	Default Value
<code>after</code>	A comment ID to seek to. Comments up to and including the specified ID are excluded from the results and do not count towards <code>max</code> . Useful for pagination. Commonly set to the <code>lastSeen</code> property from a previous query.	integer	query	No	
<code>max</code>	Maximum number of comments to return. This does not guarantee that <code>max</code> comments are returned. It does guarantee that the number of comments returned won't exceed <code>max</code> .	integer	query	No	100

Parameter	Description	Type	Parameter Type	Required	Default Value
ignoreArchived	Prevents archived comments from being returned.	boolean	query	No	
tasksOnly	Returns only comments that have been flagged as tasks.	boolean	query	No	
taskState	Limit the returned comments to ones that match the provided task state (one or more of open , closed , verified , or comment).	array (of strings)	query	No	
fields	An optional comma-separated list (or array) of fields to show for each comment. Omitting this parameter or passing an empty value shows all fields.	string	query	No	

Example usage

Get all comments for a review

Get all of the comments for review 885:

```
curl -u "username:ticket" "https://my-swarm-host/api/v10/comments/reviews/885"
```

Swarm responds with the comments for the specified review:

Tip

By default, this is limited to the 100 most recent comments. This can be changed by adding a **max** parameter value to the request.

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 51,
        "topic": "reviews/885",
        "context": [],
        "attachments": [],
        "flags": [],
        "taskState": "comment",
        "likes": [],
        "user": "bruno"
        "time": 1619823600,
        "updated": 1619823600,
        "edited": null,
        "body": "This is a really good comment."
        "readBy": []
      },
      {
        "id": 63,
        "topic": "reviews/885",
        "context": [],
        "attachments": [],
```

```
    "flags": [],
    "taskState": "comment",
    "likes": [],
    "user": "alice"
    "time": 1620799560,
    "updated": 1620799560,
    "edited": null,
    "body": "This is my comment.",
    "readBy": []
  },
  {
    "id": 95,
    "topic": "reviews/885",
    "context": [],
    "attachments": [],
    "flags": [],
    "taskState": "comment",
    "likes": [],
    "user": "bob"
    "time": 1620813784,
    "updated": 1620813784,
    "edited": null,
    "body": "This is another comment.",
    "readBy": []
  },
],
"lastSeen": 95,
"totalCount": 3
}
```

Get a list of the first two comments for a review

Get a list of the first two comments for review 885:

```
curl -u "username:ticket" "https://my-swarm-
host/api/v10/comments/reviews/885&max=2"
```

Swarm responds with a list of the first two comments for review 885 and a **lastSeen** value for pagination:

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 51,
        "topic": "reviews/885",
        "context": [],
        "attachments": [],
        "flags": [],
        "taskState": "comment",
        "likes": [],
        "user": "bruno"
        "time": 1619823600,
        "updated": 1619823600,
        "edited": null,
        "body": "This is a really good comment.",
        "readBy": []
      },
      {
        "id": 63,
        "topic": "reviews/885",
        "context": [],
        "attachments": [],
        "flags": [],
        "taskState": "comment",
        "likes": [],
        "user": "alice"
        "time": 1620799560,
        "updated": 1620799560,
        "edited": null,
        "body": "This is my comment.",
        "readBy": []
      }
    ]
  }
}
```

```

    },
  ],
  "lastSeen": 63,
  "totalCount": 2
}
}

```

Paginating a comment listing

To get the next page of a comments list (based on the example above):

```
curl -u "username:ticket" "https://my-swarm-host/api/v10/comments/reviews/885&max=2&after=63"
```

Swarm responds with the second page of results, if any comments are present after the last seen comment:

```

{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "topic": "reviews/885",
        "context": [],
        "attachments": [],
        "flags": [],
        "taskState": "comment",
        "likes": [],
        "user": "bob"
        "time": 1620813784,
        "updated": 1620813784,
        "edited": null,
        "body": "This is another comment.",
        "readBy": []
      },
    ],
    "lastSeen": 95,
    "totalCount": 1
  }
}

```

```
}
}
```

If a request fails

<error code>:

- **401** you are not authorized to view the comments in the specified review
- **404** one of the following:
 - the **topic_subject** specified is not valid, it must be a **reviews**, **changes**, or **jobs**
 - the **topic_id** id specified does not exist

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Get a comment by ID

Summary

Get comment details by comment id.

GET /api/v10/comments/<id>

Description

Get comment details by comment id if the user has permission to view it. If the comment is a reply, the parent comment id is also returned.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project

restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).

- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Example usage

Get comment id 1627

```
curl -u "username:ticket" "https://my-swarm-host/api/v10/comments/1627"
```

Swarm responds with the comment details:

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": {
      "id": 1627,
      "topic": "reviews/2004",
      "context": {
        "file": "//projects/acme/main/src/xml.c",
        "leftLine": 46,
        "rightLine": 48,
        "content": [
          " ",
          "    // Commodo magnis posuere lobortis.",
          "- tincidunt = cumLigula * 100;",
          "- data++;",
          "    int suspendisse = 0;"
        ],
        "type": "text",
        "review": 2004,
        "version": 1,
        "change": null,
        "md5": "e66e0d0d88b2f39fb721132adda7ecf5",
        "name": "xml.c",
```



```
    "line": 48
  },
  "attachments": [
    17,
    18,
    107,
    108,
    110,
    111
  ],
  "flags": [],
  "taskState": "comment",
  "likes": [],
  "user": "aaron.gleber",
  "time": 1649147498,
  "updated": 1650640678,
  "edited": 1650638878,
  "body": "This comment gives some examples of images you can
use.",
  "readBy": []
}
}
```

Get the parent comment of comment id 525

```
curl -u "username:ticket" "https://my-swarm-
host/api/v10/comments/525"
```

Swarm responds with comment id 501 which is the parent of comment id 525:

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      "id": 525,
      "topic": "reviews/12373",
```

```

    "context": {
      "comment": 501,
      "version": 2
    },
    "attachments": [],
    "flags": [],
    "taskState": "comment",
    "likes": [],
    "user": "alice"
    "time": 1622712940,
    "updated": 1622712940,
    "edited": null,
    "body": "dao save test 1",
    "readBy": [],
    "notify": "delayed"
  }
}

```

If a request fails

<error code>:

- **401** you are not authorized to view the comments
- **404** the comment **id** specified does not exist

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```

Add a comment

Summary

Add a comment to a review, change, or job.

POST /api/v10/comments/<topic_subject>/<topic_id>

Description

Add a comment to a review, change, or job. The user that makes the API request is set as comment author.

Valid **topic_subject** values are **reviews**, **changes**, and **jobs**.

For example:

- POST /api/v10/comments/reviews/1234
- POST /api/v10/comments/changes/1234
- POST /api/v10/comments/jobs/job001234

Parameters

Parameter	Description	Type	Parameter Type	Required
body	The text content of the comment.	string	form	Yes
taskState	Optional task state of the comment. Valid values when adding a comment are comment (default) and open . This creates a plain comment or opens a task, respectively. If the taskState parameter is not in the request, the task state is set to comment .	string	form	No

Parameter	Description	Type	Parameter Type	Required
notify	<p>Comment notifications can be delayed, see "Comment notification delay" on page 355. The notify parameter enables you to set when the notification for the comment is sent.</p> <p>Valid options are:</p> <ul style="list-style-type: none">▪ delayed (default): notifications for this comment are delayed by the time set in "Comment notification delay" on page 585▪ immediate: the notification for this comment is sent immediately▪ silent: no notification is sent for this comment	string	query	No

Parameter	Description	Type	Parameter Type	Required
context	<p>Used to add the comment to a file in the review.</p> <p>Valid fields are:</p> <ul style="list-style-type: none">▪ file mandatory unless attribute or comment are set: File to comment on. Valid only for changes and reviews topics. For example, <code>//depot/main/README.txt</code>.▪ leftline optional, but if specified, you must also specify the rightline and content parameters. Integer: Left-side diff line number to attach the inline comment to. Valid only for changes and reviews topics.▪ rightline optional, but if specified, you must also specify the leftline and content parameters. Integer: Right-side diff line number to attach the inline comment to. Valid only for changes and reviews topics.	array	form	No

Parameter	Description	Type	Parameter Type	Required
	<ul style="list-style-type: none"> ■ content optional, but if specified, you must also specify the leftline and rightline parameters. Array of strings: Provide the content of the codeline the comment is on and the four preceding codelines. This is used to specify a short excerpt of context in case the lines being commented on change during the review. When set to null, Swarm makes an effort to build the content on its own. Because this involves file operations, it might be slow. ■ version optional, integer: With a reviews topic, this field specifies which version to attach the comment to. ■ attribute optional: Set to description to comment on the review description. ■ comment optional, integer: Set to the comment id this comment is replying to. 			

Example usage

Add a comment to a review

Add a comment to review 885, make the comment an open task, and send the notification immediately:

```
curl -X POST -H "Content-Type: application/json" -u
"username:ticket" -d "@mybodyfilename.txt" "https://my-swarm-
host/api/v10/comments/reviews/885?notify=immediate"
```

The "mybodyfilename.txt" file contains the authenticated user's comment and makes the comment an open task:

```
{
  "body": "This is another comment.",
```

```
"taskState": "open"
}
```

Swarm responds the comment entity:

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 95,
        "topic": "reviews/885",
        "context": [],
        "attachments": [],
        "flags": [],
        "taskState": "open",
        "likes": [],
        "user": "bob"
        "time": 1620813784,
        "updated": 1620813784,
        "edited": null,
        "body": "This is another comment.",
        "readBy": [],
        "notify": "immediate"
      },
    ],
    "userDelayedNotificationCount": 1
  }
}
```

Add a comment to a file in a review

Add a comment to the end of the `//depot/main/README.txt` file on review 123:

```
curl -X POST -H "Content-Type: application/json" -u
"username:ticket" -d "@mybodyfilename.txt" "https://my-swarm-
host/api/v10/comments/reviews/123"
```

The "mybodyfilename.txt" file contains the authenticated user's comment for the file:

```
{
  "body": "This is a comment on a file.",
  "context": {
    "file": "//depot/main/README.txt"
  }
}
```

Swarm responds the comment entity:

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 105,
        "topic": "reviews/123",
        "context": {
          "file": "//depot/main/README.txt",
          "leftLine": null,
          "rightLine": null,
          "content": [],
          "version": 1,
          "change": 122,
          "review": 123,
          "md5": "6af0e07c2d96afacaf2da234242c3cbd",
          "name": "README.txt"
          "line": null
        },
        "attachments": [],
        "flags": [],
        "taskState": "open",
        "likes": [],
        "user": "bob"
        "time": 1620813784,
        "updated": 1620813784,
        "edited": null,

```



```

        "body": "This is a comment on a file.",
        "readBy": [],
        "notify": "delayed"
    },
],
"userDelayedNotificationCount": 1
}
}

```

Add a comment to a line in a file and include context for the comment

Tip

To populate the **content** field, you need to know the content of the line you are commenting on and the content of the four lines before the line you are commenting on

Add a comment to a line in a file on review 110 with the following information:

- Full filepath and filename of the file to comment on
- Left and right line numbers of the diff to comment on
- Text of the codeline you are commenting on and the four codelines before that to provide context for the comment
- Review version to add the comment to

```
curl -X POST -H "Content-Type: application/json" -u
"username:ticket" -d "@mybodyfilename.txt" "https://my-swarm-
host/api/v10/comments/reviews/110"
```

The "mybodyfilename.txt" file contains the details for the new comment:

```

{
  "body": "This is a comment on a line in a file with context.",
  "context": {
    "file": "//depot/main/README.txt"
    "rightLine": 54,
    "leftLine": null,
    "content": {
      "FORTRAN default = \"\" ;\n",
      "LIBDIR default = /boot/develop/libraries ;\n",
      "-LINK default = $(CC) ;\n",
      "+LINK default = mwld ;\n",
    }
  }
}

```

```

    "+LINKFLAGS default = \"\" ;\n"
  },
  "version": 3
}
}

```

Swarm responds the comment entity:

```

{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 143,
        "topic": "reviews/110",
        "context": {
          "file": "//depot/main/README.txt"
          "leftLine": null,
          "rightLine": 54,
          "content": [
            "FORTRAN default = \"\" ;\n",
            "LIBDIR default = /boot/develop/libraries ;\n",
            "-LINK default = $(CC) ;\n",
            "+LINK default = mwld ;\n",
            "+LINKFLAGS default = \"\" ;\n"
          ],
          "version": 3,
          "change": 109,
          "review": 110,
          "md5": "6af0e07c2d96afacaf2da234242c3cbd",
          "name": "README.txt",
          "line": 54
        },
        "attachments": [],
        "flags": [],
        "taskState": "open",

```

```
    "likes": [],
    "user": "bob"
    "time": 1620813784,
    "updated": 1620813784,
    "edited": null,
    "body": "This is a comment on a line in a file with context.",
    "readBy": [],
    "notify": "delayed"
  },
],
"userDelayedNotificationCount": 1
}
```

Add a comment to a review description

Add a comment to the description of review 885 and delay the notification:

```
curl -X POST -H "Content-Type: application/json" -u
"username:ticket" -d "@mybodyfilename.txt" "https://my-swarm-
host/api/v10/comments/reviews/885?notify=delayed"
```

The "mybodyfilename.txt" file contains the authenticated user's comment and adds it to the review description:

```
{
  "body": "This is a comment on the description.",
  "context":{
    "attribute": "description"
  },
}
```

Swarm responds the comment entity:

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 99,
```

```

    "topic": "reviews/885",
    "context": {
      "attribute": "description"
    },
    "attachments": [],
    "flags": [],
    "taskState": "comment",
    "likes": [],
    "user": "bob"
    "time": 1620813784,
    "updated": 1620813784,
    "edited": null,
    "body": "This is a comment on the description.",
    "readBy": [],
    "notify": "delayed"
  },
],
"userDelayedNotificationCount": 1
}
}

```

Reply to a comment

You have two options when replying to a comment. Although the URLs for the options are different, they both do the same job:

- ["Reply to a comment using POST /api/v10/comments/<topic_subject>/<topic_id>" below](#)
- ["Reply to a comment using POST /api/v10/comments/<id>" on page 1066](#)

Reply to a comment using POST /api/v10/comments/<topic_subject>/<topic_id>

Add a reply to comment id 99 in review 885:

```
curl -X POST -H "Content-Type: application/json" -u
"username:ticket" -d "@mybodyfilename.txt" "https://my-swarm-
host/api/v10/comments/reviews/885"
```

The "mybodyfilename.txt" file contains the authenticated user's comment text and the comment the reply is for:

```

{
  "body": "This is a reply to a comment.",

```

```
"context": {
  "comment": 99
}
}
```

Swarm responds the comment entity:

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 109,
        "topic": "reviews/885",
        "context": {
          "comment": 99,
          "version": 2
        },
        "attachments": [],
        "flags": [],
        "taskState": "comment",
        "likes": [],
        "user": "bruno"
        "time": 1622712940,
        "updated": 1622712940,
        "edited": null,
        "body": "This is a reply to a comment.",
        "readBy": [],
        "notify": "delayed"
      }
    ],
    "userDelayedNotificationCount": 1
  }
}
```

Reply to a comment using POST /api/v10/comments/<id>

Add a reply to comment id 99 in review 885:

```
curl -X POST -H "Content-Type: application/json" -u
"username:ticket" -d "@mybodyfilename.txt" "https://my-swarm-
host/api/v10/comments/99"
```

The "mybodyfilename.txt" file contains the authenticated user's comment text and the comment the reply is for:

```
{
  "body": "This is a reply to a comment.",
  "topic": "reviews/885"
}
```

Swarm responds the comment entity:

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 120,
        "topic": "reviews/885",
        "context": {
          "comment": 99,
          "version": 2
        },
        "attachments": [],
        "flags": [],
        "taskState": "comment",
        "likes": [],
        "user": "bruno"
        "time": 1622712940,
        "updated": 1622712940,
        "edited": null,
        "body": "This is another reply to a comment.",
        "readBy": [],
        "notify": "delayed"
      }
    ]
  }
}
```

```
    }
  ]
  "userDelayedNotificationCount": 1
}
}
```

If a request fails

<error code>:

- **400** one of the following:
 - the **id** specified must be an integer greater than 0
 - invalid **context** field specified, valid fields are **file**, **leftline**, **rightline**, **content**, **version**, **attribute**, and **comment**
 - invalid **notify** value specified, value must be **delayed**, **immediate**, or **silent**
 - invalid **taskState** value specified, value must be **open** or **comment**
- **403**: you are not authorized to create this comment
- **404** the review, change, job, review version, file line number, or file specified does not exist
- **409** invalid **topic_subject** specified, value must be **reviews**, **changes**, or **jobs**
- **500** internal error, request failed

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Edit a comment

Summary

Edit a comment.

POST /api/v10/comments/{comment_id}/edit

Description

Important

Only the author of the original comment can edit it. Users with *admin* or *super* permissions cannot edit a comment by another user.

Edit a comment. Only the parameters specified in the API request are updated for the comment, the rest of the comment parameters are left unchanged.

When a comment is edited:

- The comment time stamp is marked as (*edited*) to show that the comment has been changed.
- All of the likes for the comment are removed.
- If users have marked the comment as read, the comment read status is reset to unread for those users.
- Swarm sends a notification to everyone involved in the review, including the review author and the reviewers, but not the editor of the comment.

Parameters

Parameter	Description	Type	Parameter Type	Required
body	The text content of the comment.	string	form	No
taskState	<p>The task state of the comment. Valid values are comment, open, addressed, and verified.</p> <p>Certain transitions (such as moving from open to verified) are not possible without an intermediate step (addressed, in this case).</p>	string	form	No

Parameter	Description	Type	Parameter Type	Required
notify	<p>Comment notifications can be delayed, see "Comment notification delay" on page 355. The notify parameter enables you to set when the notification for the comment is sent.</p> <p>Valid options are:</p> <ul style="list-style-type: none"> ▪ delayed (default): notifications for this comment are delayed by the time set in "Comment notification delay" on page 585 ▪ immediate: the notification for this comment is sent immediately ▪ silent: no notification is sent for this comment 	string	query	No

Edit a comment

Edit the text of comment id 95, set its task state to addressed, and send the comment notification immediately:

```
curl -X POST -H "Content-Type: application/json" -u
"username:ticket" -d "@mybodyfilename.txt" "https://my-swarm-
host/api/v10/comments/95/edit?notify=immediate"
```

The "mybodyfilename.txt" file contains the comment author's edited text and the new task state value:

```
{
  "body": "This is an edit to this comment.",
  "taskState": "addressed"
}
```

Swarm responds with:

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 95,
        "topic": "reviews/885",
        "context": [],
        "attachments": [],
        "flags": [],
        "taskState": "addressed",
        "likes": [],
        "user": "bob"
        "time": 1620813784,
        "updated": 1620813784,
        "edited": true,
        "body": "This is an edit to this comment.",
        "readBy": [],
        "notify": "immediate"
      },
    ],
    "userDelayedNotificationCount": 1
  }
}
```

If a request fails

<error code>:

- **400** one of the following:
 - the **comment-id** specified must be an integer greater than 0
 - invalid **notify** value specified, value must be **delayed**, **immediate**, or **silent**
 - invalid **taskState** value specified, value must be **open**, **comment**, **addressed**, or **verified**
- **403**: you are not authorized to edit this comment, you must be the comment author
- **404** the **comment_id** specified does not exist

```
HTTP/1.1 <response error code>

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Mark a comment as read

Summary

Mark a comment as read for the authenticated user.

POST /api/v10/comments/{id}/read

Description

Mark a comment as read for the authenticated user.

Example usage

Mark comment id **1626** as read for **Bruno**:

```
curl -X POST -u "bruno:ticket" "https://my-swarm-host/api/v10/comments/1626/read"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 1626,
        "topic": "reviews/2004",

```

```
"context": {
  "file": "//projects/acme/main/src/xml.c",
  "leftLine": null,
  "rightLine": 36,
  "version": 1,
  "content": [
    "/*",
    " * Fringilla dui feugiat natoque. Mauris fames cursus,
adipiscing justo",
    " * arcu consectetur. Condimentum sed hymenaeos sit,
pellentesque",
    "+ * arcu consectetur. Condimentum sed hymenaeos sit,
pellentesque",
    "+ * arcu consectetur. Condimentum sed hymenaeos sit,
pellentesque"
  ],
  "change": null,
  "review": null,
  "md5": "e66e0d0d88b2f39fb721132adda7ecf5",
  "name": "xml.c",
  "line": 36
},
"attachments": [
  14,
  15,
  16
],
"flags": [
  "closed"
],
"taskState": "comment",
"likes": [],
"user": "aaron.gleber",
"time": 1649147456,
"updated": 1650981576,
```

```
    "edited": 1649147475,
    "body": "We have lots of repeated comments.",
    "readBy": [
      "bruno"
    ],
    "notify": "delayed"
  }
]
```

If a request fails

<error code>:

- **400** the `comment-id` specified must be an integer greater than 0
- **404** the `comment_id` specified does not exist

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Mark a comment as unread

Summary

Mark a comment as unread for the authenticated user.

POST /api/v10/comments/{id}/unread

Description

Mark a comment as unread for the authenticated user.

Example usage

Mark comment id **1626** as unread for **Bruno**:

```
curl -X POST -u "bruno:ticket" "https://my-swarm-host/api/v10/comments/1626/unread"
```

Swarm responds with:

```
HTTP/1.1 200 OK
```

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 1626,
        "topic": "reviews/2004",
        "context": {
          "file": "//projects/acme/main/src/xml.c",
          "leftLine": null,
          "rightLine": 36,
          "version": 1,
          "content": [
            "/*",
            " * Fringilla dui feugiat natoque. Mauris fames cursus,
adipiscing justo",
            " * arcu consetetur. Condimentum sed hymenaeos sit,
pellentesque",
            "+ * arcu consetetur. Condimentum sed hymenaeos sit,
pellentesque",
            "+ * arcu consetetur. Condimentum sed hymenaeos sit,
pellentesque"
          ],
          "change": null,
          "review": null,
          "md5": "e66e0d0d88b2f39fb721132adda7ecf5",
          "name": "xml.c",
```

```
        "line": 36
      },
      "attachments": [
        14,
        15,
        16
      ],
      "flags": [
        "closed"
      ],
      "taskState": "comment",
      "likes": [],
      "user": "aaron.gleber",
      "time": 1649147456,
      "updated": 1650981576,
      "edited": 1649147475,
      "body": "We have lots of repeated comments.",
      "readBy": [],
      "notify": "delayed"
    }
  ]
}
```

If a request fails

<error code>:

- **400** the `comment-id` specified must be an integer greater than 0
- **404** the `comment_id` specified does not exist

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  ]
}
```

Mark all comments read

Summary

Mark all comments on a review as read for the authenticated user.

POST /api/v10/reviews/{id}/comments/read

Description

Mark all comments on a review as read for the authenticated user.

Example usage

Mark comments on review 2004 as read for Vernon Renno:

```
curl -X POST -u "vernon.renno:ticket" "https://my-swarm-host/api/v10/reviews/2004/comments/read"
```

Swarm responds with:

HTTP/1.1 200 OK

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 1626,
        "topic": "reviews/2004",
        "context": {
          "file": "//projects/acme/main/src/xml.c",
          "leftLine": null,
          "rightLine": 36,
          "version": 1,
          "content": [
            "/*",
```



```
        " * Fringilla duis feugiat natoque. Mauris fames cursus,
adipiscing justo",
        " * arcu consectetur. Condimentum sed hymenaeos sit,
pellentesque",
        "+ * arcu consectetur. Condimentum sed hymenaeos sit,
pellentesque",
        "+ * arcu consectetur. Condimentum sed hymenaeos sit,
pellentesque"
    ],
    "change": null,
    "review": null,
    "md5": "e66e0d0d88b2f39fb721132adda7ecf5",
    "name": "xml.c",
    "line": 36
},
"attachments": [
    14,
    15,
    16
],
"flags": [
    "closed"
],
"taskState": "comment",
"likes": [],
"user": "aaron.gleber",
"time": 1649147456,
"updated": 1650981576,
"edited": 1649147475,
"body": "We have lots of repeated comments.",
"readBy": [
    "vernon.renno"
],
"notify": "delayed"
},
```

```

    {
      ...
      <other comment ids formatted as above>
      ...
    }
  ]
}
}

```

If a request fails

<error code>:

- **400** the review **id** specified must be an integer greater than 0
- **404** the review **id** specified does not exist

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```

Mark all comments unread

Summary

Mark all comments on a review as unread for the authenticated user.

POST /api/v10/reviews/{id}/comments/unread

Description

Mark all comments on a review as unread for the authenticated user.

Example usage

Mark comments on review **2004** as unread for **Vernon Renno**:

```
curl -X POST -u "vernon.renno:ticket" "https://my-swarm-host/api/v10/reviews/2004/comments/unread"
```

Swarm responds with:

```
HTTP/1.1 200 OK
```

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 1626,
        "topic": "reviews/2004",
        "context": {
          "file": "//projects/acme/main/src/xml.c",
          "leftLine": null,
          "rightLine": 36,
          "version": 1,
          "content": [
            "/*",
            " * Fringilla dui feugiat natoque. Mauris fames cursus,
adipiscing justo",
            " * arcu consetetur. Condimentum sed hymenaeos sit,
pellentesque",
            "+ * arcu consetetur. Condimentum sed hymenaeos sit,
pellentesque",
            "+ * arcu consetetur. Condimentum sed hymenaeos sit,
pellentesque"
          ],
          "change": null,
          "review": null,
          "md5": "e66e0d0d88b2f39fb721132adda7ecf5",
          "name": "xml.c",
          "line": 36
        },
        "attachments": [
```

```

        14,
        15,
        16
    ],
    "flags": [
        "closed"
    ],
    "taskState": "comment",
    "likes": [],
    "user": "aaron.gleber",
    "time": 1649147456,
    "updated": 1650981576,
    "edited": 1649147475,
    "body": "We have lots of repeated comments.",
    "readBy": [],
    "notify": "delayed"
},
{
    ...
    <other comment ids formatted as above>
    ...
}
]
}
}

```

If a request fails

<error code>:

- **400** the review **id** specified must be an integer greater than 0
- **404** the review **id** specified does not exist

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{

```

```
"code" : "<code string>",
"text" : "<error message>"
}],
"data" : null
}
```

Archive a comment

Summary

Archive a comment.

POST /api/v10/comments/{id}/archive

Description

Archive a comment. The Archive option is only available for top level comments, if the comment has replies they are archived with the parent comment.

Example usage

Archive comment id **1626**:

```
curl -X POST -u "bruno:ticket" "https://my-swarm-
host/api/v10/comments/1626/archive"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 1626,
        "topic": "reviews/2004",
        "context": {
          "file": "//projects/acme/main/src/xml.c",
          "leftLine": null,
          "rightLine": 36,

```

```
    "version": 1,
    "content": [
      "**",
      " * Fringilla dui feugiat natoque. Mauris fames cursus,
adipiscing justo",
      " * arcu consetetuer. Condimentum sed hymenaeos sit,
pellentesque",
      "+ * arcu consetetuer. Condimentum sed hymenaeos sit,
pellentesque",
      "+ * arcu consetetuer. Condimentum sed hymenaeos sit,
pellentesque"
    ],
    "change": null,
    "review": null,
    "md5": "e66e0d0d88b2f39fb721132adda7ecf5",
    "name": "xml.c",
    "line": 36
  },
  "attachments": [
    14,
    15,
    16
  ],
  "flags": [
    "closed"
  ],
  "taskState": "comment",
  "likes": [],
  "user": "aaron.gleber",
  "time": 1649147456,
  "updated": 1650988005,
  "edited": 1649147475,
  "body": "We have lots of repeated comments.",
  "readBy": [
    "vernon.renno"
  ]
}
```

```
    ],
    "notify": "delayed"
  },
  {
    "id": 1628,
    "topic": "reviews/2004",
    "context": {
      "file": "//projects/acme/main/src/xml.c",
      "leftLine": null,
      "rightLine": 36,
      "version": 1,
      "content": [
        "/*",
        " * Fringilla duis feugiat natoque. Mauris fames cursus,
adipiscing justo",
        " * arcu consectetur. Condimentum sed hymenaeos sit,
pellentesque",
        "+ * arcu consectetur. Condimentum sed hymenaeos sit,
pellentesque",
        "+ * arcu consectetur. Condimentum sed hymenaeos sit,
pellentesque"
      ],
      "change": null,
      "review": null,
      "md5": "e66e0d0d88b2f39fb721132adda7ecf5",
      "name": "xml.c",
      "line": 36,
      "comment": 1626
    },
    "attachments": [
      19
    ],
    "flags": [
      "closed"
    ]
  },
```

```

    "taskState": "comment",
    "likes": [],
    "user": "aaron.gleber",
    "time": 1649147599,
    "updated": 1650988005,
    "edited": 1649147611,
    "body": "This is a reply comment,",
    "readBy": []
  }
]
}

```

If a request fails

<error code>:

- **400** the **comment-id** specified must be an integer greater than 0
- **404** the **comment_id** specified does not exist

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```

Restore an archived comment

Summary

Restore an archived comment.

POST /api/v10/comments/{id}/unarchive

Description

Restore an archived comment. If the archived comment has replies, the replies are also restored.

Example usage

Restore archived comment id 1626:

```
curl -X POST -u "bruno:ticket" "https://my-swarm-host/api/v10/comments/1626/unarchive"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    {
      "id": 1626,
      "topic": "reviews/2004",
      "context": {
        "file": "//projects/acme/main/src/xml.c",
        "leftLine": null,
        "rightLine": 36,
        "version": 1,
        "content": [
          "/*",
          " * Fringilla duis feugiat natoque. Mauris fames cursus,
adipiscing justo",
          " * arcu consectetur. Condimentum sed hymenaeos sit,
pellentesque",
          "+ * arcu consectetur. Condimentum sed hymenaeos sit,
pellentesque",
          "+ * arcu consectetur. Condimentum sed hymenaeos sit,
pellentesque"
        ],
        "change": null,
        "review": null,

```

```

    "md5": "e66e0d0d88b2f39fb721132adda7ecf5",
    "name": "xml.c",
    "line": 36
  },
  "attachments": [
    14,
    15,
    16
  ],
  "flags": [],
  "taskState": "comment",
  "likes": [],
  "user": "aaron.gleber",
  "time": 1649147456,
  "updated": 1650989315,
  "edited": 1649147475,
  "body": "We have lots of repeated comments.",
  "readBy": [],
  "notify": "delayed"
},
{
  "id": 1628,
  "topic": "reviews/2004",
  "context": {
    "file": "//projects/acme/main/src/xml.c",
    "leftLine": null,
    "rightLine": 36,
    "version": 1,
    "content": [
      "/*",
      " * Fringilla dui feugiat natoque. Mauris fames cursus,
adipiscing justo",
      " * arcu consectetur. Condimentum sed hymenaeos sit,
pellentesque",
      "+ * arcu consectetur. Condimentum sed hymenaeos sit,

```

```

pellentesque",
    "+ * arcu consectetur. Condimentum sed hymenaeos sit,
pellentesque"
    ],
    "change": null,
    "review": null,
    "md5": "e66e0d0d88b2f39fb721132adda7ecf5",
    "name": "xml.c",
    "line": 36,
    "comment": 1626
  },
  "attachments": [
    19
  ],
  "flags": [],
  "taskState": "comment",
  "likes": [],
  "user": "aaron.gleber",
  "time": 1649147599,
  "updated": 1650989315,
  "edited": 1649147611,
  "body": "This is a reply comment,",
  "readBy": []
}
]
}
}

```

If a request fails

<error code>:

- **400** the `comment-id` specified must be an integer greater than 0
- **404** the `comment_id` specified does not exist

HTTP/1.1 <response error code>

```
{
```

```
"error": <error code>,  
"messages": [{  
  "code" : "<code string>",  
  "text" : "<error message>"  
}],  
"data" : null  
}
```

Add a reaction

Summary

Add a reaction.

POST /api/v10/comments/{id}/addReaction

Description

Add a thumbs up to a comment.

Parameters

Parameter	Description	Type	Parameter Type	Required
reactions	Used to add a reaction icon to a comment. Valid fields are: <ul style="list-style-type: none">▪ thumbsup mandatory, added to a comment.	array	form	Yes

Example usage

Add a thumbs up to the comment id 28:

```
curl -X POST -H "Content-Type: application/json" -u  
"username:ticket" -d "@mybodyfilename.txt" "https://my-swarm-  
host/api/v10/comments/28/addReaction"
```

The "mybodyfilename.txt" file contains the username and reaction icon:

```
{  
  "reactions": {
```

```
    "thumbsup": "bruno"
  }
}
```

Swarm responds with:

HTTP/1.1 200 OK

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 28,
        "topic": "reviews/12625",
        "context": {
          "file": "//jam/main/src/make.c",
          "version": 1,
          "leftLine": 5,
          "rightLine": 5,
          "content": [
            " * Copyright 1993, 1995 Christopher Seiwald.",
            " *",
            "- * This file is part of Jam - see jam.c for Copyright
information. This is change 1 and this is change 2 and change 3 change
4",
            "+ * This file is part of Jam - see jam.c for Copyright
information. This is change 1 and this is change 2 and change 3 change
4 and change 4",
            " */"
          ],
          "change": null,
          "review": null,
          "md5": "8e79caff1ae9cdbd21b1581b554331ba",
          "name": "make.c",
          "line": 5
        }
      ]
    }
  }
}
```

```

    },
    "attachments": [],
    "flags": [],
    "taskState": "comment",
    "likes": [],
    "user": "bruno",
    "time": 1690801379,
    "updated": 1690801556,
    "edited": null,
    "body": "THis is test reaction comment",
    "readBy": [],
    "notify": "delayed",
    "reactions": {
      "thumbsup": [
        "bruno"
      ]
    }
  }
]
}
}

```

If a request fails

<error code>:

- **400** the `comment-id` specified must be an integer greater than 0
- **404** the `comment_id` specified does not exist

Remove a reaction

Summary

Remove a reaction.

DELETE /api/v10/comments/{id}/removeReaction

Description

Removes the thumbs up from a comment.

Parameters

Parameter	Description	Type	Parameter Type	Required
reactions	Used to remove a reaction icon from a comment. Valid fields are: <ul style="list-style-type: none">▪ thumbsup mandatory, removed from a comment.	array	form	Yes

Example usage

Remove the thumbs up from the comment id 28:

```
curl -X DELETE -H "Content-Type: application/json" -u "username:ticket" -d "@mybodyfilename.txt" "https://my-swarm-host/api/v10/comments/28/removeReaction"
```

The "mybodyfilename.txt" file contains the username and the reaction icon to be removed from the comment:

```
{
  "reactions": {
    "thumbsup": "bruno"
  }
}
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 28,
        "topic": "reviews/12625",
        "context": {
          "file": "//jam/main/src/make.c",
```

```
    "version": 1,
    "leftLine": 5,
    "rightLine": 5,
    "content": [
      " * Copyright 1993, 1995 Christopher Seiwald.",
      " *",
      "- * This file is part of Jam - see jam.c for Copyright
information. This is change 1 and this is change 2 and change 3 change
4",
      "+ * This file is part of Jam - see jam.c for Copyright
information. This is change 1 and this is change 2 and change 3 change
4 and change 4",
      " */"
    ],
    "change": null,
    "review": null,
    "md5": "8e79caff1ae9cdbd21b1581b554331ba",
    "name": "make.c",
    "line": 5
  },
  "attachments": [],
  "flags": [],
  "taskState": "comment",
  "likes": [],
  "user": "bruno",
  "time": 1690801379,
  "updated": 1690801556,
  "edited": null,
  "body": "THis is test reaction comment",
  "readBy": [],
  "reactions": [],
  "notify": "delayed"
}
]
```



```
}
```

If a request fails

<error code>:

- **400** the `comment-id` specified must be an integer greater than 0
- **404** the `comment_id` specified does not exist

Add a comment with an attachment to a review

Summary

Add a comment with an attachment to a review.

POST `/api/v10/comments/reviews/{id}`

Description

Add a comment with an attachment to a review. This is a two-step process:

1. Create a stand-alone attachment.
Once the attachment is created, Swarm assigns a unique id to the attachment. The unique attachment id is returned in the response. See ["Create an attachment" on page 1027](#).
2. Pass the unique attachment id as a parameter using the **POST** `/api/v10/comments/reviews/{id}` API endpoint.
You have now successfully added a comment with an attachment to a review.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Review ID.	integer	path	Yes
<code>body</code>	The text content of the comment.	string	form	Yes

Parameter	Description	Type	Parameter Type	Required
taskState	<p>Optional task state of the comment. Valid values when adding a comment are comment (default) and open. This creates a plain comment or opens a task, respectively.</p> <p>If the taskState parameter is not in the request, the task state is set to comment.</p>	string	form	No
notify	<p>Comment notifications can be delayed, see "Comment notification delay" on page 355. The notify parameter enables you to set when the notification for the comment is sent.</p> <p>Valid options are:</p> <ul style="list-style-type: none"> ▪ delayed (default): notifications for this comment are delayed by the time set in "Comment notification delay" on page 585 ▪ immediate: the notification for this comment is sent immediately ▪ silent: no notification is sent for this comment 	string	query	No

Parameter	Description	Type	Parameter Type	Required
context	<p>Used to add the comment to a file in the review.</p> <p>Valid fields are:</p> <ul style="list-style-type: none"> ▪ file mandatory unless attribute or comment are set: File to comment on. Valid only for changes and reviews topics. For example, <code>//depot/main/README.txt</code>. ▪ type: the filetype of the file, text, binary, symlink, unicode, utf8, utf16, apple, or resource. For more information on filetypes, see the Base filetypes section of the <i>Helix Core Command-Line (P4) Reference</i>. ▪ version optional, integer: With a reviews topic, this field specifies which version to attach the comment to. ▪ leftline optional, but if specified, you must also specify the rightline and content parameters. Integer: Left-side diff line number to attach the inline comment to. Valid only for changes and reviews topics. 	array	form	No

Parameter	Description	Type	Parameter Type	Required
	<ul style="list-style-type: none"> ▪ rightline optional, but if specified, you must also specify the leftline and content parameters. Integer: Right-side diff line number to attach the inline comment to. Valid only for changes and reviews topics. ▪ content optional, but if specified, you must also specify the leftline and rightline parameters. Array of strings: Provide the content of the codeline the comment is on and the four preceding codelines. This is used to specify a short excerpt of context in case the lines being commented on change during the review. When set to null, Swarm makes an effort to build the content on its own. Because this involves file operations, it might be slow. 			
attachment s	Attachment id generated using the "Create an attachment" on page 1027 API endpoint.	integer	form	Yes

Example usage

Create a comment with attachment id 12, make the comment an open task, and delay sending the comment notification:

```
curl -X POST -H "Content-Type: application/json" -u "username:ticket" -d "@mybodyfilename.txt" "https://my-swarm-host/api/v10/comments/reviews/885"
```

The "mybodyfilename.txt" file contains the authenticated user's comment, the attachment id 12, and makes the comment an open task:

```
{
  "body": "This is another comment.",
  "taskState": "open"
  "context" : {
    "file": "//depot/Jam/MAIN/src/jamgram.c",
    "type": "text",
    "version": 0,
    "leftLine": 13,
    "rightLine": 13,
    "content": [
      "#define yyerrok (yyerrflag=0)",
      "#define YYRECOVERING (yyerrflag!=0)",
      "-/* cfront 1.2 defines \"c_plusplus\" instead of \"__cplusplus\" */",
      "+/* cfront 1.2 defines \"c_plusplus\" instead of \"__cplusplus\" */ change 1",
      "#ifdef c_plusplus"
    ]
  },
  "attachments": [12]
}
```

Swarm responds the comment entity:

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 123,
        "topic": "reviews/12923",
        "context": {
          "file": "//depot/Jam/MAIN/src/jamgram.c",
          "version": 0,
          "leftLine": 13,
```

```

    "rightLine": 13,
    "content": [
      "#define yyerrok (yyerrflag=0)",
      "#define YYRECOVERING (yyerrflag!=0)",
      "-/* cfront 1.2 defines \"c_plusplus\" instead of \"__cplusplus\" */",
      "+/* cfront 1.2 defines \"c_plusplus\" instead of \"__cplusplus\" */ change 1",
      "#ifdef c_plusplus"
    ]
  }
  "attachments": [
    12
  ],
  "flags": [],
  "taskState": "open",
  "likes": [],
  "user": "bruno",
  "time": 1699003426,
  "updated": 1699003426,
  "edited": null,
  "body": "This is another comment.",
  "readBy": [],
  "notify": "delayed"
}
],
"userDelayedNotificationCount": 2
}
}

```

If a request fails

<error code>:

- **400** one of the following:
 - the **id** specified must be an integer greater than 0
 - invalid **context** field specified, valid fields are **file**, **leftline**, **rightline**, **content**, **version**, **attribute**, and **comment**
 - invalid **notify** value specified, value must be **delayed**, **immediate**, or **silent**
 - invalid **taskstate** value specified, value must be **open** or **comment**
- **403**: you are not authorized to create this comment
- **404** the attachment does not exist
- **500** internal error, request failed

Files: Swarm files

Get file data

Summary

Get file data for the specified file

GET /api/v10/files/{id}

Description

Get file data for the specified file.

Important

You must be authenticated to view the file information.

Parameters

Parameter	Description	Type	Parameter Type	Required
id	File ID, this is the full filepath including the filename encoded to URL safe Base64 . Various tools are available online to encode the id to URL safe Base64 , for example: https://www.base64encode.org	string	path	Yes

Parameter	Description	Type	Parameter Type	Required
fileRevision	<p>Specify the revision of the file using one of the following options:</p> <ul style="list-style-type: none"> Specific file revision number: <code>fileRevision=#nnn</code> File revision in a pending or submitted changelist: <code>fileRevision=@=nnn</code> File revision in a submitted changelist only: <code>fileRevision=@nnn</code> <p>Where <code>nnn</code> is the revision or changelist number you want.</p>	string	query	No

Example usage

Get file data for a specific revision of a file

Get file data for `//depot/main/myfile.txt` revision 3. The full filepath and filename must be encoded in **URL safe Base64**.

```
curl -u "username:ticket" "https://myswarm-url/api/v10/files/Ly9kZXBvdC9tYWluL215ZmlsZS50eHQ?fileRevision=#3"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "filename": "//depot/main/myfile.txt"
    "contentLink": "http://myswarm-url/view/depot/main/myfile.txt",
    "fileRevision": "#3",
    "contentType": "text/plain",
```



```
    "changeId": "12871"
  }
}
```

If a request fails

<error code>:

- **400** Invalid file revision specified. Must be one of the following: @ or @= followed by the changelist number, or # followed by the file revision number.
- **403** Insufficient permissions to access the file
- **404** File does not exist

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Get a file diff

Summary

Get a file diff for the specified file

GET /api/v10/files/{id}/diff

Description

Get a file diff for the specified file

Parameters

Parameter	Description	Type	Parameter Type	Required
id	<p>File ID, this is the full filepath including the filename for the diff encoded to URL safe Base64.</p> <p>Various tools are available online to encode the id to URL safe Base64, for example: https://www.base64encode.org</p>	string	path	Yes
from	<p>Specify the file revision to compare against in the form %23{revision-number}. This is the revision that would be displayed on the left side of a Swarm diff.</p> <p>If the from parameter is not specified, the head revision is used.</p>	string	query	No
to	<p>Specify the file revision being compared in the form %23{revision-number}. This is the revision that would be displayed on the right side of a Swarm diff.</p>	string	query	Yes
type	<p>Specify the file type for the diff:</p> <ul style="list-style-type: none"> ▪ file: All file types except stream specs. This is the default if the parameter is unset. ▪ stream: Stream spec files. 	string	query	No
lines	<p>Specify the number of lines of context around each diff returned in the response.</p>	integer	query	No

Parameter	Description	Type	Parameter Type	Required
<code>ignoreWs</code>	<p>Specify the whitespace options for the diff:</p> <ul style="list-style-type: none">▪ Whitespace characters such as space, tab, and newline characters are returned in the response. This is the default if the parameter is unset.▪ <code>0</code>: Whitespace characters such as space, tab, and newline characters are not returned in the response.▪ <code>1</code>: Whitespace changes are not highlighted. Used for file types where whitespace changes are not important.▪ <code>2</code>: Whitespace changes are highlighted. Used for file types where whitespace is important.	integer	query	No

Parameter	Description	Type	Parameter Type	Required
maxSize	<p>Specify the maximum sized file that can be returned as a diff:</p> <ul style="list-style-type: none"> ▪ -1: Unlimited file size. ▪ nnnn: where nnnn is the maximum file size in bytes. ▪ 0: Use the file size specified by the max_size configuration setting, see "max_size" on page 622. This is the default if the parameter is unset. <p>The lowest max_size value defaults to 50 bytes. If you set the max_size value lower than 50 bytes then Swarm will ignore it.</p> <p>When a file is in the edit state, depending on the max_size value, the following two flags are set in the response:</p> <ul style="list-style-type: none"> ▪ fetchContent: This flag determines if the size of a file is larger than the max_size value and accordingly enables or disables the show more context buttons in the file diff panel. For more information about the show more context buttons, see "File diff panel" on page 459. 	integer	query	No

Parameter	Description	Type	Parameter Type	Required
	<p>Files larger than the max_size value sets the fetchContent flag to false. When the fetchContent flag is set to false, all the show more context buttons in the file diff panel are disabled. When the fetchContent flag is set to true, all the show more context buttons in the file diff panel are enabled.</p> <ul style="list-style-type: none"> ■ isCut: This flag determines if a file diff size is larger than the max_size value and must be truncated. <p>File diff sizes larger than the max_size value sets the isCut flag value to the max_size value which truncates the diff. When the file diff size is smaller than the max_size value, the isCut flag is set to false.</p>			

Parameter	Description	Type	Parameter Type	Required
maxDiffs	<p>Specify the maximum number of diff sections (chunks) to return in the response:</p> <ul style="list-style-type: none"> ▪ -1: Returns an unlimited number of diff sections (chunks) in the response. ▪ nnnn: where nnnn is the maximum number of diff sections (chunks) returned in the response. ▪ 0: Return the number of diffs specified by the max_diffs configuration setting in the response, see "max_diffs" on page 602. This is the default if the parameter is unset. 	integer	query	No
offset	<p>Specify the diff section (chunk) to start displaying from. The number of diff sections displayed from the offset parameter value is governed by the maxDiffs parameter value specified.</p> <p>For example, if you just want to return the third diff section, set offset=2 and maxDiffs=1.</p>	integer	query	No
fromFile	<p>If the file has been moved or renamed, and edited, use the fromFile parameter to specify the original filepath including the filename encoded to URL safe Base64.</p> <p>Various tools are available online to encode the fromFile to URL safe Base64, for example: https://www.base64encode.org</p>	string	query	No

Example usage

Get a file diff

Get the file diff between revisions 233 and 234 of

`//projects/acme/main/src/aliquitaq/quosquo/do/53.txt` and don't return whitespace characters. The full filepath and filename must be encoded in **URL safe Base64**.

```
curl -u "username:ticket" "https://myswarm-url/api/v10
/files/Ly9wcm9qZWNOcy9hY2l1L21haW4vc3JjL2FsaXFpdGFxL3F1b3NxdW8vZG8v
NTMudHh0/diff?from=%233&to=%234&ignoreWs=0&type=file"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "isCut": false,
    "isSame": false,
    "header": "---
a///projects/acme/main/src/aliquitaq/quosquo/do/53.txt#3\n+++
b///projects/acme/main/src/aliquitaq/quosquo/do/53.txt#4\n",
    "summary": {
      "adds": 1,
      "deletes": 0,
      "updates": 0
    },
    "diffs": [
      "@@ -30,10 +30,12 @@\n このリリースは、バージョン1.0.1g の OpenSSL ライブラリ
をリンクすること\n によって、Heartbleed CVE-2014-0160 脆弱性に対処しています。 \n \n
-----
-----\n \n+Hello\n+\n 以前のリリースとの互換性\n -----\n
\n 1. PERFORCE ブローカは、2007.2以降にリリースされたPERFORCE クライアント及び\n
サーバと互換性があります。 \n"
    ],
    "paging": {
      "diffs": 1,
```

```

    "offset": null
  }
}
}

```

Get a file diff when the file size is greater than max_size value

Get the file diff between revisions 232 and 40 of `//jam/main/src/compile3.c` and don't return whitespace characters. The full filepath and filename must be encoded in **URL safe Base64**. In the response code, the **fetchContent** flag is set to false as the file size is greater than the **max_size** value.

```
curl -u "username:ticket" "https://myswarm-url/api/v10
/files/Ly9qYW0vbWFpbi9zcmMvY29tcGlsZTMuYw/diff?from=%232&to=%40%3D1
3444&ignoreWs=0&type=file"
```

Swarm responds with:

```
HTTP/1.1 200 OK
```

```

{
  "error": null,
  "messages": [],
  "data": {
    "isCut": 1000,
    "isSame": false,
    "header": "--- a///jam/main/src/compile3.c#2\n+++
b///jam/main/src/compile3.c@=13444\n",
    "summary": {
      "adds": 1,
      "deletes": 3,
      "updates": 3
    },
    "fetchContent": false,
    "diffs": [
      "@@ -23357,11 +23357,10 @@\n  *\tcompile_set() - compile the
\"set variable\" statement\n  *\tcompile_setcomp() - support for
`rule` - save parse tree\n  *\tcompile_setexec() - support for
`actions` - save execution string\n  *\tcompile_settings() - compile
the \"on =\" (set variable on exec) statement\n  *\tcompile_switch() -

```



```

compile 'switch' rule\n- * * External routines:\n * \n * \tcompile_
foreach() - compile the \"for x in y\" statement\n * \tcompile_if() -
compile 'if' rule\n * \tcompile_include() - support for 'include' -
call include() on file\n * \tcompile_local() - declare (and set) local
variables\n",
    @@ -24981,24 +24980,10 @@\n * \tcompile_set() - compile the
\"set variable\" statement\n * \tcompile_setcomp() - support for
`rule` - save parse tree \n * \tcompile_setexec() - support for
`actions` - save execution string \n * \tcompile_settings() - compile
the \"on =\" (set variable on exec) statement\n * \tcompile_switch() -
compile 'switch' rule\n- * * External routines:\n- * \n- * \tcompile_
foreach() - compile the \"for x \"
    ],
    "paging": {
        "diffs": 2,
        "offset": null
    }
}
}
}

```

Get a file diff when the file size is not greater than max_size value

Get the file diff between revisions 232 and 40 of `//jam/main/src/compile3.c` and don't return whitespace characters. The full filepath and filename must be encoded in **URL safe Base64**. In the response code, the **fetchContent** flag is set to true as the file size is not greater than the **max_size** value.

```

curl -u "username:ticket" "https://myswarm-url/api/v10
/files/Ly9qYW0vbWFpbi9zcmMvY29tcGlsZTMuYw/diff?from=%232&to=%40%3D1
3444&ignoreWs=0&type=file"

```

Swarm responds with:

```

HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "isCut": false,

```

```

    "isSame": false,
    "header": "--- a///jam/main/src/compile3.c#2\n+++
b///jam/main/src/compile3.c@=13444\n",
    "summary": {
        "adds": 1,
        "deletes": 3,
        "updates": 3
    },
    "fetchContent": true,
    "diffs": [
        "@@ -23357,11 +23357,10 @@\n  *\tcompile_set() - compile the
\"set variable\" statement\n  *\tcompile_setcomp() - support for
`rule` - save parse tree\n  *\tcompile_setexec() - support for
`actions` - save execution string\n  *\tcompile_settings() - compile
the \"on =\" (set variable on exec) statement\n  *\tcompile_switch() -
compile 'switch' rule\n- * * External routines:\n  *\n  *\tcompile_
foreach() - compile the \"for x in y\" statement\n  *\tcompile_if() -
compile 'if' rule\n  *\tcompile_include() - support for 'include' -
call include() on file\n  *\tcompile_local() - declare (and set) local
variables\n",
        "@@ -24981,24 +24980,10 @@\n  *\tcompile_set() - compile the
\"set variable\" statement\n  *\tcompile_setcomp() - support for
`rule` - save parse tree\n  *\tcompile_setexec() - support for
`actions` - save execution string\n  *\tcompile_settings() - compile
the \"on =\" (set variable on exec) statement\n  *\tcompile_switch() -
compile 'switch' rule\n- * * External routines:\n- *\n- *\tcompile_
foreach() - compile the \"for x in y\" statement\n- *\tcompile_if() -
compile 'if' rule\n- *\tcompile_include() - support for 'include' -
call include() on file\n- *\tcompile_local() - declare (and set) local
variables\n- *\tcompile_null() - do nothing -- a stub for parsing\n-
*\tcompile_rule() - compile a single user defined rule\n- *\tcompile_
rules() - compile a chain of rules\n- *\tcompile_set() - compile the
\"set variable\" statement\n- *\tcompile_setcomp() - support for
`rule` - save parse tree\n- *\tcompile_setexec() - support for
`actions` - save execution string\n- *\tcompile_settings() - compile

```

```

the \"on =\" (set variable on exec) statement\n- *\tcompile_switch() -
compile 'switch' rule\n */\n /\n * Copyright 1993, 1995 Christopher
Seiwald.\n * * External routines:\n *\n",
    @@ -36086,10 +36071,18 @@\n *\tcompile_set() - compile the
\"set variable\" statement\n *\tcompile_setcomp() - support for
`rule` - save parse tree\n *\tcompile_setexec() - support for
`actions` - save execution string\n *\tcompile_settings() - compile
the \"on =\" (set variable on exec) statement\n *\tcompile_switch() -
compile 'switch' rule\n+ **\n+ ** Adding new block\n+ ** Adding new
block\n+ ** Adding new block\n+ ** Adding new block\n+ **
Adding new block\n+ ** Adding new block\n+ ** Adding new
block\n * * External routines:\n *\n *\tcompile_foreach() - compile
the \"for x in y\" statement\n *\tcompile_if() - compile 'if' rule\n
*\tcompile_include() - support for 'include' - call include() on
file\n",
    @@ -57336,23 +57329,10 @@\n *\tcompile_setexec() - support
for `actions` - save execution string\n *\tcompile_settings() -
compile the \"on =\" (set variable on exec) statement\n *\tcompile_
switch() - compile 'switch' rule\n * * External routines:\n *\n-
*\tcompile_foreach() - compile the \"for x in y\" statement\n-
*\tcompile_if() - compile 'if' rule\n- *\tcompile_include() - support
for 'include' - call include() on file\n- *\tcompile_local() - declare
(and set) local variables\n- *\tcompile_null() - do nothing -- a stub
for parsing\n- *\tcompile_rule() - compile a single user defined
rule\n- *\tcompile_rules() - compile a chain of rules\n- *\tcompile_
set() - compile the \"set variable\" statement\n- *\tcompile_setcomp()
- support for `rule` - save parse tree\n- *\tcompile_setexec() -
support for `actions` - save execution string\n- *\tcompile_settings
() - compile the \"on =\" (set variable on exec) statement\n-
*\tcompile_switch() - compile 'switch' rule\n- * * External
routines:\n *\n *\tcompile_foreach() - compile the \"for x in y\"
statement\n *\tcompile_if() - compile 'if' rule\n *\tcompile_include
() - support for 'include' - call include() on file\n *\tcompile_
local() - declare (and set) local variables\n",
    @@ -99842,11 +99822,11 @@\n *\tcompile_setcomp() - support

```

```

for `rule` - save parse tree \n  *\tcompile_setexec() - support for
`actions` - save execution string \n  *\tcompile_settings() - compile
the "\"on =\"" (set variable on exec) statement\n  *\tcompile_switch() -
compile 'switch' rule\n  * * External routines:\n- *\n+ *****\n
*\tcompile_foreach() - compile the "\"for x in y\"" statement\n
*\tcompile_if() - compile 'if' rule\n  *\tcompile_include() - support
for 'include' - call include() on file\n  *\tcompile_local() - declare
(and set) local variables\n  *\tcompile_null() - do nothing -- a stub
for parsing\n",
    "@@ -99870,11 +99850,34 @@\n  *\tcompile_setcomp() - support
for `rule` - save parse tree \n  *\tcompile_setexec() - support for
`actions` - save execution string \n  *\tcompile_settings() - compile
the "\"on =\"" (set variable on exec) statement\n  *\tcompile_switch() -
compile 'switch' rule\n  * * External routines:\n- *\n+ * **\n+ ***\n+
***\n+ ***\n+  * **\n+ ***\n+ ***\n+ ***\n+  * **\n+ ***\n+ ***\n+
***\n+  * **\n+ ***\n+ ***\n+ ***\n+  * **\n+ ***\n+ ***\n+  *
**\n+ ***\n+ ***\n+ ***\n  *\tcompile_foreach() - compile the "\"for x
in y\"" statement\n  *\tcompile_if() - compile 'if' rule\n  *\tcompile_
include() - support for 'include' - call include() on file\n
*\tcompile_local() - declare (and set) local variables\n  *\tcompile_
null() - do nothing -- a stub for parsing\n",
    "@@ -99968,11 +99971,11 @@\n  *\tcompile_setcomp() - support
for `rule` - save parse tree \n  *\tcompile_setexec() - support for
`actions` - save execution string \n  *\tcompile_settings() - compile
the "\"on =\"" (set variable on exec) statement\n  *\tcompile_switch() -
compile 'switch' rule\n  * * External routines:\n- *\n+ *****\n
*\tcompile_foreach() - compile the "\"for x in y\"" statement\n
*\tcompile_if() - compile 'if' rule\n  *\tcompile_include() - support
for 'include' - call include() on file\n  *\tcompile_local() - declare
(and set) local variables\n  *\tcompile_null() - do nothing -- a stub
for parsing\n"
  ],
  "paging": {
    "diffs": 7,
    "offset": null
  }

```

```
    }  
  }  
}
```

If a request fails

<error code>:

- **400** Invalid file revision specified.
- **404** File does not exist or you do not have permission to view it

HTTP/1.1 <response error code>

```
{  
  "error": <error code>,  
  "messages": [{  
    "code" : "<code string>",  
    "text" : "<error message>"  
  }],  
  "data" : null  
}
```

File edit

Important

- This API endpoint only supports text files.
- File edit is a Technology Preview feature.

Features offered in Technology Preview are experimental and not guaranteed to always work as expected. If you have feedback and functionality suggestions, email techpreview@perforce.com.

Summary

Edit the file content and submit or shelve the file

PUT /api/v10/files/{id}

Description

Edit the file content and submit or shelve the file.

Note**Requirements:**

- You must have permissions to edit the file.
- If the file version changes while you are editing it or if you are not editing the Head revision of the file, committing your file changes will overwrite the current head revision of the file.
If you are unsure, shelving your change is a safer option because it will not overwrite the head revision of the file and you can do a manual resolve when you commit the shelf.
- File edit is enabled by default but can be disabled by your Swarm administrator, see ["allow_edits"](#) on page 623.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	File ID, this is the full filepath including the filename encoded to URL safe Base64 . Various tools are available online to encode the <code>id</code> to URL safe Base64 , for example: https://www.base64encode.org	string	path	Yes
<code>content</code>	Specify the entire file content to be shelved or submitted. Lines in the content must be separated with backslash and n characters <code>\n</code> .	string	body	Yes

Parameter	Description	Type	Parameter Type	Required
description	Specify a changelist description. Optional: you can create a review for your change or add it to an existing review by including a review keyword in the Change description . For instruction on creating a review and adding a changelist to a review using review keywords, see "Create a review" on page 682 and "Add a changelist to a review" on page 682 .	string	body	No
action	Specify the file action as either shelve or submit .	string	body	Yes

Example usage

Shelve file with edited content

Shelve `//depot/main/myfile.txt` file with edited content. The full filepath and filename must be encoded in **URL safe Base64**.

```
curl -X PUT -H "Content-Type: application/json" -u
"username:ticket" -d "@mybodyfilename.txt" "https://myswarm-
url/api/v10/files/Ly9kZXBvdC9tYWluL215ZmlsZS50eHQ"
```

The "mybodyfilename.txt" file contains the file content, description and the shelve action:

```
{
  "content": "My original line\nMy nice new line\nAnother original
line"
  "description": "Added a new line."
  "action": "shelve"
}
```

Swarm responds with:

```
HTTP/1.1 200 OK
```

```
{
```

```

"error": null,
"messages": [],
"data" : {
  "filename": "//depot/main/myfile.txt"
  "contentLink": "http://myswarm-url/view/depot/main/myfile.txt",
  "fileRevision": "@=12896",
  "contentType": "text/plain",
  "changeId": "12896"
}
}

```

If a request fails

<error code>:

- **400** Invalid file revision specified. Must be one of the following: @ or @= followed by the changelist number, or # followed by the file revision number.
- **403** Insufficient permissions to submit or shelve the file
- **404** File does not exist
- **500** Submit command failed because a newer revision of the file exists, you must resolve and submit your changes or revert your changes.

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```


Groups: Swarm groups

Get a list of groups

Summary

Get a list of groups.

GET /api/v10/groups

Description

Get a list of groups.

Parameters

Parameter	Description	Type	Parameter Type	Required
[ids]	An optional array of group ids to return in the response. Omitting this parameter or passing an empty value shows all groups.	string	query	No
[fields]	An optional comma-separated list (or array) of fields to show for each group. Omitting this parameter or passing an empty value shows all fields.	string	query	No

Parameter	Description	Type	Parameter Type	Required
[expand]	<p>An optional parameter to specify whether details of any subgroups in a group are returned in the response. Omitting the expand parameter or passing an empty value will not return subgroup details.</p> <ul style="list-style-type: none">▪ true: Returns details of any subgroups for each group returned. This is done recursively, so if a subgroup has a subgroup, its details are also returned.▪ false: Details of any subgroups in the groups are not returned.	boolean	query	No

Example usage

Get a list of all groups

```
curl -u "username:ticket" "https://myswarm-url/api/v10/groups"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "groups": [
      {
        "id": "Mercury",
```

```
    "description": "Group for project Mercury that includes
Development, QA, and Docs",
    "maxResults": null,
    "maxScanRows": null,
    "maxLockTime": null,
    "maxOpenFiles": "unset",
    "timeout": 43200,
    "passwordTimeout": null,
    "ldapConfig": null,
    "ldapSearchQuery": null,
    "ldapUserAttribute": null,
    "subgroups": [
        "mercury-dev",
        "mercury-docs",
        "mercury-qa"
    ],
    "owners": [
        "allison.clayborne",
        "jack.boone"
    ],
    "users": [
        "allison.clayborne",
        "jack.boone",
        "claire.brevia",
        "alex.randolph"
    ],
    "name": "Mercury",
    "emailFlags": [],
    "group_notification_settings": null,
    "emailAddress": null,
    "useMailingList": "mercury.group@nowhere.xxjzzj.com"
},
{
    ...
<Other group ids formatted as above>
```

```

    ...
  }
]
}
}

```

Limiting returned fields

To limit the returned fields when fetching groups:

```
curl -u "username:ticket" "https://myswarm-
url/api/v10/groups?fields=id,owners,users"
```

Swarm responds with:

```

HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "groups": [
      {
        "id": "Mercury",
      },
      "owners": [
        "allison.clayborne",
        "jack.boone"
      ],
      "users": [
        "allison.clayborne",
        "jack.boone",
        "claire.brevia",
        "alex.randolph"
      ],
    },
  },
  {
    ...
    <Other group ids formatted as above>
  }
}

```

```
    ...
  }
]
}
}
```

Limiting groups returned and expanding subgroup details

Return the `id`, `name`, and `subgroups` fields for the `Mercury` group and return subgroup details:

```
curl -u "username:ticket" "https://myswarm-  
url/api/  
v10/groups?fields=id,name,subgroups&ids=Mercury&expand=true"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "groups": [
      {
        "id": "Mercury",
        "name": "Mercury",
        "subgroups": [
          "mercury-dev",
          "mercury-docs",
          "mercury-qa"
        ]
      },
      {
        "id": "mercury-dev",
        "name": "mercury-dev",
        "subgroups": []
      },
      {
        "id": "mercury-docs",
        "name": "mercury-docs",
        "subgroups": []
      }
    ]
  }
}
```

```
    },  
    {  
      "id": "mercury-qa",  
      "name": "mercury-qa",  
      "subgroups": []  
    }  
  ]  
}  
}
```

Jobs: Swarm Perforce jobs

Get a list of Perforce jobs

Summary

Get a list of Perforce jobs.

GET /api/v10/jobs

Description

Get a list of Perforce jobs.

Parameters

Parameter	Description	Type	Parameter Type	Required
filter	Filter Perforce jobs by using a combination of a job field and value. The job fields and values available depend on your Perforce job configuration.	string	query	No

Parameter	Description	Type	Parameter Type	Required
max	Maximum number of Performce jobs to return. This does not guarantee that max jobs are returned. It does guarantee that the number of jobs returned won't exceed max . If not specified, the 50 most recent jobs are returned.	integer	query	No

Example usage

Get all Performce jobs

```
curl -u "username:ticket" "https://myswarm-url/api/v10/jobs"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "jobs": [
      {
        "job": "job001002",
        "link": "/jobs/job001002",
        "fixStatus": "open",
        "description": "The description is blank when not specified
for schedules in release 2.0. Please resolve.\n",
        "descriptionMarkdown": "<span class=\"first-line\">The
description is blank when not specified for schedules in release
2.0.</span><br><span class=\"more-lines\">\n &nbsp;Please
resolve.</span>"
      },
      {
        ...
      }
    ]
  }
}
```

```

    <Other job ids formatted as above>
    ...
  }
]
}
}

```

Get Perforce job id job000050

```
curl -u "username:ticket" "https://myswarm-
url/api/v10/jobs?filter=id=job000050"
```

Swarm responds with:

```

HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "jobs": [
      {
        "job": "job000050",
        "link": "/jobs/job000050",
        "fixStatus": "fixed",
        "description": "Update API call for status. Use the new
reponse style.\n",
        "descriptionMarkdown": "<span class=\"first-line\">Update API
call for status.</span><br><span class=\"more-lines\">\n &nbsp;Use the
new reponse style.</span>"
      }
    ]
  }
}

```

Get the 10 most recent Perforce jobs with a status of "fixed"

```
curl -u "username:ticket" "https://myswarm-
url/api/v10/jobs?filter=status=fixed&max=10"
```


Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "jobs": [
      {
        "job": "job000999",
        "link": "/jobs/job000999",
        "fixStatus": "fixed",
        "description": "Copyright text is incomplete. Please
resolve.\n",
        "descriptionMarkdown": "<span class=\"first-line\">Copyright
text is incomplete.</span><br><span class=\"more-lines\">\n
&nbsp;Please resolve.</span>"
      },
      {
        ...
        <Other job ids formatted as above>
        ...
      }
    ]
  }
}
```

If a request fails

<error code>:

500 Invalid field name specified, the error text contains the invalid field name.

```
HTTP/1.1 <response error code>

{
  "error": <error code>,
  "messages": [{
```

```
"code" : "<code string>",
"text" : "<error message>"
}],
"data" : null
}
```

Login: Swarm login

Get Swarm user session details

Summary

Get Swarm session details of a user.

GET /api/v10/session

Description

Get the Swarm session details of a user.

Example usage

Get the session details for bruno:

```
curl -u "bruno:<ticket>" https://my-swarm-host/api/v10/session
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [{
    "code": "user-login-successful",
    "text": "User logged in."
  }],
  "data" : {
    "user": {
      "User": "bruno",
      "Type": "standard",
      "Email": "bruno@swarm.local",
```

```
    "FullName": "bruno",
    "isAdmin": true,
    "isSuper": false
  }
}
```

If a request fails

<error code>:

- **401** Incorrect or missing credentials
- **403** Unauthorized

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Create a new Swarm user session

Summary

Create a new Swarm user session.

POST /api/v10/session

Description

Create a new Swarm user session.

When creating a new Swarm user session with a valid "[sso](#)" on [page 196](#) user, you must call "[Get SSO authentication URL](#)" on [page 1135](#) API endpoint to get the navigation URL. This navigation URL must be opened in a web browser to login the user. Once the user is logged in, a login successful message for this user is displayed in the "[Create a new Swarm user session](#)" [above](#) API endpoint.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>method</code>	<p>Login service type identifier. The default value is <code>basic</code>.</p> <ul style="list-style-type: none">▪ <code>basic</code>: uses basic authentication to login to Swarm▪ <code>sso</code>: uses "<code>sso</code>" on page 196 configurable as follows:<ul style="list-style-type: none">• <code>enabled</code> all users must use Helix Authentication Service to log in to Swarm. <p>The login page displays an input box to enter the email or username. Once the input box is filled, the Log in with SSO button is activated. See "Log in with SSO" on page 406.</p> <p>If <code>SAML</code> is configured when the <code>sso</code> configurable is enabled, the user can only log in to Swarm using the Log in with SSO button.</p>	string	body	No

Parameter	Description	Type	Parameter Type	Required
	<ul style="list-style-type: none"> • optional <p>Helix Authentication Service is available for users to log in to Swarm but is not enforced.</p> <p>The login page displays an input box to enter the email or username. Once the input box is filled, the Log in with SSO button is enabled. To log in with SSO, see "Log in with SSO" on page 406.</p> <p>The user can also manually log in to Swarm using the Log in with credentials button. To log in manually, see "Log in with a password" on page 405.</p> <p>If you switch to Log in with credentials and then decide to use single sign-on, click Log in with SSO.</p>			

Parameter	Description	Type	Parameter Type	Required
	<p>If SAML is configured when the sso configurable is set to optional, the user can either log in to Swarm using the Log in with SSO button or the Log in with credentials button.</p> <ul style="list-style-type: none"> • disabled Helix Authentication Service is not available to Swarm. This is the default value. <p>The login page displays two input boxes to enter the email or username, and to enter the password. Once the input boxes are filled, the Log in with credentials button is activated. See "Log in with a password" on page 405.</p>			
username	Username	string	body	Yes

Parameter	Description	Type	Parameter Type	Required
<code>password</code>	User ticket (recommended) or password	string	body	Yes
<div> Important If your Helix Core Server is set to security level 3 or higher, you must use a ticket value for the <code>password</code> parameter. </div>				
<code>remember</code>	<ul style="list-style-type: none"> ▪ <code>true</code>: User stays logged in between browser restarts. ▪ <code>false</code>: User does not stay logged in between browser restarts. This is the default if the <code>remember</code> parameter is not in the body or if it does not have a value. 	Boolean	body	No

Example usage

Create a new Swarm session for bruno using the `basic` method

```
curl -X POST -H "Content-Type: application/json" -d "@mybodyfilename.txt" "https://my-swarm-host/api/v10/session"
```

The "mybodyfilename.txt" file contains the login details for bruno:

```
{
  "method": "basic",
  "username": "bruno",
  "password": "A1AFB97F0F218DF7B122F229C7DECA46",
  "remember": false
}
```

Swarm responds with:


```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [
    {
      "code": "user-login-successful",
      "text": "User logged in."
    }
  ],
  "data" : {
    "user": {
      "User": "bruno",
      "Type": "standard",
      "Email": "bruno@swarm.local",
      "FullName": "bruno",
      "isAdmin": true,
      "isSuper": false
    }
  }
}
```

Create a new Swarm session for bruno using the sso method:

```
curl -X POST -H "Content-Type: application/json" -d
"@mybodyfilename.txt" "https://my-swarm-host/api/v10/session"
```

The "mybodyfilename.txt" file contains the login details for bruno:

```
{
  "method": "sso",
  "username": "bruno",
  "remember": false
}
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
```

```

"error": null,
"messages": [
  {
    "code": "user-login-successful",
    "text": "User logged in."
  }
],
"data" : {
  "user": {
    "User": "bruno",
    "Type": "standard",
    "Email": "bruno@swarm.local",
    "FullName": "bruno",
    "isAdmin": true,
    "isSuper": false
  }
}
}

```

If a request fails

<error code>:

- **401** Incorrect or missing credentials
- **403** Unauthorized

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```

Get SSO authentication URL

Summary

Get Swarm SSO authentication URL.

GET /api/v10/get-sso-auth-url

Description

Get Swarm SSO authentication URL.

Example usage

Get Swarm SSO authentication URL for bruno.

```
curl -u "bruno:<ticket>" https://my-swarm-host/api/v10/get-sso-auth-url
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "data": {
    "url":
    "https://10.153.120.29:3000/oidc/login/01HR6QMK0MPC5J506C2WG5ZZ0R?instanceId=none",
    "isValid": true,
    "errors": null
  },
  "messages": null,
  "error": null
}
```

If a request fails

<error code>:

- **401** Incorrect or missing credentials
- **403** Unauthorized

```
HTTP/1.1 <response error code>
```

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Login to Swarm with SAML

Important

Swarm now supports Helix Authentication Service (HAS) as a Single Sign-On (SSO) provider. This helps to simplify configuration and create a more robust SSO solution. The [SAML](#) configuration will be removed from Swarm in a later release.

If [SAML](#) is configured when the "sso" on [page 196](#) configurable is set to optional, the user can either log in to Swarm using the [Log in with SSO](#) button or the [Log in with credentials](#) button. If [SAML](#) is configured when the "sso" on [page 196](#) configurable is enabled, the user can only log in to Swarm using the [Log in with SSO](#) button.

Summary

Login to Swarm with SAML

POST /api/v10/saml/login

Description

Login to Swarm with SAML.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>method</code>	Login service type identifier is set to use saml configuration to login to Swarm.	string	body	No

Parameter	Description	Type	Parameter Type	Required
redirect	Options are: <ul style="list-style-type: none"> ▪ true or not specified: Swarm redirects the user to the HTTP_REFERER url or to the specified custom "logout_url" on page 612 if it has been set. ▪ false: Swarm does not redirect the user. 	string	query	No

Example usage

Login in to Swarm with SAML

Login in to Swarm with SAML.

```
curl -X POST -u "super:<ticket>" "https://myswarm-url/api/v10/saml/login"
```

On successful login Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "isValid": "true",
    "url": "<url to redirect to>"
  }
}
```

Login in to Swarm with SAML and redirect=false

Login in to Swarm with SAML and **redirect=false**

```
curl -X POST -u "super:<ticket>" "https://myswarm-url/api/v10/saml/login?redirect=false"
```

On successful login Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "isValid": "true"
  }
}
```

If a request fails

<error code>:

400 Error occurred with the SAML login

```
HTTP/1.1 <response error code>

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Logout of a Swarm user session

Summary

Logout of a Swarm user session.

DELETE /api/v10/session

Description

Logout of a Swarm user session.

Example usage

Logout bruno of a Swarm user session.

```
curl -X DELETE -u "bruno:<ticket>" https://my-swarm-host/api/v10/session
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [
    {
      "code": "user-logged-out",
      "text": "Successful Logout."
    }
  ],
  "data": {
    "url": "/"
  }
}
```

If a request fails

<error code>:

- **401** Incorrect or missing credentials
- **403** Unauthorized

```
HTTP/1.1 <response error code>

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Menus: Swarm main menu items

Get main menu details

Summary

Get details of the Swarm main menu items.

GET /api/v10/menus

Description

Get details of the Swarm main menu items including [custom menu items](#).

Note

Any users can get the Swarm main menu details, the user does not need to be authenticated.

Parameter	Description	Type	Parameter Type	Required
project	Specify a project in the request to return the menu items for the project.	string	query	No

Example usage

Fetch Swarm menu items

Get details of the Swarm main menu items:

```
curl "https://myswarm-url/api/v10/menus"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "menu": [
      {
        "id": "dashboard",
```



```
    "enabled": true,  
    "target": "/",  
    "cssClass": "component",  
    "title": "dashboard",  
    "priority": 100,  
    "roles": null  
  },  
  {  
    "id": "activity",  
    "enabled": true,  
    "target": "/activity/",  
    "cssClass": "activity",  
    "title": "activity",  
    "priority": 120,  
    "roles": null  
  },  
  {  
    "id": "reviews",  
    "enabled": true,  
    "target": "/reviews/",  
    "cssClass": "component",  
    "title": "reviews",  
    "priority": 130,  
    "roles": null  
  },  
  {  
    "id": "projects",  
    "enabled": true,  
    "target": "/projects/",  
    "cssClass": "projects",  
    "title": "projects",  
    "priority": 140,  
    "roles": null  
  },  
  {  

```

```
    "id": "files",
    "enabled": true,
    "target": "/files/",
    "cssClass": "files",
    "title": "files",
    "priority": 150,
    "roles": null
  },
  {
    "id": "changes",
    "enabled": true,
    "target": "/changes/",
    "cssClass": "changes",
    "title": "changes",
    "priority": 160,
    "roles": null
  },
  {
    "id": "jobs",
    "enabled": true,
    "target": "/jobs/",
    "cssClass": "jobs",
    "title": "jobs",
    "priority": 170,
    "roles": null
  },
  {
    "id": "groups",
    "enabled": true,
    "target": "/groups/",
    "cssClass": "groups",
    "title": "groups",
    "priority": 180,
    "roles": null
  },
}
```

```
{
  "id": "workflows",
  "enabled": true,
  "target": "/workflows/",
  "cssClass": "workflows",
  "title": "workflows",
  "priority": 190,
  "roles": null
},
{
  "id": "tests",
  "enabled": true,
  "target": "/testdefinitions/",
  "cssClass": "tests",
  "title": "tests",
  "priority": 200,
  "roles": null
},
{
  "id": "jenkins_blue",
  "enabled": true,
  "target": "http://my-jenkins.instance.blue.com",
  "cssClass": "custom_menu",
  "title": "Jenkins Build",
  "priority": 210,
  "roles": "authenticated"
}
],
"contextMetadata": []
}
```

Fetch menu items for a project (authenticated and project owner)

Get menu items for project Jam when the authenticated user owns the project:

```
curl -u "username:ticket" "https://myswarm-  
url/api/v10/menus?project=jam"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "menu": [
      {
        "id": "activity",
        "enabled": true,
        "target": "/projects/jam/activity/",
        "cssClass": "activity",
        "title": "activity",
        "priority": 120,
        "roles": null
      },
      {
        "id": "reviews",
        "enabled": true,
        "target": "/projects/jam/reviews/",
        "cssClass": "component",
        "title": "reviews",
        "priority": 130,
        "roles": null
      },
      {
        "id": "files",
        "enabled": true,
        "target": "/projects/jam/files/",
        "cssClass": "files",
        "title": "files",
        "priority": 150,
        "roles": null
      },
      {
```

```
    "id": "changes",
    "enabled": true,
    "target": "/projects/jam/changes/",
    "cssClass": "changes",
    "title": "changes",
    "priority": 160,
    "roles": null
  },
  {
    "id": "settings",
    "enabled": true,
    "target": "/projects/jam/settings/",
    "cssClass": "settings",
    "title": "settings",
    "priority": 200,
    "roles": null
  }
],
"contextMetadata": {
  "contextName": "jam"
}
}
```

Fetch menu items for a project (unauthenticated or not project owner)

Get menu items for project Jam without authentication or when the authenticated user is not a project owner:

```
curl "https://myswarm-url/api/v10/menus?project=jam"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
```

```
"menu": [  
  {  
    "id": "activity",  
    "enabled": true,  
    "target": "/projects/jam/activity/",  
    "cssClass": "activity",  
    "title": "activity",  
    "priority": 120,  
    "roles": null  
  },  
  {  
    "id": "reviews",  
    "enabled": true,  
    "target": "/projects/jam/reviews/",  
    "cssClass": "component",  
    "title": "reviews",  
    "priority": 130,  
    "roles": null  
  },  
  {  
    "id": "files",  
    "enabled": true,  
    "target": "/projects/jam/files/",  
    "cssClass": "files",  
    "title": "files",  
    "priority": 150,  
    "roles": null  
  },  
  {  
    "id": "changes",  
    "enabled": true,  
    "target": "/projects/jam/changes/",  
    "cssClass": "changes",  
    "title": "changes",  
    "priority": 160,  
  }  
]
```

```
        "roles": null
      }
    ],
    "contextMetadata": {
      "contextName": "Jam"
    }
  }
}
```

If a request fails

<error code>:

- **404** Project specified does not exist

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Projects : Swarm Projects

Get List of Projects

Summary

Get List of Projects and metadata

GET /api/v10/projects

Description

Returns a list of projects in Swarm that are visible to the current user.

Note

The branch **paths** field is always returned empty for all project branches returned with the project list endpoint. This is because calculating the paths for all of the project branches is an expensive operation when there are a lot of projects and paths. To check the branch **paths** for a specific project, use the ["Get Project Information" on page 1152](#) endpoint.

Tip

- **tests** and **deploy** fields:
 - **If a project has an owner:** only project owners and users with *super* user permissions can view the **tests** and **deploy** fields when fetching projects.
 - **If a project does not have an owner:** only project members and users with *super* user permissions can view the **tests** and **deploy** fields when fetching projects.
- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
fields	An optional comma-separated list or array of top level fields to return for the projects. Omitting this parameter or passing an empty value shows all fields.	string	query	No

Example usage**Fetch a list of projects**

To fetch all projects:

```
curl -u "username:ticket" "https://my-swarm-host/api/v10/projects"
```

Pagination is not currently supported by this endpoint. Swarm responds with a list of projects:


```
{
  "error": null,
  "messages": [],
  "data": {
    "projects": [
      {
        "id": "jam",
        "branches": [
          {
            "id": "main",
            "name": "Main",
            "workflow": null,
            "paths": [],
            "defaults": {
              "reviewers": []
            },
            "minimumUpVotes": null,
            "retainDefaultReviewers": false,
            "moderators": [],
            "moderators-groups": []
          },
          ...
          ...
        ],
        "defaults": {
          "reviewers": []
        },
        "deleted": false,
        "deploy": {
          "enabled": false,
          "url": null
        },
        "description": "This is the project Jam description",
        "emailFlags": {
          "change_email_project_users": "1",
```

```

        "review_email_project_members": "1"
    },
    "jobview": null,
    "members": [
        "Bruno",
        "raj"
    ],
    "minimumUpVotes": null,
    "name": "Jam",
    "owners": [],
    "private": true,
    "retainDefaultReviewers": false,
    "subgroups": [],
    "tests": {
        "enabled": false,
        "url": null,
        "postBody": null,
        "postFormat": "url"
    },
    "workflow": null,
},
{
    ...
    <other projects formatted as above>
    ...
},
],
}
}

```

Fetch specified fields for all projects

To fetch the **id**, **description**, and **members** fields for all projects, use one of the following methods:

- **Using a comma-separated list of top level fields:**

```
curl -u "username:ticket" "https://my-swarm-host/api/v10/projects?fields=id,description,members"
```

- **Using an array of top level fields:**

```
curl -u "username:ticket" "https://my-swarm-host/api/v10/projects?fields[]=id&fields[]=description&fields[]=members"
```

Pagination is not currently supported by this endpoint. Swarm responds with a list of all projects:

```
{
  "error": null,
  "messages": [],
  "data": {
    "projects": [
      {
        "id": "blue-book",
        "description": "This is a private project for use only by
users with sufficient clearance.",
        "members": [
          "alex.randolph",
          "allison.clayborne"
        ],
      },
      {
        "id": "jam",
        "description": "This the project Jam description.",
        "members": [
          "Bruno",
          "raj"
        ],
      },
      {
        "id": "jplugin",
        "description": "A Java plugin for continuous integration.",
        "members": [
          "allison.clayborne",
          "jack.boone",

```

```

        "steve.russell"
      ],
    },
  ],
}
}

```

If a request fails

<error code>:

400 metadata request invalid, boolean required

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```

Get Project Information

Summary

Get Project Information

GET /api/v10/projects/{id}

Description

Retrieve all information about a project or specify which fields you want to return for the project.

Tip

- **tests** and **deploy** fields:
 - **If a project has an owner:** only project owners and users with *super* user permissions can view the **tests** and **deploy** fields when fetching projects.

- **If a project does not have an owner:** only project members and users with *super* user permissions can view the **tests** and **deploy** fields when fetching projects.
- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
id	Project ID	string	path	Yes
fields	An optional comma-separated list or array of top level fields to return for the project. Omitting this parameter or passing an empty value shows all fields.	string	query	No

Example usage

Fetch all fields for a project

To fetch all fields for project jam:

```
curl -u "username:ticket" "https://my-swarm-host/api/v10/projects/jam"
```

Swarm responds with a project entity:

```
{
  "error": null,
  "messages": [],
  "data": {
    "projects": [
      {
        "id": "jam",
        "branches": [
```

```
{
  "id": "main",
  "name": "Main",
  "workflow": null,
  "paths": [
    "//depot/Jam/MAIN/...",
    "//jam/main/..."
  ],
  "defaults": {
    "reviewers": []
  },
  "minimumUpVotes": null,
  "retainDefaultReviewers": false,
  "moderators": [],
  "moderators-groups": []
},
...
...
],
"defaults": {
  "reviewers": []
},
"deleted": false,
"deploy": {
  "enabled": false,
  "url": ""
},
"description": "This the project Jam description",
"emailFlags": {
  "change_email_project_users": "1",
  "review_email_project_members": "1"
},
"jobview": "",
"members": [
  "Bruno",
```

```
        "raj"
      ],
      "minimumUpVotes": null,
      "name": "Jam",
      "owners": [],
      "private": false,
      "retainDefaultReviewers": false,
      "subgroups": [],
      "tests": {
        "enabled": false,
        "url": "",
        "postBody": "",
        "postFormat": "url"
      },
      "workflow": null
    },
  ],
}
```

Fetch specified fields for a project

To fetch the **id**, **description**, and **members** fields for project Jam, use one of the following methods:

- Using a comma-separated list of top level fields:

```
curl -u "username:ticket" "https://my-swarm-host/api/v10/projects/jam?fields=id,description,members"
```

- Using an array of top level fields:

```
curl -u "username:ticket" "https://my-swarm-host/api/v10/projects/jam?fields[]=id&fields[]=description&fields[]=members"
```

Swarm responds with a project entity containing the requested fields:

```
{
  "error": null,
  "messages": [],
  "data": {
```

```

    "projects": [
      {
        "id": "jam",
        "description": "This the project Jam description",
        "members": [
          "Bruno",
          "raj"
        ],
      },
    ],
  }
}

```

If a request fails

<error code>:

404 Project does not exist or you do not have permission to view it

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```

Get a list of Perforce jobs on a project

Summary

Get a list of Perforce jobs on a project.

GET /api/v10/projects/{project_id}/jobs

Description

Get a list of Perforce jobs on a project.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
filter	Filter Perforce jobs by using a combination of a job field and value. The job fields and values available depend on your Perforce job configuration. If you have a job filter selected for the project you can also limit jobs returned by field name and value, see Job filter in Project settings .	string	query	No
max	Maximum number of Perforce jobs to return. This does not guarantee that max jobs are returned. It does guarantee that the number of jobs returned won't exceed max . If not specified, the 50 most recent jobs are returned.	integer	query	No

Example usage**Get all Perforce jobs on the Acme project**

```
curl -u "username:ticket" "https://myswarm-url/api/v10/projects/acme/jobs"
```

Swarm responds with:

```

HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "jobs": [
      {
        "job": "job001002",
        "link": "/jobs/job001002",
        "fixStatus": "open",
        "subsystem": "backend",
        "description": "The description is blank when not specified
for schedules in release 2.0. Please resolve.\n",
        "descriptionMarkdown": "<span class=\"first-line\">The
description is blank when not specified for schedules in release
2.0.</span><br><span class=\"more-lines\">\n &nbsp;Please
resolve.</span>"
      },
      {
        ...
        <Other job ids formatted as above>
        ...
      }
    ]
  }
}

```

Get the 10 most recent Perforce jobs in project Acme

```
curl -u "username:ticket" "https://myswarm-
url/api/v10/projects/acme/jobs?max=10"
```

Swarm responds with:

```

HTTP/1.1 200 OK

{

```

```
"error": null,
"messages": [],
"data" : {
  "jobs": [
    {
      "job": "job000999",
      "link": "/jobs/job000999",
      "fixStatus": "fixed",
      "subsystem": "docs",
      "description": "Copyright text is incomplete. Please
resolve.\n",
      "descriptionMarkdown": "<span class=\"first-line\">Copyright
text is incomplete.</span><br><span class=\"more-lines\">\n
&nbsp;Please resolve.</span>"
    },
    {
      ...
      <Other job ids formatted as above>
      ...
    }
  ]
}
```

Get Performer jobs in the range of job000050 to job000059 that are part of project Acme

```
curl -u "username:ticket" "https://myswarm-
url/api/v10/projects/acme/jobs?filter=job00005*"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "jobs": [
```

```

    {
      "job": "job000050",
      "link": "/jobs/job000050",
      "fixStatus": "fixed",
      "subsystem": "frontend",
      "description": "Update API call for status. Use the new
reponse style.\n",
      "descriptionMarkdown": "<span class=\"first-line\">Update API
call for status.</span><br><span class=\"more-lines\">\n &nbsp;Use the
new repsonse style.</span>"
    }
  ]
}

```

Get all Performer jobs with a Job filter matching subsystem=frontend for project Acme

```
curl -u "username:ticket" "https://myswarm-
url/api/v10/projects/acme/jobs?filter=subsystem=frontend"
```

Swarm responds with:

```

HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "jobs": [
      {
        "job": "job000050",
        "link": "/jobs/job000050",
        "fixStatus": "fixed",
        "subsystem": "frontend",
        "description": "Update API call for status. Use the new
reponse style.\n",
        "descriptionMarkdown": "<span class=\"first-line\">Update API
call for status.</span><br><span class=\"more-lines\">\n &nbsp;Use the

```

```
new response style.</span>"
    }
  ]
}
}
```

If a request fails

<error code>:

- **404** Project does not exist or you do not have permission to view it
- **500** Invalid field name specified, the error text contains the invalid field name.

```
HTTP/1.1 <response error code>
```

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Create a new project

Summary

Create a new Swarm project.

POST /api/v10/projects

Description

Create a new Swarm project.

Parameters

Parameter	Description	Type	Parameter Type	Required
id	Project ID	string	path	Yes

Example usage

Create project Saturn:

```
curl -X POST -u "username:ticket" -H "Content-Type: application/json" -d "@mybodyfilename.txt" "https://my-swarm-host/api/v10/projects"
```

The "mybodyfilename.txt" file contains details for project Saturn:

```
{
  "id": "saturn",
  "name": "Saturn",
  "description": "Project Saturn is our moonshot project.",
  "jobView": "project=swarm",
  "members": ["bruno", "alice", "bob"],
  "defaults": {
    "reviewers": ["bob", "alice"]
  },
  "tests": {
    "enabled": false
  },
  "branches": [
    {
      "id": "main",
      "name": "main",
      "paths": ["/depot/saturn/..."],
      "defaults": {
        "reviewers": ["bob"]
      }
    }
  ]
}
```

Swarm responds with:

```
{
  "error": null,
  "messages": [],
  "data": {
    "projects": [
```

```
{
  "id": "saturn",
  "name": "Saturn",
  "defaults": {
    "reviewers": {
      "bob",
      "alice"
    }
  },
  "description": "Project Saturn is our moonshot project.",
  "members": [
    "bruno"
  ],
  "subgroups": [],
  "owners": [],
  "branches": [
    {
      "id": "main",
      "name": "main",
      "paths": [
        "//depot/saturn/..."
      ],
      "defaults": {
        "reviewers": {
          "Aruna_Gupta": []
        }
      },
      "moderators": [],
      "moderators-groups": [],
      "workflow": null,
      "retainDefaultReviewers": false,
      "minimumUpVotes": null
    }
  ],
  "jobview": null,
```

```

    "emailFlags": {
      "change_email_project_users": "1",
      "review_email_project_members": "1"
    },
    "tests": {
      "enabled": false,
      "url": null,
      "postBody": null,
      "postFormat": "url"
    },
    "deploy": {
      "enabled": false,
      "url": null
    },
    "deleted": false,
    "private": false,
    "workflow": null,
    "retainDefaultReviewers": false,
    "minimumUpVotes": null
  }
]
}
}

```

If a request fails

<error code>:

400 Invalid data specified for the project or invalid references such as references to tests and workflows.

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }]
}

```



```
  }],  
  "data" : null  
}
```

Update a project

Summary

Update a Swarm project.

PATCH /api/v10/projects/{{PROJECTID}}

Description

Update details for a Swarm project.

Tip

If sharing is switched off for a workflow, any projects or branches that were associated with the workflow while it was shared will remain associated with it.

If you edit a project or branch associated with that unshared workflow, you will still see the name of the workflow in the **Workflow** field, even if you don't own that workflow. If you remove that workflow from the project/branch you will not be able to see the workflow in the **Workflow** dropdown list unless you own it.

Note

By default, any member of a project can edit the project's configuration. Administrators can configure Swarm to prevent changes to the project's name and branch definition(s). See ["Projects" on page 670](#) for details.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Project ID	string	path	Yes

Example usage

Update project Saturn:

```
curl -X PATCH -u "username:ticket" "https://my-swarm-host/api/v10/projects/saturn"
```

```
{"workflow": "1"}
```

Swarm responds with the details of the updated project.

```
{
  "error": null,
  "messages": [],
  "data": {
    "projects": [
      {
        "name": "Saturn",
        "defaults": {
          "reviewers": ["bob", "alice"]
        },
        "description": "",
        "members": [
          "bruno", "alice", "bob"
        ],
        "subgroups": [],
        "owners": [
          "bob"
        ],
        "branches": [
          {
            "id": "main",
            "name": "main",
            "paths": ["/depot/saturn/..."],
            "defaults": {
              "reviewers": ["bob"]
            }
          }
        ],
        "jobview": "",
        "emailFlags": {
          "change_email_project_users": "1",
          "review_email_project_members": "1"
        },
        "tests": {
          "enabled": false,

```

```
    "url": "",
    "postBody": "",
    "postFormat": "url"
  },
  "deploy": {
    "enabled": false,
    "url": ""
  },
  "deleted": true,
  "private": false,
  "workflow": 1,
  "retainDefaultReviewers": false,
  "minimumUpVotes": null,
  "id": "saturn"
}
]
```

If a request fails

<error code>:

404 Project does not exist or you do not have permission to view it

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Delete a project

Summary

Delete a Swarm project.

DEL /api/v10/projects/{{PROJECTID}}

Description

Delete a Swarm project.

Note

- Users with *super* or *admin* privileges in Helix Core Server can delete projects.
- Project owners can delete projects that they own.
- If a project has no owners, any member of the project can delete the project.

When you delete a project:

- The deleted project is removed from the Swarm project list on the dashboard, and from project searches.
- The deleted project is removed from the profile page of the project owners, members, moderators, and followers.
- The project name is removed from the **Project** column on the "[Reviews list](#)" on [page 428](#) page.
- Reviews that belong to the deleted project are not changed. The open and closed reviews remain accessible, their [review states](#), [comments](#), and [tasks](#) can be modified as normal.
- Reviews that belong to the deleted project, the project branch name link in the [review heading](#) is replaced with a link to the common depot location that contains the files included in the review.

Note

The deleted project name cannot be reused for a new project. This behavior is not case sensitive, this means that if you delete a project called **Project B** you cannot create a new project called **project b**.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Project ID	string	path	Yes

Example usage

Delete project Saturn:

```
curl -X DEL -u "username:ticket" "https://my-swarm-host/api/v10/projects/saturn"
```

```
{
  "id": "saturn",
  "name": "Saturn",
  "description": "Project Saturn is our moonshot project.",
  "jobView": "project=swarm",
  "members": ["bruno", "alice", "bob"],
  "defaults": {
    "reviewers": ["bob", "alice"]
  },
  "tests": {
    "enabled": false
  },
  "workflow": 12,
  "branches": [
    {
      "id": "main",
      "name": "main",
      "paths": ["//depot/saturn/..."],
      "defaults": {
        "reviewers": ["bob"]
      }
    }
  ]
}
```

Swarm responds with the details of the deleted project.

```
{
  "error": null,
  "messages": [],
  "data": {
    "projects": [
      {
```

```
"name": "Saturn",
"defaults": {
  "reviewers": ["bob", "alice"]
},
"description": "",
"members": [
  "bruno", "alice", "bob"
],
"subgroups": [],
"owners": [
  "bob"
],
"branches": [
  {
    "id": "main",
    "name": "main",
    "paths": ["/depot/saturn/..."],
    "defaults": {
      "reviewers": ["bob"]
    }
  }
],
"jobview": "",
"emailFlags": {
  "change_email_project_users": "1",
  "review_email_project_members": "1"
},
"tests": {
  "enabled": false,
  "url": "",
  "postBody": "",
  "postFormat": "url"
},
"deploy": {
  "enabled": false,
```

```
        "url": ""
      },
      "deleted": true,
      "private": false,
      "workflow": null,
      "retainDefaultReviewers": false,
      "minimumUpVotes": null,
      "id": "saturn"
    }
  ]
}
```

If a request fails

<error code>:

- **401** This operation is limited to project or group owners.
- **404** Project does not exist or you do not have permission to view it

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Undelete a Project

Summary

Undelete a Swarm project.

POST /api/v10/projects/{{PROJECTID}}/undelete

Description

Undelete a Swarm project.

Note

- Users with *super* or *admin* privileges in Helix Core Server can undelete projects.
- Project owners can undelete projects that they own.
- If a project has no owners, any member of the project can undelete the project.

Parameters

Parameter	Description	Type	Parameter Type	Required
id	Project ID	string	path	Yes

Example usage

Undelete project Saturn that has two branches:

```
curl -X POST -u "username:ticket" "https://my-swarm-host/api/v10/projects/saturn/undelete"
```

Swarm responds with the details of the undeleted project.

```
{
  "error": null,
  "messages": [],
  "data": {
    "projects": [
      {
        "name": "Saturn",
        "defaults": {
          "reviewers": []
        },
        "description": null,
        "members": [
          "mei"
        ],
        "subgroups": [],
        "owners": [],
```



```
"branches": [  
  {  
    "name": "branch1",  
    "paths": [  
      "//depot/paths/..."  
    ],  
    "id": "branch1",  
    "moderators": [],  
    "moderators-groups": [],  
    "defaults": {  
      "reviewers": []  
    },  
    "workflow": null,  
    "retainDefaultReviewers": false,  
    "minimumUpVotes": null  
  },  
  {  
    "name": "branch2",  
    "paths": [  
      "//gwt/..."  
    ],  
    "id": "branch2",  
    "moderators": [],  
    "moderators-groups": [],  
    "defaults": {  
      "reviewers": []  
    },  
    "workflow": null,  
    "retainDefaultReviewers": false,  
    "minimumUpVotes": null  
  }  
],  
"jobview": null,  
"emailFlags": {  
  "change_email_project_users": "1",
```

```

        "review_email_project_members": "1"
      },
      "deleted": false,
      "private": false,
      "workflow": null,
      "retainDefaultReviewers": false,
      "minimumUpVotes": null,
      "id": "saturn"
    }
  ]
}

```

If a request fails

<error code>:

- **401** This operation is limited to project or group owners.
- **404** Project does not exist or you do not have permission to view it.

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```

Reviews: Swarm reviews

Get a list of reviews

Summary

Get a list of reviews.

GET /api/v10/reviews

Description

List the reviews the user has permission to view.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>max</code>	Specify the maximum number of results to return as an integer greater than 0. If the <code>max</code> parameter is not specified, the 50 most recent reviews are returned.	integer	query	No
<code>project[]</code>	To limit the results to a list of reviews in a specific project or projects, specify the projects you want reviews returned for.	string or an array (of strings)	query	No

Parameter	Description	Type	Parameter Type	Required
after	<p>A review ID to seek to. Reviews up to and including the specified id are excluded from the results and do not count towards max. Useful for pagination. Commonly set to the lastSeen property from a previous query.</p> <p>Reviews are returned in the order they were created in, starting with the most recently created review.</p> <div>Note You cannot use the after and afterUpdated parameters in the same request, they are mutually exclusive.</div>	string	query	No

Parameter	Description	Type	Parameter Type	Required
<code>afterUpdated</code>	<p>Return reviews updated on the day before the date/time specified in the <code>afterUpdated</code> parameter. The parameter value must be set in seconds since epoch.</p> <p>The seconds since epoch value for the day reviews are returned for is included at the end of the response. Reviews are returned in the order they were updated in, most recent first.</p> <p>Used by the Swarm team to step backwards one day at a time, loading reviews in an iterative process by using the value returned in a response as the <code>afterUpdated</code> parameter value for the next request.</p> <p>For example:</p> <ul style="list-style-type: none"> ▪ Requested: <code>afterUpdated=1606233362</code> (Tuesday November 24th, 2020, 15:56:02) ▪ Response fetched for: any reviews updated on Monday November 23rd, 2020. <code>1606089600</code> (Monday, 23 November 2020 00:00:00) is added to the end of the response 	string	query	No
<div>Note</div>				

Parameter	Description	Type	Parameter Type	Required
	You cannot use the after and afterUpdated parameters in the same request, they are mutually exclusive.			
state[]	<p>To limit the results to a list of reviews in a specific state or states, specify the states you want reviews returned for.</p> <p>Valid review states are:</p> <ul style="list-style-type: none"> ▪ needsRevision review needs revision ▪ needsReview review needs review ▪ approved review is approved ▪ approved:isPending review is approved but not committed ▪ committed review is committed ▪ approved:commit review is approved and committed ▪ approved:notPending review is approved and committed ▪ rejected review is rejected ▪ archived review is archived 	string or an array (of strings)	query	No

Parameter	Description	Type	Parameter Type	Required
keywords	To limit the results to a list of reviews that contain a specified keyword in a specified review field, use the keywords parameter in conjunction with the keywordsFields parameter. Multiple keywordFields can be specified.	string	query	No but required if keywordsFields is used
keywordsFields []	<p>To limit the results to a list of reviews that contain a specified keyword in a specified review field, use the keywordsFields parameter in conjunction with the keywords parameter. Multiple keywordFields can be specified.</p> <p>Valid keywordFields indexed on a review are:</p> <ul style="list-style-type: none"> ▪ changes ▪ author ▪ participants ▪ hasReviewer ▪ description ▪ updated ▪ projects ▪ state ▪ testStatus ▪ pending ▪ groups ▪ id 	string or an array (of strings)	query	No but required if keywords is used

Parameter	Description	Type	Parameter Type	Required
fields[]	To limit the fields returned for the reviews, add the fields parameter to the request. Multiple fields can be specified.	string or an array (of strings)	query	No

Example usage

Get a list of reviews

```
curl -u "username:ticket" "https://myswarm-url/api/v10/reviews"
```

Swarm responds with the review entities:

Tip

By default, this is limited to the 50 most recent reviews. This can be changed by adding a **max** parameter value to the request.

```
HTTP/1.1 200 OK
```

```
{
  "error": null,
  "messages": [],
  "data": {
    "reviews": [{
      "id": 12345,
      "type": "default",
      "changes": [
        12344,
      ],
      "commits": [],
      "author": "bruno",
      "approvals": null,
      "participants": ["bruno"],
      "participantsData": {
        "bruno": []
      }
    }]
  }
}
```



```
    },
    "hasReviewer": 0,
    "description": "Review for my code change",
    "created": 1594265070,
    "updated": 1594266228,
    "projects": {
      "myproject": ["main"]
    },
    "state": "needsReview",
    "stateLabel": "Needs Review",
    "testStatus": "fail",
    "testDetails": [],
    "deployStatus": null,
    "deployDetails": [],
    "pending": true,
    "commitStatus": [],
    "groups": [],
    "complexity": {
      "files_modified": 1,
      "lines_added": 0,
      "lines_edited": 1,
      "lines_deleted": 0
    },
  },
  ...
  <other review ids formatted as above>
  ...
],
"totalCount": 1,
"lastSeen": 12345
}
```

Get a list of the 25 most recent reviews

```
curl -u "username:ticket" "https://myswarm-  
url/api/v10/reviews?max=25"
```

Swarm responds with the 25 most recent review entities:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "reviews": [{
      "id": 12345,
      "type": "default",
      "changes": [
        12344,
      ],
      "commits": [],
      "author": "bruno",
      "approvals": null,
      "participants": ["bruno"],
      "participantsData": {
        "bruno": []
      },
      "hasReviewer": 0,
      "description": "Review for my code change",
      "created": 1594265070,
      "updated": 1594266228,
      "projects": {
        "myproject": ["main"]
      },
      "state": "needsReview",
      "stateLabel": "Needs Review",
      "testStatus": "fail",
      "testDetails": [],
      "deployStatus": null,
      "deployDetails": [],
      "pending": true,
      "commitStatus": [],
      "groups": [],
```

```
    "complexity": {
      "files_modified": 1,
      "lines_added": 0,
      "lines_edited": 1,
      "lines_deleted": 0
    },
  },
  ...
  <other review ids formatted as above>
  ...
],
"totalCount": 1,
"lastSeen": 12345
}
}
```

Get a list of reviews for the myproject project

```
curl -u "username:ticket" "https://myswarm-  
url/api/v10/reviews?project=myproject"
```

Swarm responds with the review entities for the **myproject** project:

Tip

By default, this is limited to the 50 most recent reviews. This can be changed by adding a **max** parameter value to the request.

HTTP/1.1 200 OK

```
{
  "error": null,
  "messages": [],
  "data": {
    "reviews": [{
      "id": 12345,
      "type": "default",
      "changes": [
        12344,
      ],
    },
  ],
}
```

```

    "commits": [],
    "author": "bruno",
    "approvals": null,
    "participants": ["bruno"],
    "participantsData": {
      "bruno": []
    },
    "hasReviewer": 0,
    "description": "Review for my code change",
    "created": 1594265070,
    "updated": 1594266228,
    "projects": {
      "myproject": ["main"]
    },
    "state": "needsReview",
    "stateLabel": "Needs Review",
    "testStatus": "fail",
    "testDetails": [],
    "deployStatus": null,
    "deployDetails": [],
    "pending": true,
    "commitStatus": [],
    "groups": [],
    "complexity": {
      "files_modified": 1,
      "lines_added": 0,
      "lines_edited": 1,
      "lines_deleted": 0
    },
  },
  ...
  <other review ids formatted as above>
  ...
],
"totalCount": 1,

```

```
    "lastSeen": 12345
  }
}
```

Get a list of reviews for the myproject and gemini projects

```
curl -u "username:ticket" "https://myswarm-  
url/api/v10/reviews?project[]=myproject&project[]=gemini"
```

Swarm responds with the review entities for the **myproject** and **gemini** projects:

Tip

By default, this is limited to the 50 most recent reviews. This can be changed by adding a **max** parameter value to the request.

```
HTTP/1.1 200 OK
```

```
{
  "error": null,
  "messages": [],
  "data": {
    "reviews": [{
      "id": 12345,
      "type": "default",
      "changes": [
        12344,
      ],
      "commits": [],
      "author": "bruno",
      "approvals": null,
      "participants": ["bruno"],
      "participantsData": {
        "bruno": []
      },
      "hasReviewer": 0,
      "description": "Review for my code change",
      "created": 1594265070,
      "updated": 1594266228,
      "projects": {
```

```
    "myproject": ["main"]
  },
  "state": "needsReview",
  "stateLabel": "Needs Review",
  "testStatus": "fail",
  "testDetails": [],
  "deployStatus": null,
  "deployDetails": [],
  "pending": true,
  "commitStatus": [],
  "groups": [],
  "complexity": {
    "files_modified": 1,
    "lines_added": 0,
    "lines_edited": 1,
    "lines_deleted": 0
  },
},
  "id": 12356,
  "type": "default",
  "changes": [
    12350,
  ],
  "commits": [],
  "author": "jsmith",
  "approvals": null,
  "participants": ["jsmith"],
  "participantsData": {
    "jsmith": []
  },
  "hasReviewer": 0,
  "description": "Another one of my reviews",
  "created": 1594265070,
  "updated": 1594266228,
  "projects": {
```

```

        "gemini": ["main"]
    },
    "state": "needsReview",
    "stateLabel": "Needs Review",
    "testStatus": "pass",
    "testDetails": [],
    "deployStatus": null,
    "deployDetails": [],
    "pending": true,
    "commitStatus": [],
    "groups": [],
    "complexity": {
        "files_modified": 1,
        "lines_added": 0,
        "lines_edited": 1,
        "lines_deleted": 0
    },
    },
    ...
    <other review ids formatted as above>
    ...
],
"totalCount": 2,
"lastSeen": 12356
}
}

```

Get a list of reviews that have jam in the projects or description fields

```
curl -u "username:ticket" "https://myswarm-
url/api/v10/reviews?keywords=jam&keywordsFields
[]=projects&keywordsFields[]=description"
```

Swarm responds with the review entities with "jam" in either the **projects** or **description** fields:

Tip

By default, this is limited to the 50 most recent reviews. This can be changed by adding a **max** parameter value to the request.

HTTP/1.1 200 OK

```
{
  "error": null,
  "messages": [],
  "data": {
    "reviews": [{
      "id": 12340,
      "type": "default",
      "changes": [
        12339,
      ],
      "commits": [],
      "author": "bruno",
      "approvals": null,
      "participants": ["bruno"],
      "participantsData": {
        "bruno": []
      },
      "hasReviewer": 0,
      "description": "Review for my code change",
      "created": 1594265070,
      "updated": 1594266228,
      "projects": {
        "jam": ["main"]
      },
      "state": "needsReview",
      "stateLabel": "Needs Review",
      "testStatus": "fail",
      "testDetails": [],
      "deployStatus": null,
      "deployDetails": [],
      "pending": true,
      "commitStatus": [],
      "groups": [],
```



```
    "complexity": {
      "files_modified": 1,
      "lines_added": 0,
      "lines_edited": 1,
      "lines_deleted": 0
    },
  },
  "id": 12360,
  "type": "default",
  "changes": [
    12359,
  ],
  "commits": [],
  "author": "jsmith",
  "approvals": null,
  "participants": ["jsmith"],
  "participantsData": {
    "jsmith": []
  },
  "hasReviewer": 0,
  "description": "Another one of my reviews and this one has jam
in the description",
  "created": 1594265070,
  "updated": 1594266228,
  "projects": {
    "gemini": ["main"]
  },
  "state": "needsReview",
  "stateLabel": "Needs Review",
  "testStatus": "pass",
  "testDetails": [],
  "deployStatus": null,
  "deployDetails": [],
  "pending": true,
  "commitStatus": [],
```

```

    "groups": [],
    "complexity": {
      "files_modified": 1,
      "lines_added": 0,
      "lines_edited": 1,
      "lines_deleted": 0
    },
  },
  ...
  <other review ids formatted as above>
  ...
],
"totalCount": 2,
"lastSeen": 12360
}
}

```

Get a review associated with a change

Get details of a review associated with a specific change 12345. The curl example is using the **&fields=id** parameter to reduce the size of the output.

```

curl https://mywarm-
url/api/v11/reviews?keywords=12345&keywordsFields
[]=changes&fields=id

```

Swarm responds with the review associated with the change 12345 which is 12346.

```

HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "reviews": [
      {
        "id": 12346
      }
    ],
    "totalCount": 1,

```

```
"lastSeen": 12344
}
}
```

Limit fields returned in a list of reviews

Get a list of reviews and limit the fields returned to `id`, `author`, and `pending`:

```
curl -u "username:ticket" "https://myswarm-  
url/api/v10/reviews?fields[]=id&fields[]=pending&fields[]=author"
```

Swarm responds with the review entities:

Tip

By default, this is limited to the 50 most recent reviews. This can be changed by adding a `max` parameter value to the request.

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "reviews": [{
      "id": 12345,
      "author": "bruno",
      "pending": true,
    },
    ...
    <other review ids formatted as above>
    ...
  ],
  "totalCount": 7,
  "lastSeen": 12352
}
```

Get a list of reviews created after review 12344

```
curl -u "username:ticket" "https://myswarm-url/api/v10/reviews?after=12344"
```

Swarm responds with the review entities:

```
HTTP/1.1 200 OK
```

```
{
  "error": null,
  "messages": [],
  "data": {
    "reviews": [{
      "id": 12345,
      "type": "default",
      "changes": [
        12345,
      ],
      "commits": [],
      "author": "bruno",
      "approvals": null,
      "participants": ["bruno"],
      "participantsData": {
        "bruno": []
      },
      "hasReviewer": 0,
      "description": "Review for my code change",
      "created": 1594265070,
      "updated": 1594266228,
      "projects": {
        "myproject": ["main"]
      },
      "state": "needsReview",
      "stateLabel": "Needs Review",
      "testStatus": "fail",
      "testDetails": [],
      "deployStatus": null,
      "deployDetails": [],
    }
  ]
}
```

```

    "pending": true,
    "commitStatus": [],
    "groups": [],
    "complexity": {
      "files_modified": 1,
      "lines_added": 0,
      "lines_edited": 1,
      "lines_deleted": 0
    },
  },
  ...
  <other review ids formatted as above>
  ...
],
"totalCount": 7,
"lastSeen": 12352
}
}

```

Get a list of reviews updated during the day before the set date

Request reviews updated on the day before **1606233362** (Tuesday November 24th, 2020, 15:56:02)

```
curl -u "username:ticket" "https://myswarm-  
url/api/v10/reviews?afterUpdated=1606233362"
```

Swarm return all reviews that were updated on the day before the date/time specified in your request. In this example, Swarm responds with any reviews that were updated on Monday November 23rd 2020. The seconds since epoch value Swarm used is added to the end of the response:

```

HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "reviews": [{
      "id": 12345,
      "type": "default",
      "changes": [

```

```

    12345,
  ],
  "commits": [],
  "author": "bruno",
  "approvals": null,
  "participants": ["bruno"],
  "participantsData": {
    "bruno": []
  },
  "hasReviewer": 0,
  "description": "Review for my code change",
  "created": 1594265070,
  "updated": 1606119044,
  "projects": {
    "myproject": ["main"]
  },
  "state": "needsReview",
  "stateLabel": "Needs Review",
  "testStatus": "fail",
  "testDetails": [],
  "deployStatus": null,
  "deployDetails": [],
  "pending": true,
  "commitStatus": [],
  "groups": [],
  "complexity": {
    "files_modified": 1,
    "lines_added": 0,
    "lines_edited": 1,
    "lines_deleted": 0
  },
},
...
<other review ids formatted as above>
...
```

```
    ],
    "totalCount": 25,
    "lastSeen": "0",
    "afterUpdated": "1606089600"
  }
}
```

Get a list of reviews that need review

```
curl -u "username:ticket" "https://myswarm-  
url/api/v10/reviews?state=needsReview"
```

Swarm responds with the review entities in the **needsReview** state:

Tip

By default, this is limited to the 50 most recent reviews. This can be changed by adding a **max** parameter value to the request.

```
HTTP/1.1 200 OK
```

```
{
  "error": null,
  "messages": [],
  "data": {
    "reviews": [{
      "id": 12345,
      "type": "default",
      "changes": [
        12344,
      ],
      "commits": [],
      "author": "bruno",
      "approvals": null,
      "participants": ["bruno"],
      "participantsData": {
        "bruno": []
      },
      "hasReviewer": 0,
      "description": "Review for my code change",
    ]
  }
}
```

```

    "created": 1594265070,
    "updated": 1594266228,
    "projects": {
      "myproject": ["main"]
    },
    "state": "needsReview",
    "stateLabel": "Needs Review",
    "testStatus": "fail",
    "testDetails": [],
    "deployStatus": null,
    "deployDetails": [],
    "pending": true,
    "commitStatus": [],
    "groups": [],
    "complexity": {
      "files_modified": 1,
      "lines_added": 0,
      "lines_edited": 1,
      "lines_deleted": 0
    },
  },
  ...
<other review ids formatted as above>
  ...
],
"totalCount": 1,
"lastSeen": 12345
}
}

```

Get a list of reviews that are open

Reviews are considered open when they are in any one of three states:

- **needsReview**
- **needsRevision**
- **approved:isPending**


```
curl -u "username:ticket" "https://myswarm-  
url/api/v10/reviews?state[]=approved:isPending&state  
[]=needsReview&state[]=needsRevision"
```

Swarm responds with the review entities for open reviews:

Tip

By default, this is limited to the 50 most recent reviews. This can be changed by adding a **max** parameter value to the request.

```
HTTP/1.1 200 OK
```

```
{  
  "error": null,  
  "messages": [],  
  "data": {  
    "reviews": [{  
      "id": 12345,  
      "type": "default",  
      "changes": [  
        12344,  
      ],  
      "commits": [],  
      "author": "bruno",  
      "approvals": null,  
      "participants": ["bruno"],  
      "participantsData": {  
        "bruno": []  
      },  
      "hasReviewer": 0,  
      "description": "Review for my code change",  
      "created": 1594265070,  
      "updated": 1594266228,  
      "projects": {  
        "myproject": ["main"]  
      },  
      "state": "needsReview",  
      "stateLabel": "Needs Review",  
    }  
  ]  
}
```

```
"testStatus": "fail",
"testDetails": [],
"deployStatus": null,
"deployDetails": [],
"pending": true,
"commitStatus": [],
"groups": [],
"complexity": {
  "files_modified": 1,
  "lines_added": 0,
  "lines_edited": 1,
  "lines_deleted": 0
},
},
"id": 12356,
"type": "default",
"changes": [
  12350,
],
"commits": [],
"author": "jsmith",
"approvals": null,
"participants": ["jsmith"],
"participantsData": {
  "jsmith": []
},
"hasReviewer": 0,
"description": "Another one of my reviews",
"created": 1594265070,
"updated": 1594266228,
"projects": {
  "gemini": ["main"]
},
"state": "needsRevision",
"stateLabel": "Needs Revision",
```

```
    "testStatus": "pass",
    "testDetails": [],
    "deployStatus": null,
    "deployDetails": [],
    "pending": true,
    "commitStatus": [],
    "groups": [],
    "complexity": {
      "files_modified": 1,
      "lines_added": 0,
      "lines_edited": 1,
      "lines_deleted": 0
    },
  },
  ...
  <other review ids formatted as above>
  ...
],
"totalCount": 2,
"lastSeen": 12356
}
```

If a request fails

<error code>:

- **400** one of the following:
 - the input specified must be an integer greater than 0
 - invalid **keywordsFields** field specified
- **401** you are not authorized to view reviews

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
```

```
"text" : "<error message>"
}},
"data" : null
}
```

Get review information

Summary

Get review information and metadata.

GET /api/v10/reviews/{id}

Description

Retrieve all of the information and metadata for a review.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Review ID	integer	path	Yes
<code>transitions</code>	To limit the results to a list of the state transitions that are allowed for the review, add the <code>transitions</code> parameter to the request.	query	path	No

Parameter	Description	Type	Parameter Type	Required
fields	To limit the fields returned for the review, add the fields parameter to the request. Multiple fields can be specified.	string or an array (of strings)	query	No

Example usage

Fetch information for review number 12345

```
curl -u "username:ticket" "https://myswarm-url/api/v10/reviews/12345"
```

Swarm responds with a review entity:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "reviews": [{
      "id": 12345,
      "type": "default",
      "changes": [1,3],
      "commits": [],
      "author": "normal-user",
      "approvals": null,
      "participants": ["normal-user"],
      "participantsData": {
        "normal-user": []
      },
      "hasReviewer": 0,
      "description": "Added a file\n",
      "created": 1576595749,
      "updated": 1576595754,
```

```

"projects": [],
"state": "needsReview",
"stateLabel": "Needs Review",
"transitions": {
  "needsRevision": "Needs Revision",
  "approved": "Approve",
  "approved:commit": "Approve and Commit",
  "rejected": "Reject",
  "archived": "Archive"
},
"testStatus": null,
"testDetails": [],
"deployStatus": null,
"deployDetails": [],
"pending": true,
"commitStatus": [],
"groups": [],
"versions": [
  {
    "difference": 1,
    "stream": "//jam/main",
    "streamSpecDifference": 0,
    "change": 2,
    "user": "normal-user",
    "time": 1576595752,
    "pending": true,
    "addChangeMode": "replace",
    "archiveChange": 3,
    "testRuns": []
  },
  {
    ...
    <other "versions" formatted as above>
    ...
  }
]

```

```
    ],
    "description-markdown": "<span class=\"first-line\">Added a
file</span>"
  }]
}
```

Fetch a list of the state transitions that are allowed for review 12345

```
curl -u "username:ticket" "https://myswarm-
url/api/v10/reviews/12345/transitions"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "transitions": {
      "needsRevision": "Needs Revision",
      "approved": "Approve",
      "approved:commit": "Approve and Commit",
      "rejected": "Reject",
      "archived": "Archive"
    }
  }
}
```

Limit fields returned for review number 12345

Limit the fields return for review 12345 to the review author and the committed status of the review.

```
curl -u "username:ticket" "https://myswarm-
url/api/v10/reviews/12345?fields[]=commitStatus&fields[]=author"
```

Swarm responds with a review entity:

```
HTTP/1.1 200 OK

{
```

```

"error": null,
"messages": [],
"data": {
  "reviews": [
    {
      "author": "normal-user",
      "commitStatus": []
    }
  ]
}

```

If a request fails

<error code>:

404 Review does not exist

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```

Get read by status for files in a review

Summary

Get the read by status for each file in a review for the authenticated user.

GET /api/v10/reviews/{id}/files/readby

Description

Get the read by status for each file in a review for the authenticated user. This is the status the user sets when they select **Mark file as read/Mark file as unread** for a file.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Example usage

Get a list of the read by status of each file in review 12345 for the authenticated user.

```
curl -u "bruno:ticket" "https://myswarm-url/api/v10/reviews/12345/files/readby"
```

With **bruno** is the authenticated user, Swarm responds with the following:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "filesInfo": [
      {
        "review": 12345,
        "depotFile": "//depot/jam/main/src/Build.com",
        "readBy": {
          "bruno": {
            "version": 1,
            "digest": "82D5161601D12DB46F184D3F0778A16D"
          }
        }
      }
    ]
  }
}
```

If a request fails

<error code>:

404 Review does not exist

```

HTTP/1.1 <response error code>

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Get transitions and items that are blocking approval for a review

Summary

Get a list of transitions and items that are blocking approval for a review.

GET /api/v10/reviews/{{REVIEWID}}/transitions

Description

Get a list of transitions and details about the items that are blocking approval for a review.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Pending moderator approvals required for a review

Provides a list of branch moderators that are blocking approval for a review. Depending on the configuration of `moderator_approval` configurable in the `SWARM_ROOT/data/config.php` file, the branch moderator approval is required.

Changing the state of any review associated with a moderated branch is restricted as follows:

- Only moderators can approve or reject the review. Moderators can also transition a review to any other state.

Important

Moderators prevent the automatic approval of reviews, for more information about automatically approving reviews using workflow rules see ["Workflow rules" on page 526](#).

Note

By default, when a review spans multiple branches that have different moderators, only one moderator from any one of the branches needs to approve the review.

Swarm can be configured to require that one moderator from each branch must approve the review, this is a global setting. If a moderator belongs to more than one of the branches spanned by the review, their approval will count for each of the branches they belong to. For instructions on how to configure moderator behavior, see ["Moderator behavior when a review spans multiple branches" on page 692](#).

- The review's author, when not a moderator, can change the review's state to **Needs review**, **Needs revision**, **Archived**, and can attach committed changelists.

Normally, the review's author cannot change the review's state to **Approved** or **Rejected** on moderated branches. However, authors that are also moderators have moderator privileges, and may approve or reject their own review.

When `disable_self_approve` is enabled, authors who are moderators (or even users with admin privileges) cannot approve their own reviews.
- Project members can change the review's state to **Needs review** or **Needs revision**, and can attach committed changelists. Project members cannot change the review's state to **Approved**, **Rejected**, or **Archived**.
- Users that are not project members, moderators, or the review's author cannot transition the review's state.
- For the review's author and project members, if a review is not in one of their permitted states, for example if the review's state is **Rejected**, they cannot transition the review to another state.

These restrictions have no effect on who can start a review.

Minimum Up votes required for a review

Provides the number of **Minimum Up votes** required for each project that you have access to. This includes the maximum number of **Minimum Up votes** required for all branches in a project blocking approval for a review.

For [private project](#), the maximum number of **Minimum Up votes** for all private projects and their branches that is blocking approval for a review is returned.

A review cannot be approved until all of the [Required reviewers](#) have voted up the review and the **Minimum up votes** specified has been satisfied.

- If a review spans projects/branches, the Minimum up votes for **each** of the projects and branches must be satisfied before you can approve the review.
- Required reviewers are included when up votes are counted.
- When **Count votes up from** is set to **Members** for a workflow associated with a project/branch, only the up votes of members of the project contribute to satisfying the **Minimum up votes** for a project/branch. For more information about the **Count votes up from** rule, see ["Workflow rules" on page 526](#).

Important

If the Workflow feature is disabled, all votes are counted not just votes from project members.

Important

If the **Minimum up votes** required is set higher than the number of reviewers that exist for a review, approval will be blocked for that review. This is true even if all the reviewers on the review have voted up the review.

Tests blocking approval for a review

Provides a list of workflow tests that have **Failed** or have an **In progress** status that is blocking approval for a review. For more information about a failed test, see [Tests](#).

Example usage

[Example for review 12345](#) shows a list of transitions and pending moderator approvals for two projects

GET /api/v10/reviews/12345/transitions

Swarm responds with:

```
HTTP/1.1 200 OK

{
  error: null,
  messages: [ ],
  data: {
    transitions: {
      "needsRevision": "Needs Revision",
```

```
    "rejected": "Reject",
    "archived": "Archive"
  },
  blocked: {
    approved: {
      moderators: {
        projects: {
          {
            id: "project1",
            name: "Project1",
            branches: {
              {
                id: "main",
                name: "Main Line",
                Users: {
                  {
                    id: "earl",
                    name: "Earl Ashby",
                    vote: 0,
                  },
                },
              },
            },
          },
        },
      },
    },
    {
      id: "project2",
      name: "Project2",
      branches: {
        {
          id: "main",
          name: "Main Line",
          Users: {
            {
              id: "quinn",
              name: "Quinn Cass",
```

```

    },
    {
      "id": "project1",
      "name": "Project1",
      "votes": {
        "approved": {
          "votes": {
            "projects": {
              "project1": {
                "id": "project1",
                "name": "Project1",
                "vote": 0,
              }
            }
          }
        }
      }
    }
  ],
  "messages": [ ],
  "error": null,
  "data": {
    "transitions": {
      "needsRevision": "Needs Revision",
      "rejected": "Reject",
      "archived": "Archive"
    },
    "blocked": {
      "approved": {
        "votes": {
          "projects": {
            "project1": {
              "id": "project1",
              "name": "Project1",
              "vote": 0,
            }
          }
        }
      }
    }
  }
}

```

Example for review 12345 shows a list of transitions and two Minimum Up votes that are blocking approval

GET /api/v10/reviews/12345/transitions

Swarm responds with:

```

HTTP/1.1 200 OK

{
  error: null,
  messages: [ ],
  data: {
    transitions: {
      "needsRevision": "Needs Revision",
      "rejected": "Reject",
      "archived": "Archive"
    },
    blocked: {
      approved: {
        votes: {
          projects: {
            project1: {
              id: "project1",
              name: "Project1",
              vote: 0,
            }
          }
        }
      }
    }
  }
}

```

```
        minimumUpVotes: 2
      }
    }
  }
}
```

Example for review 12345 shows a list of transitions and required Minimum Up votes that are blocking approval for a public project and a private project

GET /api/v10/reviews/12345/transitions

Swarm responds with:

HTTP/1.1 200 OK

```
{
  error: null,
  messages: [ ],
  data: {
    transitions: {
      "needsReview": "Needs Review",
      "rejected": "Reject",
      "archived": "Archive"
    },
    blocked: {
      approved: {
        votes: {
          projects: {
            jam: {
              id: "jam",
              name: "Jam",
              minimumUpVotes: 2
            }
          },
          privateProjects: 3
        }
      }
    }
  }
}
```

```

    }
  }
}
}

```

Example for review 12345 shows a list of transitions and three tests that are blocking approval

GET /api/v10/reviews/12345/transitions

Swarm responds with:

```

HTTP/1.1 200 OK

{
  error: null,
  messages: [ ],
  data: {
    transitions: {
      "needsReview": "Needs Review",
      "rejected": "Reject",
      "archived": "Archive"
    },
    blocked: {
      approved: {
        tests: ["1", "2", "3"],
      },
    },
  },
}

```

If a request fails

<error code>:

- **401** User must be authenticated
- **403** You do not have permission to view the review
- **404** Review does not exist


```
HTTP/1.1 <response error code>

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Get list of files that changed in a review

Summary

Get a list of files that changed between specified versions of a review.

GET /api/v10/reviews/{id}/files?from={x}&to={y}

Description

Get a list of files that changed between specified versions of a review.

The GET request returns the following:

- **root**: the common root depot path of the files that changed between the specified versions of the review.
- **limited**: Swarm limits the number of files presented for changelists based on the [max_changelist_files](#) setting for the Helix Core Server connection. By default, a maximum of **1000** files are returned:
 - **true**: the number of files in the review is \geq **max_changelist_files**. Some of the files changed in the specified review versions might not have been returned by this request.
 - **false**: the number of files in the review is $<$ **max_changelist_files**. All of the files changed in the specified review versions have been returned by this request.
- **For each file:**
 - **depotFile**: the filename and path of the file.
 - **action**: how the file changed, **Add**, **Edit**, or **Delete**.
 - **type**: the filetype of the file, **text**, **binary**, **symlink**, **unicode**, **utf8**, **utf16**, **apple**, or **resource**.

For more information on filetypes, see the [Base filetypes](#) section of the *Helix Core Command-Line (P4) Reference*.

- **rev**: the file revision number.
- **fileSize**: the file size in bytes.
- **digest**: the file digest.

Tip

- **Private projects**: when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions**: if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
id	Review ID	integer	path	Yes
from	Specify a value for from : <ul style="list-style-type: none"> ▪ 0 specifies the base version of the review. ▪ -1 specifies the current #head version of the review files in the Helix Core Server. ▪ If omitted, from defaults to 0 which is the base version of the review. ▪ Must not be greater than the latest version of the review. 	integer	query	No

Parameter	Description	Type	Parameter Type	Required
<code>to</code>	<p>Specify a value for <code>to</code>:</p> <ul style="list-style-type: none">■ If omitted, <code>to</code> defaults to the latest version of the review.■ Cannot be <code>0</code> or <code>-1</code>■ Must not be greater than the latest version of the review.	integer	query	No

Example usage

Get a list of files that changed between version 1 and 3 of review 12345

```
curl -u "username:ticket" "https://myswarm-url/api/v10/reviews/12345/files?from=1&to=3"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "root": "//jam/main/src",
    "limited": false,
    "files": [{
      "depotFile": "//jam/main/src/execvms.c",
      "action": "edit",
      "type": "text",
      "rev": "3",
      "fileSize": "3343",
      "digest": "82D5161601D12DB46F184D3F0778A16D"
    },
    {
      "depotFile": "//jam/main/src/jam.h",
      "action": "add",
```

```

    "type": "text",
    "rev": "2",
    "fileSize": "7364",
    "digest": "2E94B379F201AD02CF7E8EE33FB7DA99"
  }
]
}
}

```

If a request fails

<error code>:

- **400** invalid review version number requested
- **403** Insufficient permissions to access the review
- **404** Review does not exist

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```

Get activity for a review

Summary

Get a list of activities for the specified review id.

GET /api/v10/reviews/{id}/activity

Description

Get a list of activities for the specified review id, filtered by permissions.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
after	An activity ID to seek to. Activity entries up to and including the specified ID are excluded from the results and do not count towards max . Useful for pagination. Commonly set to the lastSeen property from a previous query.	integer	query	No
max	Maximum number of activity entries to return. This does not guarantee that max entries are returned. It does guarantee that the number of entries returned won't exceed max . Server-side filtering may exclude some activity entries for permissions reasons. If the max parameter is not included in the request, only the most recent 100 activities are returned.	integer	query	No
fields	An optional comma-separated list (or array) of fields to show. Omitting this parameter or passing an empty value shows all fields.	string	query	No

Example usage

Get a list of activities for review 1236

```
curl -u "username:ticket" "https://my-swarm-host/api/v10/reviews/1236/activity"
```

Swarm responds with:

Tip

By default, this is limited to the 100 most recent activities. This can be changed by adding a **max** parameter value to the request.

```
{
  "error": null,
  "messages": null,
  "data": {
    "activity": [
      {
        "id": 1646,
        "type": "review",
        "link": [
          "review",
          {
            "review": "1236",
            "version": 4
          }
        ],
        "user": null
        "action": "Automated tests reported",
        "target": "review 1236",
        "preposition": "failed tests for",
        "description": "Make a review.",
        "details": [ ],
        "topic": "reviews/1236",
        "depotFile": null,
        "time": 1620228301,
        "behalfOf": null,
        "projects": [ ],

```

```
    "followers": [
      "bruno"
      "bob"
    ],
    "streams": {
      "review-1236",
      "user-",
      "personal-",
      "personal-macy.winter",
      "personal-super"
    },
    "change": 1235
  },
  {
    Other ids formatted as above
  }
],
"totalCount": 100,
"lastSeen": 1646
}
```

Activity pagination

To get the second page of activities for the review limited to a max of 1 and showing specified fields only (based on the previous example):

```
curl -u "username:ticket" "https://myswarm.url/api/v10
/reviews/1236/activity?max=1&after=1646&fields=id,action,date,descr
iption,type"
```

Swarm responds with:

```
{
  "error": null,
  "messages": null,
  "data": {
    "activity": [
      {
        "id": 1645,
```

```

    "type": "review",
    "action": "Automated tests reported",
    "description": "Make a review.",
    "time": 1626945417
  }
],
"totalCount": 1,
"lastSeen": 1645
}
}

```

If a request fails

<error code>:

- **400** One of the following:
 - The **max** parameter value must be an integer greater than 0
 - The **after** parameter must be an integer greater than 0
- **403** Insufficient permissions to access the review
- **404** Review does not exist
- **500** Internal error, request failed

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```

Get a list of comments on a review

Summary

Get a list of comments on a review.

GET /api/v10/reviews/{id}/comments

Description

Get a list of comments on a review.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Review ID	integer	path	Yes

Example usage

Get a list of comments on a review

```
curl -u "username:ticket" "https://myswarm-url/api/v10/reviews/368/comments"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "84": {
          "id": 84,
          "topic": "reviews/368",
```

```
"context": {
  "attribute": "description",
  "review": 368,
  "version": false
},
"attachments": [],
"flags": [],
"taskState": "comment",
"likes": [],
"user": "steve.russell",
"time": 1594121651,
"updated": 1594121651,
"edited": null,
"body": "@paula.boyle - can you take a look?",
"readBy": []
},
"85": {
  "id": 85,
  "topic": "reviews/368",
  "context": {
    "file": "//projects/mercury/dev/src/api/client.cc",
    "leftLine": null,
    "rightLine": 7,
    "content": [
      " */",
      " ",
      " #include \"client.h\"",
      " ",
      "+#define OK \"OK\""
    ],
    "type": "text",
    "review": 368,
    "version": 1
  },
  "attachments": [],
```

```
        "flags": [],
        "taskState": "comment",
        "likes": [],
        "user": "paula.boyle",
        "time": 1594121935,
        "updated": 1594121935,
        "edited": null,
        "body": "Is this better? const char *OK=\"OK\";",
        "readBy": []
    },
    ...
    <other comments formatted as above>
    ...
}
}
]
}
```

If a request fails

<error code>:

- **403** Insufficient permissions to access the review
- **404** Review does not exist

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Get the number of delayed notifications on a review

Summary

Gets the number of [delayed comment notifications](#) for the authenticated user for a review.

GET /api/v10/reviews/{id}/notify

Description

Gets the number of [delayed comment notifications](#) for the authenticated user for a review.

Example usage

Get the number of delayed comment notifications for Bruno on review 12345:

```
curl -u "bruno:ticket" "https://myswarm-url/api/v10/reviews/12345/notify"
```

Swarm responds with the number of delayed comment notifications for Bruno on Review 12345:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "userDelayedNotificationCount": 3
  }
}
```

If a request fails

<error code>:

404 Review does not exist

```
HTTP/1.1 <response error code>

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }]
}
```

```
  }],  
  "data" : null  
}
```

Get Dashboard reviews for the authenticated user

Summary

Gets a list of reviews for the authenticated user's dashboard

GET /api/v10/reviews/dashboard

Description

Gets a list of reviews for the authenticated user's dashboard. For information about when reviews are displayed in the user dashboard, see ["My Dashboard" on page 306](#).

In addition to the normal review information, the following information is returned for the dashboard filter:

- **roles** lists the roles the authenticated user has for the review.
- **projectsForUser** lists *projectids* the authenticated user is an owner, member, or moderator of.

Example usage

Get the dashboard reviews for Bruno

```
curl -u "bruno:ticket" "https://myswarm-  
url/api/v10/reviews/dashboard"
```

Swarm responds with:

```
HTTP/1.1 200 OK  
  
{  
  "error": null,  
  "messages": [],  
  "data": {  
    "reviews": [{  
      "id": 12345,  
      "type": "default",  
      "changes": [  
        12344,
```

```
,
  "commits": [],
  "author": "bruno",
  "approvals": null,
  "participants": ["bruno"],
  "participantsData": {
    "bruno": []
  },
  "hasReviewer": 0,
  "description": "Review for my code change",
  "created": 1594265070,
  "updated": 1594266228,
  "projects": {
    "myproject": ["main"]
  },
  "state": "needsRevision",
  "stateLabel": "Needs Revision",
  "testStatus": "fail",
  "testDetails": [],
  "deployStatus": null,
  "deployDetails": [],
  "pending": true,
  "commitStatus": [],
  "groups": [],
  "complexity": {
    "files_modified": 1,
    "lines_added": 0,
    "lines_edited": 1,
    "lines_deleted": 0
  },
  "roles": [
    "author",
    "moderator"
  ],
},
```

```
    ...
    <other review ids formatted as above>
    ...
  ],
  "totalCount": 1,
  "lastSeen": 12345,
  "projectsForUser": [
    "jam"
    "mercury",
  ]
}
```

If a request fails

<error code>:

401 User must be authenticated

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Get a .zip archive of files in a review

Summary

Download a **.zip** archive of the files in the most recent version of a review.

GET /api/v10/reviews/{id}/archive

Description

Get a **.zip** archive of the files in the most recent version of a review. The **zip** command-line tool must be installed on the Swarm server to use this endpoint.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Example usage

Download a **.zip** archive of the files in review 12345.

GET /api/v10/reviews/12345/archive

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "archive": {
      "phase": "initializing",
      "error": null,
      "success": false,
      "progress": 0
    }
  }
}
```

If you repeat the request during the compression phase Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
```



```
"messages": [],
"data": {
  "archive": {
    "phase": "compressing",
    "error": null,
    "success": false,
    "progress": 50
  }
}
```

When **.zip** file has been created, Swarm downloads it.

If a request fails

<error code>:

- **401** User must be authenticated
- **404** Review does not exist
- **.Zip** command line tool not installed

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Refresh projects associated with a review

Summary

Refresh the list of projects review files are associated with.

POST /api/v10/reviews/{id}/refreshProjects

Description

Use the `refreshProjects` API endpoint to check if the review files are associated with any projects created or updated after the review was last updated. If any projects are found that are not already associated with the review, the review is linked to them.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Example usage

Refresh projects associated with review 12345

```
curl -X POST -u "username:ticket" "https://myswarm-url/api/v10/reviews/12345/refreshProjects"
```

Swarm responds with:

```
HTTP/1.1 200 OK
```

```
{
  "error": null,
  "messages": [],
  "data": {
    "reviews": [
      {
        "id": 12345,
        "type": "default",
        "changes": [
          12344
        ],
        "commits": [
          12344
        ],
        "author": "bruno",
```

```
"approvals": null,
"participants": [
  "bruno"
],
"participantsData": {
  "bruno": []
},
"hasReviewer": 0,
"description": "streamy",
"created": 1608644775,
"updated": 1612371881,
"projects": [
  "bluebook",
  "jam",
  "mercury"
],
"state": "needsRevision",
"stateLabel": "Needs Revision",
"testStatus": null,
"testDetails": [],
"deployStatus": null,
"deployDetails": [],
"pending": false,
"commitStatus": [],
"groups": [
  "swarm-project-my-precious",
  "swarm-project-jam",
  "Administrators"
],
"complexity": {
  "files_modified": 1,
  "lines_added": 0,
  "lines_edited": 0,
  "lines_deleted": 0
}
```

```

    ]
  }
}

```

If a request fails

<error code>:

- **401** User must be authenticated
- **404** Review does not exist

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```

Set vote for the authenticated user

Summary

Set the vote for the authenticated user to be up, down, or cleared

POST /api/v10/reviews/{id}/vote

Description

Set the vote for the authenticated user to be up, down, or cleared.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).

- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Review ID	integer	path	Yes
<code>vote</code>	Specify the <code>vote</code> value you want to set on the review. Valid vote values are <code>up</code> , <code>down</code> , and <code>clear</code> . If the requested vote value already exists on the latest version of the review, the vote is not changed.	string	body	Yes
<code>version</code>	Votes can only be set on the latest version of a review. The <code>version</code> parameter is used to check that you are voting on the latest version of the review. If the specified <code>version</code> is not the latest version of the review, the vote is not applied to the review and a 400 error is returned. If the <code>version</code> parameter is not included, the vote is applied to the latest version of the review.	integer	body	No

Example usage

Vote up version 1 of review 12345

```
curl -X POST -H "Content-Type: application/json" -u "username:ticket" -d "@mybodyfilename.txt" "https://myswarm-url/api/v10/reviews/12345/vote"
```

The "mybodyfilename.txt" file contains the authenticated user's `vote` and the `version` of the review they are voting on:

```
{
  "vote": "up"
  "version": "1"
}
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "vote": {
      "value": 1,
      "version": 1,
      "isStale": false
    }
  }
}
```

If a request fails

<error code>:

- **400** Version specified is not the latest version of the review
- **401** User must be authenticated
- **403** Insufficient permissions to access the review
- **404** Review does not exist
- **409** Invalid vote specified, the response error message lists valid votes.

```
HTTP/1.1 <response error code>

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
}
```

```
"data" : null
}
```

Change review state

Summary

Change the state of a review.

POST /api/v10/reviews/{id}/transitions

Description

Change the state of a review. The state that a review can be changed to depends on a number of factors including its current state. You can check that the state change you want to make is allowed for the review by making a ["Get transitions and items that are blocking approval for a review" on page 1206](#) request first. For more information on review states and restrictions on changing them, see ["States" on page 485](#)

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Review ID	integer	path	Yes

Parameter	Description	Type	Parameter Type	Required
transition	<p>Set the state transition to make to the review with the transition parameter.</p> <p>Review state transitions:</p> <ul style="list-style-type: none"> ▪ needsRevision ▪ needsReview ▪ approved only available if the voting requirements for the review are satisfied ▪ committed only available for "Pre-commit model" on page 424 reviews that have been approved ▪ approved:commit only available for unapproved "Pre-commit model" on page 424 reviews when the voting requirements for the review are satisfied ▪ rejected ▪ archived 	string	body	Yes
jobs	<p>Optional, add Perforce jobs to the review when the state is changed with the jobs parameter. Multiple jobs can be added but they must be comma separated. For information on Perforce jobs, see "Jobs" on page 335.</p>	array	body	No

Parameter	Description	Type	Parameter Type	Required
fixedStatus	<p>Optional, only used if the jobs parameter is used. Specify the Perforce job status when the state is changed with the fixedStatus parameter. Options depend on the configuration of your jobspec fields, see "Jobs" on page 335.</p> <p>The approved:commit state change updates the Perforce job description, all other state changes add the text as a comment on the job.</p>	String	body	No
cleanup	<p>Only used when the review is committed, options are:</p> <ul style="list-style-type: none"> ▪ true Swarm will attempt to automatically clean up any changelists left behind after the review has been committed, including removing any shelved files. This option can be removed by an administrator, see "Review cleanup" on page 678 for details. ▪ false Swarm will not try to automatically cleanup any changelists left behind after the review is committed. This is the default if the cleanup parameter is not present. 	Boolean	body	No

Example usage

Change the state of review 12345 to **Approved and Committed**, add Perforce jobs "job000001" and "job000015", set job status to closed, and get Swarm to attempt to cleanup any changelist that are left behind.

```
curl -X POST -H "Content-Type: application/json" -u
"username:ticket" -d "@mybodyfilename.txt" "https://myswarm-
url/api/v10/reviews/12345/transitions"
```

The "mybodyfilename.txt" file contains the state you want to change the review to and any other parameters you want to add:

```
{
  "transition": "approved:commit",
  "jobs" : ["job000001","job000015"],
  "fixStatus" : "closed",
  "cleanup" : true
}
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "state": "approved:commit",
    "stateLabel": "Approved and Committed"
  }
}
```

If a request fails

<error code>:

- **400** Specified review state does not exist
- **401** Insufficient permissions to change review state
- **403** Insufficient permissions to access the review
- **404** Review does not exist
- **409** Specified state is not a valid transition for this review. API response message lists all valid transition states for the review

```
HTTP/1.1 <response error code>

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Change review author

Summary

Change review author.

PUT /api/v10/reviews/{id}/author

Description

Change the review author.

Important

By default you cannot change the author of a review, this option must be enabled by your Swarm administrator. See ["Allow author change" on page 687](#) for details.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Review ID	integer	path	Yes

Example usage

Change the author of review 12345

```
curl -X PUT -H "Content-Type: application/json" -u "username:ticket" -d "@mybodyfilename.txt" "https://myswarm-url/api/v10/reviews/12345/author"
```

The "mybodyfilename.txt" file contains the new author username:

```
{
  "author": "asmith"
}
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "author": "asmith"
  }
}
```

If a request fails

<error code>:

- **400** Author change not allowed, set `allow_author_change` to `true`
- **401** Insufficient permissions to change author
- **403** Insufficient permissions to access the review
- **404** Review does not exist

```
HTTP/1.1 <response error code>

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Replace review description

Summary

Replace review description.

PUT /api/v10/reviews/{id}/description

Description

Replace the review description with a new description.

Optional: also replace pending changelist description

Only available if you are editing the description of a pre-commit review and you are the original author of the changelist that created the review.

To also apply your review description changes to the original changelist description, specify `'updateOriginalChangelist' : true`.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Review ID	integer	path	Yes

Example usage

Replace the description for review 12345

```
curl -X PUT -H "Content-Type: application/json" -u
"username:ticket" -d "@mybodyfilename.txt" "https://myswarm-
url/api/v10/reviews/12345/description"
```

The "mybodyfilename.txt" file contains the new description:

```
{
  "description": "This is my replacement description text and it is
better than the original description."
}
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "description": "This is my replacement description text and it is
better than the original description.",
    "description-markdown": "<span class=\"first-line\">This is my
replacement description text and it is better than the original
description.</span>"
  }
}
```

If a request fails

<error code>:

- **401** Insufficient permissions to change description
- **404** Review does not exist

```
HTTP/1.1 <response error code>

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Replace review participants

Summary

Replace review participants.

PUT /api/v10/reviews/{id}/participants

Description

Replace the review participants with a new set of participants.

Note

- The text file used for your request must contain a complete list of all of the participants you want on the review (**users** and **groups**). The existing review participants are completely replaced by the participants in the text file. For example, if you already have some group participants on the review and you do not specify any **group** roles in the text file, all of the participant groups will be removed from the review.

- Default reviewers can be removed from a review but retained default reviewers cannot be removed from a review, see ["Default reviewers" on page 518](#) and ["Retain default reviewers" on page 519](#).

If your request tries to remove a retained default reviewer, your request will succeed but with a message in the response indicating that the retained default reviewer has not been removed.

- You cannot reduce the voting requirement of a retained default reviewer but you can increase it, see ["Retain default reviewers" on page 519](#).

If your request tries to reduce the voting requirement of a retained default reviewer, your request will succeed but with a message in the response indicating that the retained default reviewer has not been changed.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Review ID	integer	path	Yes

Example usage

Replace participants for review 12345

```
curl -X PUT -H "Content-Type: application/json" -u "username:ticket" -d "@mybodyfilename.txt" "https://myswarm-url/api/v10/reviews/12345/participants"
```

The "mybodyfilename.txt" file contains the new review participants and their voting requirements:

```
{
  "participants": {
    "users": {
      "asmith": {
        "required": "yes"
      },
      "jbloggs": {
        "required": "no"
      },
      "molsen": {
        "required": "no"
      }
    }
  }
}
```



```
    }
  },
  "groups": {
    "docs": {
      "required": "none"
    },
    "qa": {
      "required": "all"
    },
    "dev": {
      "required": "one"
    },
    "testers": {
      "required": "one"
    }
  }
}
```

Swarm responds with a list of all of the participants and their voting requirements:

```
HTTP/1.1 200 OK
```

```
{
  "error": null,
  "messages": {
    "combinedReviewers": [
      "'molsen' must be a reviewer with a minimum requirement of
'required'",
      "'testers' must be a reviewer with a minimum requirement of
'require all'"
    ]
  },
  "data": {
    "participants": {
      "users": {
        "asmith": {
```

```

        "required": "yes"
    },
    "jbloggs": {
        "required": "no"
    },
    "molsen": {
        "required": "yes",
        "minimumRequired": "yes"
    }
},
"groups": {
    "dev": {
        "required": "one"
    },
    "docs": {
        "required": "none"
    },
    "qa": {
        "required": "all"
    },
    "testers": {
        "required": "all",
        "minimumRequired": "all"
    }
}
}
}
}

```

If a request fails

<error code>:

- **400** Invalid participant data, for example: unknown participant id or voting requirement
- **401** Insufficient permissions to replace participants
- **403** Insufficient permissions to edit reviewers
- **404** Review does not exist

```
HTTP/1.1 <response error code>

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Append review participants

Summary

Append review participants.

POST /api/v10/reviews/{id}/participants

Description

Append more participants to the existing review participants.

Note

Only the users and groups specified in the request are appended to the review. For example, if you don't want to append any group participants to the review there is no need to include the **groups** role in the text file.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Review ID	integer	path	Yes

Example usage

Append participants to review 12345

```
curl -X POST -H "Content-Type: application/json" -u "username:ticket" -d "@mybodyfilename.txt" "https://myswarm-url/api/v10/reviews/12345/participants"
```

The "mybodyfilename.txt" file contains the new review participants and their voting requirements. In this example, `bjones` is an optional participant:

```
{
  "participants": {
    "users": {
      "bjones": []
    }
  }
}
```

Swarm responds with a list of all of the participants and their voting requirements:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "participants": {
      "users": {
        "asmith": {
          "required": "yes"
        },
        "jbloggs": {
          "required": "no"
        },

```

```
    "molsen": {
      "required": "yes",
      "minimumRequired": "yes"
    },
    "bjones": []
  },
  "groups": {
    "dev": {
      "required": "one"
    },
    "docs": {
      "required": "none"
    },
    "qa": {
      "required": "all"
    },
    "testers": {
      "required": "all",
      "minimumRequired": "all"
    }
  }
}
```

If a request fails

<error code>:

- **400** Invalid participant data, for example: unknown participant id or voting requirement
- **401** Insufficient permissions to append participants
- **403** Insufficient permissions to edit reviewers
- **404** Review does not exist

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
```

```

"messages": [{
  "code" : "<code string>",
  "text" : "<error message>"
}],
"data" : null
}

```

Add a comment to a review

Summary

Add a comment to a review.

POST /api/v10/reviews/{id}/comments

Description

Add a comment to a review. The user that makes the API request is set as comment author.

Parameters

Parameter	Description	Type	Parameter Type	Required
body	The text content of the comment.	string	form	Yes
taskState	<p>Optional task state of the comment. Valid values when adding a comment are comment (default) and open. This creates a plain comment or opens a task, respectively.</p> <p>If the taskState parameter is not in the request, the task state is set to comment.</p>	string	form	No

Parameter	Description	Type	Parameter Type	Required
notify	<p>Comment notifications can be delayed, see "Comment notification delay" on page 355. The notify parameter enables you to set when the notification for the comment is sent.</p> <p>Valid options are:</p> <ul style="list-style-type: none">▪ delayed (default): notifications for this comment are delayed by the time set in "Comment notification delay" on page 585▪ immediate: the notification for this comment is sent immediately▪ silent: no notification is sent for this comment	string	query	No

Parameter	Description	Type	Parameter Type	Required
context	<p>Used to add the comment to a file in the review.</p> <p>Valid fields are:</p> <ul style="list-style-type: none"> ▪ file mandatory unless attribute or comment are set: File to comment on. Valid only for changes and reviews topics. For example, <code>//depot/main/README.txt</code>. ▪ leftline optional, but if specified, you must also specify the rightline and content parameters. Integer: Left-side diff line number to attach the inline comment to. Valid only for changes and reviews topics. ▪ rightline optional, but if specified, you must also specify the leftline and content parameters. Integer: Right-side diff line number to attach the inline comment to. Valid only for changes and reviews topics. 	array	form	No

Parameter	Description	Type	Parameter Type	Required
	<ul style="list-style-type: none"> ■ content optional, but if specified, you must also specify the leftline and rightline parameters. Array of strings: Provide the content of the codeline the comment is on and the four preceding codelines. This is used to specify a short excerpt of context in case the lines being commented on change during the review. When set to null, Swarm makes an effort to build the content on its own. Because this involves file operations, it might be slow. ■ version optional, integer: With a reviews topic, this field specifies which version to attach the comment to. ■ attribute optional: Set to description to comment on the review description. ■ comment optional, integer: Set to the comment id this comment is replying to. 			

Example usage

Add a comment to a review

Add a comment to review 885, make the comment an open task, and send the comment notification immediately:

```
curl -X POST -H "Content-Type: application/json" -u "username:ticket" -d "@mybodyfilename.txt" "https://my-swarm-host/api/v10/reviews/885/comments?notify=immediate"
```

The "mybodyfilename.txt" file contains the authenticated user's comment and makes the comment an open task:

```
{
  "body": "This is another comment.",
  "taskState": "open"
}
```

Swarm responds the comment entity:

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 95,
        "topic": "reviews/885",
        "context": [],
        "attachments": [],
        "flags": [],
        "taskState": "open",
        "likes": [],
        "user": "bob"
        "time": 1620813784,
        "updated": 1620813784,
        "edited": null,
        "body": "This is another comment.",
        "readBy": [],
        "notify": "immediate"
      },
    ],
    "userDelayedNotificationCount": 1
  }
}
```

Add a comment to a file in a review

Add a comment to the end of the `//depot/main/README.txt` file on review 123:

```
curl -X POST -H "Content-Type: application/json" -u
"username:ticket" -d "@mybodyfilename.txt" "https://my-swarm-
host/api/v10/reviews/123/comments"
```

The "mybodyfilename.txt" file contains the authenticated user's comment for the file:

```
{
  "body": "This is a comment on a file.",
  "context": {
    "file": "//depot/main/README.txt"
  }
}
```

Swarm responds the comment entity:

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 105,
        "topic": "reviews/123",
        "context": {
          "file": "//depot/main/README.txt",
          "leftLine": null,
          "rightLine": null,
          "content": [],
          "version": 1,
          "change": 122,
          "review": 123,
          "md5": "6af0e07c2d96afacaf2da234242c3cbd",
          "name": "README.txt"
          "line": null
        },
        "attachments": [],
        "flags": [],
        "taskState": "open",
        "likes": [],
        "user": "bob"
      }
    ]
  }
}
```

```

    "time": 1620813784,
    "updated": 1620813784,
    "edited": null,
    "body": "This is a comment on a file.",
    "readBy": [],
    "notify": "delayed"
  },
],
"userDelayedNotificationCount": 1
}
}

```

Add a comment to a line in a file and include context for the comment

Tip

To populate the `content` field, you need to know the content of the line you are commenting on and the content of the four lines before the line you are commenting on

Add a comment to a line in a file on review 110 with the following information:

- Full filepath and filename of the file to comment on
- Left and right line numbers of the diff to comment on
- Text of the codeline you are commenting on and the four codelines before that to provide context for the comment
- Review version to add the comment to

```
curl -X POST -H "Content-Type: application/json" -u
"username:ticket" -d "@mybodyfilename.txt" "https://my-swarm-
host/api/v10/reviews/110/comments"
```

The "mybodyfilename.txt" file contains the details for the new comment:

```

{
  "body": "This is a comment on a line in a file with context.",
  "context": {
    "file": "//depot/main/README.txt"
    "rightLine": 54,
    "leftLine": null,
    "content": [
      "FORTRAN default = \"\" ;\n",

```

```
    "LIBDIR default = /boot/develop/libraries ;\n",
    "-LINK default = $(CC) ;\n",
    "+LINK default = mwld ;\n",
    "+LINKFLAGS default = \"\" ;\n"
  },
  "version": 3
}
```

Swarm responds the comment entity:

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 143,
        "topic": "reviews/110",
        "context": {
          "file": "//depot/main/README.txt"
          "leftLine": null,
          "rightLine": 54,
          "content": [
            "FORTRAN default = \"\" ;\n",
            "LIBDIR default = /boot/develop/libraries ;\n",
            "-LINK default = $(CC) ;\n",
            "+LINK default = mwld ;\n",
            "+LINKFLAGS default = \"\" ;\n"
          ],
          "version": 3,
          "change": 109,
          "review": 110,
          "md5": "6af0e07c2d96afacaf2da234242c3cbd",
          "name": "README.txt",
          "line": 54
        }
      }
    ]
  }
}
```

```

    "attachments": [],
    "flags": [],
    "taskState": "open",
    "likes": [],
    "user": "bob"
    "time": 1620813784,
    "updated": 1620813784,
    "edited": null,
    "body": "This is a comment on a line in a file with context.",
    "readBy": [],
    "notify": "delayed"
  },
],
"userDelayedNotificationCount": 1
}
}

```

Add a comment to a review description

Add a comment to the description of review 885 and delay the notification:

```
curl -X POST -H "Content-Type: application/json" -u
"username:ticket" -d "@mybodyfilename.txt" "https://my-swarm-
host/api/v10/reviews/885/comments?notify=delayed"
```

The "mybodyfilename.txt" file contains the authenticated user's comment and adds it to the review description:

```

{
  "body": "This is a comment on the description.",
  "context":{
    "attribute": "description"
  },
}

```

Swarm responds the comment entity:

```

{
  "error": null,
  "messages": [],
  "data": {

```

```
"comments": [  
  {  
    "id": 99,  
    "topic": "reviews/885",  
    "context": {  
      "attribute": "description"  
    },  
    "attachments": [],  
    "flags": [],  
    "taskState": "comment",  
    "likes": [],  
    "user": "bob"  
    "time": 1620813784,  
    "updated": 1620813784,  
    "edited": null,  
    "body": "This is a comment on the description.",  
    "readBy": [],  
    "notify": "delayed"  
  },  
],  
  "userDelayedNotificationCount": 1  
}
```

Reply to a comment

You have two options when replying to a comment. Although the URLs for the options are different, they both do the same job:

- ["Reply to a comment using POST /api/v10/reviews/{id}/comments" below](#)
- ["Reply to a comment using POST /api/v10/reviews/{id}/comments/{id}" on page 1261](#)

Reply to a comment using POST /api/v10/reviews/{id}/comments

Add a reply to comment id 99 in review 885:

```
curl -X POST -H "Content-Type: application/json" -u  
"username:ticket" -d "@mybodyfilename.txt" "https://my-swarm-  
host/api/v10/reviews/885/comments"
```

```
{
  "body": "This is a reply to a comment.",
  "context": {
    "comment": 99
  }
}
```

Swarm responds the comment entity:

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 109,
        "topic": "reviews/885",
        "context": {
          "comment": 99,
          "version": 2
        },
        "attachments": [],
        "flags": [],
        "taskState": "comment",
        "likes": [],
        "user": "bruno"
        "time": 1622712940,
        "updated": 1622712940,
        "edited": null,
        "body": "This is a reply to a comment.",
        "readBy": [],
        "notify": "delayed"
      }
    ],
    "userDelayedNotificationCount": 1
  }
}
```


Reply to a comment using POST /api/v10/reviews/{id}/comments/{id}

Add a reply to comment id 99 in review 885:

```
curl -X POST -H "Content-Type: application/json" -u
"username:ticket" -d "@mybodyfilename.txt" "https://my-swarm-
host/api/v10/reviews/885/comments/99"
```

```
{
  "body": "This is another reply to a comment."
}
```

Swarm responds the comment entity:

```
{
  "error": null,
  "messages": [],
  "data": {
    "comments": [
      {
        "id": 120,
        "topic": "reviews/885",
        "context": {
          "comment": 99,
          "version": 2
        },
        "attachments": [],
        "flags": [],
        "taskState": "comment",
        "likes": [],
        "user": "bruno"
        "time": 1622712940,
        "updated": 1622712940,
        "edited": null,
        "body": "This is another reply to a comment.",
        "readBy": [],
        "notify": "delayed"
      }
    ]
    "userDelayedNotificationCount": 1
  }
}
```

```
}
}
```

If a request fails

<error code>:

- **400** one of the following:
 - the **id** specified must be an integer greater than 0
 - invalid **context** field specified, valid fields are **file**, **leftline**, **rightline**, **content**, **version**, **attribute**, and **comment**
 - invalid **notify** value specified, value must be **delayed**, **immediate**, or **silent**
 - invalid **taskstate** value specified, value must be **open** or **comment**
- **403**: you are not authorized to create this comment
- **404** the review, change, job, review version, file line number, or file specified does not exist
- **409** invalid **topic_subject** specified, value must be **reviews**, **changes**, or **jobs**
- **500** internal error, request failed

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Append change to a pre-commit review

Summary

Append a changelist to the changelist in a pre-commit review.

POST /api/v10/reviews/{id}/appendchange

Description

Append a changelist to the changelist in a pre-commit review. Not available for post-commit reviews. For information about what happens when you append a changelist to a review, see ["Append a pending changelist to a review" on page 473](#).

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Review ID	integer	path	Yes
<code>change_id</code>	Changelist number	integer	body	Yes

Example usage

Append changelist 12456 to review 12345:

```
curl -X POST -H "Content-Type: application/json" -u "username:ticket" -d "@mybodyfilename.txt" "https://my-swarm-host/api/v10/reviews/12345/appendchange"
```

```
{
  "changeId": 12456
}
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "reviews": [{
```

```
"id": 12345,
"type": "default",
"changes": [12344,12346,12456],
"commits": [],
"author": "normal-user",
"approvals": null,
"participants": ["normal-user"],
"participantsData": {
  "normal-user": []
},
"hasReviewer": 0,
"description": "Added a file\n",
"created": 1576595749,
"updated": 1576595754,
"projects": [],
"state": "needsReview",
"stateLabel": "Needs Review",
"transitions": {
  "needsRevision": "Needs Revision",
  "approved": "Approve",
  "approved:commit": "Approve and Commit",
  "rejected": "Reject",
  "archived": "Archive"
},
"testStatus": null,
"testDetails": [],
"deployStatus": null,
"deployDetails": [],
"pending": true,
"commitStatus": [],
"groups": [],
"versions": [
  {
    "difference": 1,
    "stream": "//jam/main",
```

```
    "streamSpecDifference": 0,
    "change": 2,
    "user": "normal-user",
    "time": 1576595752,
    "pending": true,
    "addChangeMode": "append",
    "archiveChange": 3,
    "testRuns": []
  },
  {
    "description-markdown": "<span class=\"first-line\">Added a
file</span>"
  }
]
```

If a request fails

<error code>:

- **401** Insufficient permissions to add change
- **403** Insufficient permissions to access the review
- **404** Review or change does not exist

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Replace change on a review

Summary

Replace the changelist on a review.

POST /api/v10/reviews/{id}/replacewithchange

Description

Replace the changelist on a review. Can be used for pre-commit and post-commit reviews. For information about what happens when you replace a changelist on a review, see ["Replace review with a pending changelist" on page 474](#) and ["Replace review with a committed changelist" on page 475](#).

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Review ID	integer	path	Yes
<code>change_id</code>	Changelist number	integer	body	Yes

Example usage

Replace changelist on review 12345 with changelist 12556:

```
curl -X POST -H "Content-Type: application/json" -u "username:ticket" -d "@mybodyfilename.txt" "https://my-swarm-host/api/v10/reviews/12345/replacewithchange"
```

```
{
  "changeId": 12256
}
```

Swarm responds with:

HTTP/1.1 200 OK

```
{
  "error": null,
  "messages": [],
  "data": {
    "reviews": [{
      "id": 12345,
      "type": "default",
      "changes": [12556],
      "commits": [],
      "author": "normal-user",
      "approvals": null,
      "participants": ["normal-user"],
      "participantsData": {
        "normal-user": []
      },
      "hasReviewer": 0,
      "description": "Added a file\n",
      "created": 1576595749,
      "updated": 1576595754,
      "projects": [],
      "state": "needsReview",
      "stateLabel": "Needs Review",
      "transitions": {
        "needsRevision": "Needs Revision",
        "approved": "Approve",
        "approved:commit": "Approve and Commit",
        "rejected": "Reject",
        "archived": "Archive"
      },
      "testStatus": null,
      "testDetails": [],
      "deployStatus": null,
      "deployDetails": [],
```

```

    "pending": true,
    "commitStatus": [],
    "groups": [],
    "versions": [
      {
        "difference": 1,
        "stream": "//jam/main",
        "streamSpecDifference": 0,
        "change": 2,
        "user": "normal-user",
        "time": 1576595752,
        "pending": true,
        "addChangeMode": "replace",
        "archiveChange": 3,
        "testRuns": []
      },
      {
        "description-markdown": "<span class=\"first-line\">Added a
file</span>"
      }
    ]
  }
}

```

If a request fails

<error code>:

- **401** Insufficient permissions to add change
- **403** Insufficient permissions to access the review
- **404** Review or change does not exist

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }]
}

```



```
  }],  
  "data" : null  
}
```

Send all delayed comment notifications on a review

Summary

Send all of the [delayed comment notifications](#) for the authenticated user for a review.

POST /api/v10/reviews/{review_id}/notify

Description

Send all of the [delayed comment notifications](#) for the authenticated user for a review.

Example usage

Send all of the delayed comment notifications for Bruno on review 12345:

```
curl POST -u "bruno:ticket" "https://myswarm-  
url/api/v10/reviews/12345/notify"
```

Swarm responds with the number of delayed notifications sent:

```
HTTP/1.1 200 OK  
  
{  
  "error": null,  
  "messages": [],  
  "data": {  
    "userDelayedNotificationCount": 3  
  }  
}
```

If a request fails

<error code>:

404 Review does not exist

```
HTTP/1.1 <response error code>  
  
{
```

```

"error": <error code>,
"messages": [{
  "code" : "<code string>",
  "text" : "<error message>"
}],
"data" : null
}

```

Delete review participants

Summary

Delete review participants.

DELETE /api/v10/reviews/{id}/participants

Description

Delete review participants from a review.

Note

- Only the users and groups specified in the request are deleted from the review. For example, if you don't want to delete any group participants from the review there is no need to include the **groups** role in the text file.
- Default reviewers can be removed from a review but retained default reviewers cannot be removed from a review, see ["Default reviewers" on page 518](#) and ["Retain default reviewers" on page 519](#).

If your request tries to remove a retained default reviewer, your request will succeed but with a message in the response indicating that the retained default reviewer has not been removed.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Review ID	integer	path	Yes

Example usage

Delete participants from review 12345

```
curl -X DELETE -H "Content-Type: application/json" -u "username:ticket" -d "@mybodyfilename.txt" "https://myswarm-url/api/v10/reviews/12345/participants"
```

The "mybodyfilename.txt" file contains the participants you want to delete from the review:

```
{
  "participants": {
    "users": {
      "jbloggs": [],
      "molsen": []
    },
    "groups": {
      "docs": []
    }
  }
}
```

Swarm responds with a list of all of the participants and their voting requirements:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": {
    "combinedReviewers": [
      "'molsen' must be a reviewer with a minimum requirement of 'required'"
    ]
  },
}
```

```
"data": {
  "participants": {
    "users": {
      "asmith": {
        "required": "yes"
      },
      "molsen": {
        "required": "yes",
        "minimumRequired": "yes"
      },
      "bjones": []
    },
    "groups": {
      "dev": {
        "required": "one"
      },
      "qa": {
        "required": "all"
      },
      "testers": {
        "required": "all",
        "minimumRequired": "all"
      }
    }
  }
}
```

If a request fails

<error code>:

- **400** Unknown participant id
- **401** Insufficient permissions to delete participants
- **403** Insufficient permissions to edit reviewers
- **404** Review does not exist

```
HTTP/1.1 <response error code>

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Obliterate a review

Summary

Obliterate a review.

DELETE /api/v10/reviews/{id}

Description

Important

Obliterate must be used with care, the review and all of its associated metadata are permanently deleted. An obliterated review cannot be reinstated, not even by Perforce Support.

Note

- By default, you must be a user with *admin* or *super* user rights to obliterate a review.
- **Optional:** Swarm can be configured to allow users to obliterate reviews that they have authored. This can be configured by your Swarm administrator. See ["Allow author obliterate review" on page 688](#).

Obliterate is used to permanently delete reviews that have been created by mistake. For instance, if a review is associated with the wrong changelist, or a review contains sensitive information that should not be openly available.

For information on what happens to a review when it is obliterated, see ["When you obliterate a review" on page 669](#).

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Review ID	integer	path	Yes

Example usage**Obliterate review number 12345**

```
curl -X DELETE -u "username:ticket" "https://myswarm-url/api/v10/reviews/12345"
```

Swarm responds with a message informing you that it has successfully Obliterated the review:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [{
    "code": "review-obiterated",
    "text": "The review with id [12345] has been obliterated."
  }],
  "data": {
    "reviews": [{
      "id": 12345,
      "type": "default",
      "changes": [1, 3],
      "commits": [],
      "author": "normal-user",
```

```
"approvals": null,
"participants": ["normal-user"],
"participantsData": {
  "normal-user": []
},
"hasReviewer": 0,
"description": "Hello\n",
"created": 1576579928,
"updated": 1576579932,
"projects": [],
"state": "needsReview",
"stateLabel": "Needs Review",
"testStatus": null,
"testDetails": [],
"deployStatus": null,
"deployDetails": [],
"pending": true,
"commitStatus": [],
"groups": []
"versions": [
  {
    "difference": 1,
    "stream": "//jam/main",
    "streamSpecDifference": 0
    "change": 2,
    "user": "normal-user",
    "time": 1576595752,
    "pending": true,
    "addChangeMode": "replace",
    "archiveChange": 3
    "testRuns": []
  },
  {
    ...
    <other "versions" formatted as above>
```

```

    ...
  }
}
}
}
}

```

Obliterate review number 56789 (with insufficient permission)

```
curl -X DELETE -u "username:ticket" "https://myswarm-
url/api/v10/reviews/56789"
```

Swarm responds with a message informing you that you do not have permission to obliterate the review:

```
HTTP/1.1 200 OK

{
  "error":403,
  "messages":[{"
    "code": "forbidden-exception",
    "text": "Failed to Obliterate the review, you need to have admin
privileges or be the review author with 'allow_author_obliterate' set
to true."
  }]
  "data": null
}
```

If a request fails

<error code>":

- **403** Insufficient permissions to obliterate review
- **404** Review does not exist

```
HTTP/1.1 <response error code>
```

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",

```



```
"text" : "<error message>"
}],
"data" : null
}
```

Enable notification for a review participant

Summary

Enable notifications for a review participant.

DELETE /api/v10/reviews/{id}/participantNotifications

Description

Enable notifications for a review participant. Once notifications are enabled, when you become a review participant, by joining the review or being [@mentioned](#) in a comment or in the review's description, you receive notifications for any events associated with the review.

Parameters

Parameter	Description	Type	Parameter Type	Required
<u>id</u>	Review ID	integer	path	Yes

Example usage

```
curl -X DELETE -u "username:ticket" "https://myswarm-url/api/v10/reviews/12345/participantNotifications"
```

Swarm responds with a list of all the review participants.

```
HTTP/1.1 200 OK
```

```
{
  "error": null,
  "messages": [],
  "data": {
    "participantsData": {
      "Aruna_Gupta": [],
      "bruno": [],
      "mei": [],
```

```

    "Sunil_Patil": [],
    "swarm-group-developers": []
  }
}

```

If a request fails

<error code>":

- **401** Unauthorized
- **404** Review does not exist

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```

Disable notification for a review participant

Summary

Disable notifications for a review participant.

POST /api/v10/reviews/{id}/participantNotifications

Description

Disable notifications for a review participant.

Once notifications are disabled, you no longer receive notifications. However, if you are [@mentioned](#) in a subsequent review comment, you do receive a notification for that comment; regular notifications remain disabled. This approach ensures that you don't miss anything that other reviewers or the review author deems important.

Parameters

Parameter	Description	Type	Parameter Type	Required
id	Review ID	integer	path	Yes

Example usage

```
curl -X POST -u "username:ticket" "https://myswarm-url/api/v10/reviews/12345/participantNotifications"
```

Swarm responds with a list of all the review participants.

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "participantsData": {
      "Aruna_Gupta": [],
      "bruno": {
        "notificationsDisabled": true
      },
      "mei": [],
      "Sunil_Patil": [],
      "swarm-group-developers": []
    }
  }
}
```

If a request fails

<error code>:

- **401** Unauthorized
- **404** Review does not exist

```
HTTP/1.1 <response error code>

{
```

```

"error": <error code>,
"messages": [{
  "code" : "<code string>",
  "text" : "<error message>"
}],
"data" : null
}

```

Join a review

Summary

Join a review.

POST /api/v10/reviews/{id}/join

Description

Important

The following individuals can be added as a reviewer:

- Users with *admin* or *super* privileges.
- If the review is [moderated](#), the moderators.
- If the review is part of a project, but not moderated, all project members.
- If the review is not part of a project, any authenticated user.

An authenticated user can add themselves as a reviewer for a review.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>id</code>	Review ID	integer	path	Yes

Example usage

Add vera to review 12345:

```
curl -X POST -u "vera:ticket" "https://myswarm-
url/api/v10/reviews/12345/join"
```

Swarm responds with a list of all the review participants.

HTTP/1.1 200 OK

```
{
  "error": null,
  "messages": [],
  "data": {
    "participants": {
      "users": {
        "Aruna_Gupta": {
          "vote": {
            "value": 1,
            "version": 2,
            "isStale": true
          }
        },
        "bruno": [],
        "earl": {
          "vote": {
            "value": -1,
            "version": 1,
            "isStale": true
          }
        },
        "mei": [],
        "vera": []
      },
      "groups": {
        "staletest": []
      }
    }
  }
}
```

If a request fails

<error code>":

- **401** Unauthorized
- **404** Review does not exist

Leave a review

Summary

Leave a review.

DEL /api/v10/reviews/{id}/leave

Description

Tip

You cannot leave a review if you are a retained default reviewer on the review or you are a member of a group that is a reviewer on the review.

An authenticated user can remove themselves as a reviewer from a review.

Parameters

Parameter	Description	Type	Parameter Type	Required
id	Review ID	integer	path	Yes

Example usage

Remove vera from review 12345:

```
curl -X DEL -u "vera:ticket" "https://myswarm-url/api/v10/reviews/12345/leave"
```

Swarm responds with a list of all the review participants.

```
HTTP/1.1 200 OK
```

```
{
  "error": null,
  "messages": [],
  "data": {
    "participants": {
      "users": {
        "Aruna_Gupta": {
```

```
    "vote": {
      "value": 1,
      "version": 2,
      "isStale": true
    },
    "bruno": [],
    "earl": {
      "vote": {
        "value": -1,
        "version": 1,
        "isStale": true
      },
      "mei": []
    },
    "groups": {
      "staletest": []
    }
  }
}
```

If a request fails

<error code>":

- **401** Unauthorized
- **404** Review does not exist

Archive inactive reviews

Summary

Archive inactive reviews.

POST /api/v10/reviews/archiveInactive

Description

Archiving reviews not updated since a specific date.

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>notUpdatedSince</code>	Not updated since date. Requires the date to be in the format YYYY-mm-dd, for example 2023-06-06	string	body	Yes
<code>max</code>	Optional, maximum number of reviews.	integer	body	No
<code>description</code>	Optional, a description that is posted as a comment for archiving.	string	body	No

Example usage

Archive maximum 50 reviews inactive since 2023/06/06.

```
curl -X POST -H "Content-Type: application/json" -u "username:ticket" -d "@mybodyfilename.txt" "https://myswarm-url/api/v10/reviews/archiveInactive"
```


The "mybodyfilename.txt" file contains the not updated since date, the maximum number of reviews, and a description:

```
{
  "notUpdatedSince": "2023-06-06",
  "max": 50,
  "description": "This is the description"
}
```

Swarm responds with a list of all the inactive reviews that are archived.

HTTP/1.1 200 OK

```
{
  "error": null,
  "messages": [],
  "data": {
    "archivedReviews": [
      {
        "id": 12450,
        "type": "default",
        "changes": [
          12443,
          12451
        ],
        "commits": [],
        "author": "bruno",
        "approvals": null,
        "participants": [
          "bruno"
        ],
        "participantsData": {
          "bruno": []
        },
        "hasReviewer": 0,
        "description": "pre commit review creation via swarm UI
success",
        "created": 1678794209,
```

```
"updated": 1683615474,
"projects": {
  "project3": [
    "branch1",
    "branch2",
    "branch3"
  ]
},
"state": "archived",
"stateLabel": "Archived",
"testStatus": null,
"previousTestStatus": "",
"testDetails": [],
"deployStatus": null,
"deployDetails": [],
"pending": true,
"commitStatus": [],
"groups": [
  "Administrators",
  "swarm-project-mod-project1",
  "swarm-project-mod-project2",
  "swarm-project-mod-project3",
  "swarm-project-mod-project4",
  "swarm-project-project-1",
  "swarm-project-project-2",
  "swarm-project-project-3",
  "swarm-project-project1",
  "swarm-project-project2",
  "swarm-project-project3",
  "swarm-project-project4",
  "swarm-project-project88",
  "swarm-project-project99",
  "swarm-project-test-project"
],
"complexity": {
```

```
    "files_modified": 1,  
    "lines_added": 0,  
    "lines_edited": 2,  
    "lines_deleted": 0  
  }  
}  
]  
}
```

If a request fails

<error code>:

- 401 Unauthorized
- 400 Transition of the review is not allowed

Create a Review

Summary

Create a review.

POST /api/v10/reviews/

Description

Pass in a changelist ID to create a review. Optionally, you can also provide a description and a list of reviewers. For more information about reviewers, individual reviewers, and group reviewers, see ["Reviewers" on page 447](#).

Tip

- **Private projects:** when a project is made private, only the project owners, moderators, and members, plus users with *super*-level privileges in the Helix Core Server, can see the project, its activity streams, and ongoing reviews. API responses honor these private project restrictions. For more information about private project restrictions, see ["Private projects" on page 387](#).
- **Permissions:** if a public review contains files that the user does not have permission to view, the review is returned but the restricted files are not displayed.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>change</code>	Change ID to create a review from	integer	form	Yes
<code>description</code>	Description for the new review (defaults to change description)	string	form	No
<code>reviewers</code>	A list of reviewers for the new review	an array (of strings)	form	No
<code>requiredReviewers</code>	A list of required reviewers for the new review	an array (of strings)	form	No
<code>reviewerGroups</code>	A list of reviewer groups or required reviewer groups for the new review	an array (of strings)	form	No

When a group is a required reviewer, it can be set to operate in one of two ways:

- **Require all:** all members of the group must up-vote the review to allow the review to be approved. To do this, in the `reviewerGroups` parameter for a group, set `"required" : "true"`.
- **Require one:** at least one member of the group must up-vote the review to allow the review to be approved. If any member of the group down-votes the review, the review cannot be approved. To do this, in the `reviewerGroups` parameter for a group, set `"required" : "true"` with `"quorum" : "1"`.

Example usage

Create a review for changelist 12345.

```
curl -X POST -H "Content-Type: application/json" -u
"username:ticket" -d "@mybodyfilename.txt" "https://myswarm-
url/api/v10/reviews/"
```

The "mybodyfilename.txt" file contains the changelist ID and a description:

```
{
  "change": "12345",
  "description": "This is the review description."
  "reviewers": [
    "raj",
    "mei"
  ],
  "requiredReviewers": [
    "vera",
    "dai"
  ],
  "reviewerGroups": [
    {"name": "WebDesigners", "required": "true", "quorum": "1"},
    {"name": "Developers", "required": "true"},
    {"name": "Administrators"}
  ]
}
```

Swarm responds with the new review entity:

```
HTTP/1.1 200 OK
```

```
{
  "error": null,
  "messages": [],
  "data": {
    "review": [
      {
        "id": 12346,
        "type": "default",
        "changes": [
```

```
12345
],
"commits": [],
"author": "bruno",
"approvals": null,
"participants": [
  "bruno"
  "dai",
  "mei",
  "raj",
  "swarm-group-Administrators",
  "swarm-group-Developers",
  "swarm-group-WebDesigners",
  "vera"
],
"participantsData": {
  "bruno": []
  "dai": {
    "required": true
  },
  "mei": [],
  "raj": [],
  "swarm-group-Administrators": [],
  "swarm-group-Developers": {
    "required": true
  },
  "swarm-group-WebDesigners": {
    "required": "1"
  },
  "vera": {
    "required": true
  }
},
"hasReviewer": 0,
"description": "This is the review description.",
```

```
    "created": 1689244995,
    "updated": 1689244995,
    "projects": [],
    "state": "needsReview",
    "stateLabel": "Needs Review",
    "testStatus": null,
    "previousTestStatus": "",
    "testDetails": [],
    "deployStatus": null,
    "deployDetails": [],
    "pending": true,
    "commitStatus": [],
    "groups": [],
    "complexity": null,
    "versions": []
  }
]
}
```

If a request fails

<error code>:

- **400** The input specified must be an integer greater than 0
- **401** Unauthorized
- **404** Cannot fetch change. Record does not exist.
- **403** You don't have permission to access the specified change

Search: Swarm search

Search Redis cache

Summary

Search the Redis cache for users, groups, projects, files, and file content.

GET /api/v10/search?term=<term>&context=user,group,project,filePath,fileContent

Description

Search the Redis cache for:

- *userids* and full names
- *groupids* and group names
- *projectids* and project names
- file path
- file content (only if Helix Search is configured for Swarm, see ["Search" on page 696](#))

Results are returned in alphabetical order and the search is case insensitive.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>term</code>	Specifies the characters to search for. Searches are case insensitive.	string	path	Yes

Parameter	Description	Type	Parameter Type	Required
context	Limit the search results, valid context values are user , group , project , filePath , and fileContent (see note below). Separate context values with a comma (,).	string	path	Yes

Note

If Helix Search is configured for Swarm, see ["Search" on page 696](#):

- The **filePath** search is carried out by Helix Search and not Swarm. This improves the performance of the search request.
- The **fileContent** search is available and carried out by Helix Search.

Parameter	Description	Type	Parameter Type	Required
<code>starts_with_only</code>	<p>Options are:</p> <ul style="list-style-type: none">■ <code>true</code>: searches for users/groups that start with the search <code>term</code>. This results in a faster search than when this parameter is set to <code>false</code>.■ <code>false</code>: searches for the <code>term</code> anywhere in users/groups. This is the search used if <code>starts_with_only</code> is not included in the request.	boolean	path	No

Parameter	Description	Type	Parameter Type	Required
<code>limit</code>	<p>Limit the combined maximum number of search results returned for users and groups. Search results are balanced across the contexts requested and will never return more results than the specified <code>limit</code>.</p> <p>Examples: based on a search that results in 8 user matches and 30 group matches before the <code>limit</code> is applied:</p> <ul style="list-style-type: none"> ▪ <code>limit=7</code> and <code>context=user</code>: search result is limited to 7 users. ▪ <code>limit=6</code> and <code>context=user,group</code>: search result is limited to 3 users and 3 groups. An even limit value splits the search result limit evenly between the specified contexts. ▪ <code>limit=7</code> and <code>context=user,group</code>: search result is limited to 4 users and 3 groups. An odd limit value splits the search result limit unevenly between the specified contexts with the higher number applied to the first context in the list. ▪ <code>limit=20</code> and 	integer	path	No

Parameter	Description	Type	Parameter Type	Required
	<p><code>context=user,group</code>, search result is limited to 8 users and 12 groups.</p> <p>Search results are not limited when the <code>limit</code> parameter is excluded from the request.</p>			
<code>ignoreExcludeList</code>	<p>Options are:</p> <ul style="list-style-type: none"> ▪ <code>true</code>: the <code>user_exclude_list</code> and <code>group_exclude_list</code> filters are ignored and are not applied to the search results. ▪ <code>false</code>: applies the <code>user_exclude_list</code> and <code>group_exclude_list</code> filters to the search results. This is the search used if <code>ignoreExcludeList</code> is not included in the request. 	boolean	path	No

Example usage

Search users, groups, and projects starting with "jo" with a limit of 4 results returned

Search users and groups starting with "jo" with a limit of 4 results returned.

```
curl -u "username:ticket" "https://myswarm-url/api/v10/search?starts_with_
only=true&term=jo&context=user,group,project&limit=4"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
```

```
"messages": [],
"data": {
  "user": {
    "results": [
      {
        "id": "joakly",
        "name": "Jane Oakly"
      },
      {
        "id": "jsmith",
        "name": "John Smith"
      }
    ]
  },
  "group": {
    "results": [
      {
        "id": "jogd",
        "name": "JOGD"
      }
    ]
  },
  "project": {
    "results": [
      {
        "id": "johns-project",
        "name": "Johns Project"
      }
    ]
  },
  "limit": 4
}
```

Search groups for "admin" anywhere in group names and *groupids*

Search for "admin" anywhere in group names and *groupids* with no limit on the number of results returned by the search.

```
curl -u "username:ticket" "https://myswarm-url/api/v10/search?context=group&term=admin"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "group": {
      "results": [
        {
          "id": "swarm-admin",
          "name": "Swarm-admin"
        },
        {
          "id": "jplugin-admin",
          "name": "JPlugin-Admin"
        },
        {
          "id": "administrators",
          "name": "Administrators"
        },
        {
          "id": "sysadmins",
          "name": "Sysadmins"
        }
      ]
    }
  }
}
```

Search files for "myfile" anywhere in file path

Search for "myfile" anywhere in file path with no limit on the number of results returned by the search.

```
curl -u "username:ticket" "https://myswarm-url/api/v10/search?context=filePath&term=myfile"
```

The response returned depends on whether Helix Search is configured for Swarm or not:

- **If Helix Search is not configured for Swarm, Swarm responds with:**

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "filePath": {
      "filesCount": 6,
      "maxScore": 100,
      "results": [
        {
          "type": "file",
          "change": null,
          "depotFile": "//depot/thisIsMyFile.txt",
          "fileName": "thisIsMyFile.txt",
          "action": null,
          "fileType": null,
          "rev": null,
          "fileSize": false
        },
        ....
        <results for other files formatted as above>
        ....
      ],
    }
  }
},
```

- **If Helix Search is configured for Swarm, Swarm responds with:**

```

HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "filePath": {
      "filesCount": 6,
      "maxScore": 100,
      "results": [
        {
          "type": "revision",
          "change": 10032,
          "depotFile": "//depot/thisIsMyFile.txt",
          "fileName": "thisIsMyFile.txt",
          "action": "branch",
          "fileType": "text",
          "rev": 1,
          "fileSize": 1714
        },
        ....
        <results for other files formatted as above>
        ....
      ],
    }
  }
},

```

Search file content for "foobar" anywhere in the file content

Note

The file content search is only available if Helix Search is configured for Swarm, see ["Search" on page 696](#)

Search for "foobar" anywhere in file content with no limit on the number of results returned by the search.


```
curl -u "username:ticket" "https://myswarm-url/api/v10/search?context=fileContent&term=foobar"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "fileContent": {
      "filesCount": 6,
      "maxScore": null,
      "results": [
        {
          "type": "content",
          "change": 12250,
          "depotFile": "//depot/jim",
          "fileName": "jim",
          "action": 'edit',
          "fileType": 'text',
          "rev": 3,
          "fileSize": 703
        },
        ....
        <results for other files formatted as above>
        ....
      ],
    }
  }
},
```

If a request fails

<error code>:

422 Parameter not specified or invalid, the error message in the response contains the error details.

```
HTTP/1.1 <response error code>

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Specs: Spec fields

Get spec information

Summary

Get spec file fields.

GET /api/v10/specs/{spec}/ fields

Description

Get fields for a spec type.

Parameter	Description	Type	Parameter Type	Required
spec	<p>Spec type. Valid spec types and location in the Helix Core Server spec depot:</p> <ul style="list-style-type: none"> ▪ branch: //spec/branch/ ▪ change: the change spec is not stored in the Helix Core Server spec depot ▪ client: //spec/client/ ▪ depot: //spec/depot/ ▪ group: //spec/group/ ▪ job: //spec/job/ ▪ label: //spec/label/ ▪ license: //spec/license.p4s ▪ protect: //spec/protect.p4s ▪ spec: //spec/spec/ ▪ stream: //spec/stream/ ▪ triggers: //spec/triggers.p4s ▪ typemap: //spec/typemap.p4s ▪ user: //spec/user/ 	string	path	Yes

Parameter	Description	Type	Parameter Type	Required
fields	<p>Omitting the fields parameter or including the fields parameter with no fields specified, returns all of the fields in the spec including any custom fields.</p> <p>Limit the response to specific fields by including a list of comma-separated fields in the request. The fields available depend on the spec type requested. Invalid fields do not cause an error and are ignored.</p>	array of strings	path	No

Example usage

Fetch all fields for the job spec

Fetching all fields for the job spec

```
curl -u "username:ticket" "https://myswarm-url/api/v10/specs/job/fields"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "job": {
      "Job": {
        "code": "101",
        "dataType": "word",
        "displayLength": "32",
        "fieldType": "required"
      },
      "Status": {
        "code": "102",
        "dataType": "select",
```

```
    "displayLength": "10",
    "fieldType": "required",
    "options": [
      "open",
      "suspended",
      "fixed",
      "closed"
    ],
    "default": "open"
  },
  "Created_by": {
    "code": "103",
    "dataType": "word",
    "displayLength": "32",
    "fieldType": "once",
    "default": "$user"
  },
  "Description": {
    "code": "105",
    "dataType": "text",
    "displayLength": "0",
    "fieldType": "required",
    "default": "$blank"
  },
  "Date_modified": {
    "code": "130",
    "dataType": "date",
    "displayLength": "20",
    "fieldType": "always",
    "default": "$now"
  }
}
```

Limit results to job and status fields for the job spec

Limit the results to the job and status fields for the job spec.

```
curl -u "username:ticket" "https://myswarm-url/api/v10/specs/job/fields?fields=status,job"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "job": {
      "Job": {
        "code": "101",
        "dataType": "word",
        "displayLength": "32",
        "fieldType": "required"
      },
      "Status": {
        "code": "102",
        "dataType": "select",
        "displayLength": "10",
        "fieldType": "required",
        "options": [
          "open",
          "suspended",
          "fixed",
          "closed"
        ],
        "default": "open"
      }
    }
  }
}
```

Limit results to depot and type fields for the depot spec

Limit the results to the depot and type fields for the depot spec.

```
curl -u "username:ticket" "https://myswarm-url/api/v10/specs/depot/fields?fields=depot,type"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "depot": {
      "Depot": {
        "code": "251",
        "dataType": "word",
        "displayLength": "32",
        "fieldType": "key"
      },
      "Type": {
        "code": "255",
        "dataType": "word",
        "displayLength": "10",
        "fieldType": "required"
      }
    }
  }
}
```

Limit results to status and type fields for the change spec

Limit the results to the status and type fields for the change spec.

```
curl -u "username:ticket" "https://myswarm-url/api/v10/specs/change/fields?fields=status,type"
```

Swarm responds with:

```
HTTP/1.1 200 OK
```

```
{
  "error": null,
  "messages": [],
  "data": {
    "change": {
      "Status": {
        "code": "205",
        "dataType": "word",
        "displayLength": "10",
        "fieldType": "always",
        "order": "5",
        "position": "R"
      },
      "Type": {
        "code": "211",
        "dataType": "select",
        "displayLength": "10",
        "fieldType": "optional",
        "order": "6",
        "position": "L",
        "options": [
          "public",
          "restricted"
        ]
      }
    }
  }
}
```

If a request fails

<error code>:

- **400** Unknown spec type requested
- **401** Insufficient permissions to fetch spec fields

HTTP/1.1 <response error code>


```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Testdefinitions: Swarm Test definitions

Get a list of test definitions

Summary

Get a list of test definitions and test fields

GET /api/v10/testdefinitions

Description

Returns a list of test definitions in Swarm that are visible to the current user.

Tip

Shared test definitions only display the test **url**, **body**, **headers**, and associated fields if you are the test definition owner. By default, this private data is visible in key data for users with list-level access. To hide key data, see ["Hiding Swarm storage from regular users" on page 238](#).

Example usage

Get a list of Swarm test definitions:

```
curl -u "username:ticket" "https://my-swarm-
host/api/v10/testdefinitions"
```

Swarm responds with a list of all of its test definitions:

In the example response below, you own testdefinition **id1** but you do not own testdefinition **id2**. The test **url**, **body**, **headers**, and associated fields are not returned for **id2** because you do not own that testdefinition.

```
HTTP/1.1 200 OK
```

```
{
  "error": null,
  "messages": [],
  "data" : {
    testdefinitions: [
      {
        "id": 1,
        "name": "Main code test",
        "headers": {
          "BasicAuth": "YWxpY2U6QUJDRDEyMzQ1Njc4Cg=="
        },
        "encoding": "url",
        "body": "status={status}&change={change}&update={update}",
        "url": "http://jenkins_host:8080/job/{branches}-review-
test/review/build",
        "timeout": 20,
        "owners": [
          "allison"
        ],
        "shared": false,
        "description": "Test for main code",
        "iterate": true
      },
      {
        "id": 2,
        "name": "Doc tests",
        "owners": [
          "bruno"
        ],
        "shared": true,
        "description": "Tests to check that the docs build.",
        "iterate": false
      }
    ]
  }
}
```

```
}
```

If the request fails

<error code>:

501 Workflows must be enabled for testdefinitions. To enable workflows, see ["workflow" on page 752](#)

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Get test definition information

Summary

Get the details of a test definition

GET /api/v10/testdefinitions/{id}

Description

Returns the specified test definition if it is visible to the current user.

Tip

Shared test definitions only display the test **url**, **body**, **headers**, and associated fields if you are the test definition owner. By default, this private data is visible in key data for users with list-level access. To hide key data, see ["Hiding Swarm storage from regular users" on page 238](#).

Parameters

Parameter	Description	Type	Parameter Type	Required
id	Testdefinition ID	integer	path	Yes

Example usage

Get test definition 1:

```
curl -u "username:ticket" "https://my-swarm-host/api/v10/testdefinitions/1"
```

Swarm responds with details for the test definition:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    testdefinitions: [
      {
        "id": 1,
        "name": "Main code test",
        "headers": {
          "BasicAuth": "YWxpY2U6QUJDRDEyMzQ1Njc4Cg=="
        },
        "encoding": "url",
        "body": "status={status}&change={change}&update={update}",
        "url": "http://jenkins_host:8080/job/{branches}-review-
test/review/build",
        "timeout": 20,
        "owners": [
          "allison"
        ],
        "shared": false,
        "description": "Test for main code",
        "iterate": true
      }
    ]
  }
}
```

If the request fails

<error code>:

- **404** Test definition does not exist or you do not have permission to view it.
- **501** Workflows must be enabled for testdefinitions. To enable workflows, see ["workflow" on page 752](#)

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Create a test definition

Summary

Create a test definition

POST /api/v10/testdefinitions

Description

Create a test definition.

Parameters

Parameter	Description	Type	Parameter Type	Required
name	Provide a name for your test. This should be a unique name that is 1 to 32 characters in length.	string	body	Yes
description	Provide a description for the test.	string	body	No

Parameter	Description	Type	Parameter Type	Required
headers	Specify name=value pairs to pass to the test suite.	array of strings	body	No
encoding	<p>Specify the encoding method for the request body parameter:</p> <ul style="list-style-type: none"> ▪ url for URL Encoded: POST parameters are parsed into name=value pairs. This is the default. ▪ json for JSON Encoded: parameters are passed raw in the POST body. ▪ xml for XML Encoded: parameters are passed in XML format in the POST body. 	string	body	No (required if body parameter used)
body	Specify parameters that your test suite requires in the request body parameter. Encoding is set with the encoding parameter.	string	body	No

Parameter	Description	Type	Parameter Type	Required
<code>url</code>	<p>Specify the test request URL that will trigger your test suite using the <code>url</code> parameter. 1 to 1024 characters.</p> <p>Special arguments are also available to pass details from Swarm to your test suite, see "Pass special arguments to your test suite" on page 548.</p> <div> <p>Note Helix Plugin for Jenkins 1.10.11 and later: Swarm must send the parameters for the build to Jenkins as a POST request. To do this, specify the parameters in the <code>body</code> and specify <code>url</code> for <code>encoding</code>.</p> </div>	string	body	Yes
<code>timeout</code>	Specify a timeout in seconds. The timeout is used when Swarm connects to your test suite. If a timeout is not set, the http_client_options timeout setting is used. By default, this is 10 seconds.	integer	body	No
<code>owners</code>	Specify at least one owner for the test.	array of strings	body	Yes

Parameter	Description	Type	Parameter Type	Required
shared	<p>Specify if want other Swarm users to be able to view or use this test definition:</p> <ul style="list-style-type: none">▪ true: other Swarm users can use this test.▪ false or not or not specified: other Swarm users will not be able to see or use this test.	boolean	body	No

Parameter	Description	Type	Parameter Type	Required
<code>iterate</code>	<p>Specify if you want Swarm to create a separate test run for each branch and project the review is associated with:</p> <ul style="list-style-type: none">▪ <code>true</code>: creates a separate test run for each branch and project the review is associated with. The details generated by the <code>{projects}</code> and <code>{branches}</code> arguments in the URL or Body are used in the test name displayed in the "Tests" on page 451 dropdown on the review page.▪ <code>false</code> or not or not specified: only one test will be run for the review, even if multiple branches and projects are associated with the review.	boolean	body	No

Example usage

Create a test definition

```
curl -X POST -H "Content-Type: application/json" -u "username:ticket" -d "@mybodyfilename.txt" "https://myswarm.url/api/v10/testdefinitions"
```

The "mybodyfilename.txt" file contains the test definition details:

```
{
  "name": "My test def",
  "description": "My test definition description",
  "headers": {
    "BasicAuth": "user:password",
    "OtherHeader": "OtherValue"
  },
  "encoding": "json",
  "body": "{\\"test\\":\\"{test}\\\"}",
  "url": "http://my-test-host/api/v10/testruns/test?test={test}",
  "timeout": 10,
  "owners" : ["bruno"],
  "shared": true,
  "iterate": true
}
```

Swarm responds with the testdefinition details:

HTTP/1.1 200 OK

```
{
  "error": null,
  "messages": [],
  "data" : {
    "testdefinitions": [
      {
        "id": 3,
        "name": "My test def",
        "headers": {
          "BasicAuth": "user:password",
          "OtherHeader": "OtherValue"
        },
        "encoding": "json",
        "body": "{\\"test\\":\\"{test}\\\"}",
        "url": "http://my-test-host/api/v10/testruns/test?test={test}",

```

```
    "timeout": 10,
    "owners": [
      "bruno"
    ],
    "shared": true,
    "description": "My test definition description",
    "iterate": true
  }
]
}
```

If the request fails

<error code>:

400 one of the following:

- The test definition **name** parameter is required and must be unique.
- The test definition **url** and **owners** parameters are required and cannot be empty.
- The **encoding**, **shared**, **description**, and **iterate** parameters require a value and cannot be empty.

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Testruns: Swarm Test Integration

Get testrun details of a review version

Summary

Get the testrun details of a review version.

GET /api/v10/reviews/{reviewId}/testruns

Description

Used to get the testrun details for a review version.

Note

Any authenticated user can get test run details for a review version. The test run `uuid` is not included in the request response.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>reviewId</code>	Review ID	integer	path	Yes
<code>version</code>	An optional parameter that enables you to specify the review version you want to return the testrun details for. Omitting this parameter or passing an empty value shows testrun details for the latest version of the review.	string	query	No

Example of usage

Get details for all of the test runs for the latest version of review 12345

```
curl "https://myswarm.url/api/v10/reviews/12345/testruns"
```

Swarm responds with the testrun details:

```
HTTP/1.1 200 OK
```

```
{
```

```
"error": null,
"messages": [],
"data" : {
  "testruns" : [
    {
      "id": 706,
      "change": 12345,
      "version": 2,
      "test": "global1",
      "startTime": 1567895432,
      "completedTime": 1567895562,
      "status": "pass",
      "messages": [
        "Test completed successfully",
        "another message"
      ],
      "url": "http://my.jenkins.com/projectx/main/1224"
    },
    {
      <ids for other testruns of the review, formatted as above>
    }
  ]
}
```

Get details for all of the test runs for version 2 of review 12345

```
curl "https://myswarm.url/api/v10/reviews/12345/testruns?version=2"
```

Swarm responds with the testrun details:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "testruns" : [
```

```

{
  "id": 706,
  "change": 12345,
  "version": 2,
  "test": "global1",
  "startTime": 1567895432,
  "completedTime": 1567895562,
  "status": "pass",
  "messages": [
    "Test completed successfully",
    "another message"
  ],
  "url": "http://my.jenkins.com/projectx/main/1224"
},
{
  <ids for other testruns for version 2 of the review, formatted
as above>
}
]
"status" : "pass"
}
}

```

Tip

When you make a GET request for all of the testruns in a specific review version, Swarm includes a **"status"** for the review version as a whole in the response. This is shown after the individual testruns. This status is calculated from the **"status"** values of all of the individual testruns for that review version:

- If the test **"status"** for any of the testruns is **"running"** or if any **On Demand** tests have not been run, the overall result for that version is **"running"**.
- If no testruns are **"running"** and the test **"status"** of any of the testruns is **"fail"** the overall result for that version is **"fail"**.
- If the test **"status"** of all of the testruns is **"pass"** then the overall result for that version is **"pass"**.

If a request fails

```
HTTP/1.1 <response error code>

{
  "error": <high level description>,
  "messages": [{"code" : "<code string>", "msg" : "<error message>"}],
  "data" : null
}
```

Create a testrun for a review version

Summary

Create a testrun for a review

POST /api/v10/reviews/{reviewid}/testruns

Description

Create a testrun for a review version. After you have created a testrun for a version of a review, you can update the testrun details as the test progresses using the [PATCH](#) and [PUT](#) API endpoints or the `{pass}`, `{fail}` and `{update}` callback urls.

Important

Only *admin* users can create a test run for a review.

Parameters

Parameter	Description	Type	Parameter Type	Required	Restriction
<code>reviewid</code>	Review ID.	integer	path	Yes	
<code>version</code>	The version of the review the testrun is being run against.	integer	body	Yes	

Parameter	Description	Type	Parameter Type	Required	Restriction
test	Specifies the test used for the testrun and checks that it is a valid name.	string	body	Yes	1 to 32 characters
startTime	The time the testrun started, expressed as an epoch value.	integer	body	Yes	> 0
completedTime	The time the testrun completed, expressed as an epoch value.	integer	body	No	> 0
status	The status of the testrun, options are: pass, fail, and running	string	body	No	Valid options are pass, fail, and running

Parameter	Description	Type	Parameter Type	Required	Restriction
<code>messages</code>	An array of one or more messages for the testrun. They should be formatted in JSON in the body of the POST request. You can pass a maximum 10 messages, if you provide more than 10 messages only the first 10 are saved. Each message can contain a maximum of 80 characters, any messages with more than 80 characters will be automatically truncated.	array	body	No	
<code>url</code>	The CI system url for the testrun, this is the link to the CI that enables you to view the current testrun results.	string	body	No	1 to 1024 characters


```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "testruns" : [
      {
        "id": 706,
        "change": 12345,
        "version": 2,
        "test": "mytest",
        "startTime": 1567895432,
        "completedTime": null,
        "status": "running",
        "messages": [
          "Mytest running",
          "another message"
        ],
        "url": "http://my.jenkins.com/projectx/main/1224"
      }
    ]
  }
}
```

If the request fails

```
HTTP/1.1 <response error code>

{
  "error": <high level description>,
  "messages": [{"code" : "<code string>", "msg" : "<error message>"}],
  "data" : null
}
```

Create a testrun for a review version using a specified test

Summary

Create a testrun for a review using a specified test
POST /api/v10/reviews/{reviewid}/testruns/{test_id}

Description

Create a testrun for a review version using a specified test. Tests are configured from the Swarm **Tests** page, see ["Tests" on page 544](#). After you have created a testrun for a version of a review, you can update the testrun details as the test progresses using the [PATCH](#) and [PUT](#) API endpoints or the `{pass}`, `{fail}` and `{update}` callback urls.

Important
Only *admin* users can create a test run for a review.

Note
The `test_id` parameter specified must exist in Swarm, if it doesn't exist an error is generated when the API request is made.

Parameters

Parameter	Description	Type	Parameter Type	Required	Restriction
<code>reviewid</code>	Review ID.	integer	path	Yes	
<code>version</code>	The version of the review the testrun is being run against.	integer	body	Yes	

Parameter	Description	Type	Parameter Type	Required	Restriction
test_id	<p>Specifies the test_id used for the testrun and checks that it is a valid test id. If there is a test name in the body of the request, Swarm will ignore it.</p> <p>If the test_id parameter is omitted, Swarm uses the test name in the body of the request.</p>	string	path	Yes	1 to 32 characters
startTime	The time the testrun started, expressed as an epoch value.	integer	body	Yes	> 0
completedTime	The time the testrun completed, expressed as an epoch value.	integer	body	No	> 0
status	The status of the testrun, options are: pass, fail, and running.	string	body	No	Valid options are pass, fail, and running

Parameter	Description	Type	Parameter Type	Required	Restriction
messages	An array of one or more messages for the testrun. They should be formatted in JSON in the body of the POST request. You can pass a maximum 10 messages, if you provide more than 10 messages only the first 10 are saved. Each message can contain a maximum of 80 characters, any messages with more than 80 characters will be automatically truncated.	array	body	No	

Parameter	Description	Type	Parameter Type	Required	Restriction
<code>url</code>	The CI system url for the testrun, this is the link to the CI that enables you to view the current testrun results.	string	body	No	1 to 1024 characters
<code>uuid</code>	The uuid for the testrun. Used to enable non- <i>admin</i> users to update the testrun. If the requested uuid does not match the review an error message is returned.	string	body	No	Minimum 32 max 64 characters

Example usage

Append a testrun for a review using test id 5 for version 3 of review 12345

```
curl -X POST -H "Content-Type: application/json" -u "username:ticket" -d "@mybodyfilename.txt" "https://myswarm.url/api/v10/reviews/12345/testruns/5"
```

The "mybodyfilename.txt" file contains the testrun details:

```
{
  "change": 12345,
  "version": 3,
  "startTime": 1567895432,
  "status": "running",
```

```
"messages": [  
  "Smoke-test running",  
  "another message"  
],  
"url": "http://my.jenkins.com/projecty/main/1224",  
"uuid": "FAE4501C-E4BC-73E4-A11A-FF710601BC4F"  
}
```

Swarm responds with the testrun details:

HTTP/1.1 200 OK

```
{  
  "error": null,  
  "messages": [],  
  "data" : {  
    "testruns" : [  
      {  
        "id": 5,  
        "change": 12345,  
        "version": 3,  
        "test": "test01",  
        "startTime": 1567895432,  
        "completedTime": null,  
        "status": "running",  
        "messages": [  
          "Smoke-test running",  
          "another message"  
        ],  
        "url": "http://my.jenkins.com/projecty/main/1224"  
      }  
    ]  
  }  
}
```


If the request fails

```
HTTP/1.1 <response error code>

{
  "error": <high level description>,
  "messages": [{"code" : "<code string>", "msg" : "<error message>"}],
  "data" : null
}
```

Update details for a testrun - PATCH

Summary

Update details for a testrun

PATCH /api/v10/reviews/{reviewid}/testruns/{id}

Description

Update the details for a testrun.

Important

Only *admin* users can update testrun details for a review.

Parameters

Parameter	Description	Type	Parameter Type	Required	Restriction
Id	Testrun ID, identifies a specific testrun for the review. Automatically generated by Swarm when the testrun is created.	integer	path	Yes	
reviewid	Review ID.	integer	path	Yes	

Parameter	Description	Type	Parameter Type	Required	Restriction
version	The version of the review the testrun is being run against.	integer	body	No	
test	Specifies the test used for the testrun and checks that it is a valid name.	string	body	No	1 to 32 characters
startTime	The time the testrun started, expressed as an epoch value.	integer	body	No	> 0
completedTime	The time the testrun completed, expressed as an epoch value.	integer	body	No	> 0
status	The status of the testrun, options are: pass, fail, and running	string	body	No	Valid options are pass, fail, and running

Parameter	Description	Type	Parameter Type	Required	Restriction
<code>messages</code>	An array of one or more messages for the testrun. They should be formatted in JSON in the body of the POST request. You can pass a maximum 10 messages, if you provide more than 10 messages only the first 10 are saved. Each message can contain a maximum of 80 characters, any messages with more than 80 characters will be automatically truncated.	array	body	No	
<code>url</code>	The CI system url for the testrun, this is the link to the CI that enables you to view the current testrun results.	string	body	No	1 to 1024 characters

Parameter	Description	Type	Parameter Type	Required	Restriction
uuid	The uuid for the testrun. Used to enable non- <i>admin</i> users to update the testrun. If the requested uuid does not match the review an error message is returned.	string	body	No	Minimum 32 max 64 characters

Example usage

Update testrun 706 for review 12345

```
curl -X PATCH -H "Content-Type: application/json" -u "username:ticket" -d "@mybodyfilename.txt" "https://myswarm.url/api/v10/reviews/12345/testruns/706"
```

The "mybodyfilename.txt" file contains the testrun details you want to update.

```
{
  "completedTime": "1567895562",
  "status": "fail",
  "messages": [
    "Test has failed",
    "yet another message"
  ]
}
```

Swarm responds with the testrun details:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
```

```
"testruns" : [
  {
    "id": 706,
    "change": 12345,
    "version": 2,
    "test": "mytest",
    "startTime": 1567895432,
    "completedTime": 1567895562,
    "status": "fail",
    "messages": [
      "Test has failed",
      "yet another message"
    ],
    "url": "http://my.jenkins.com/projectx/main/1224"
  }
]
```

If the request fails

```
HTTP/1.1 <response error code>

{
  "error": <high level description>,
  "messages": [{"code" : "<code string>", "msg" : "<error message>"}],
  "data" : null
}
```

Update details for a testrun - PUT

Summary

Update details for a testrun

PUT /api/v10/reviews/{reviewid}/testruns/{id}

Description

Update the details for a testrun. All values must be provided in the request.

Important

Only *admin* users can update testrun details for a review.

Parameters

Parameter	Description	Type	Parameter Type	Required	Restriction
Id	Testrun ID, identifies a specific testrun for the review. Automatically generated by Swarm when the testrun is created.	integer	path	Yes	
reviewid	Review ID.	integer	path	Yes	
version	The version of the review the testrun is being run against.	integer	body	Yes	
test	Specifies the test used for the testrun and checks that it is a valid name.	string	body	Yes	1 to 32 characters
startTime	The time the testrun started, expressed as an epoch value.	integer	body	Yes	> 0

Parameter	Description	Type	Parameter Type	Required	Restriction
<code>completedTime</code>	The time the testrun completed, expressed as an epoch value.	integer	body	Yes	> 0
<code>status</code>	The status of the testrun, options are: pass, fail, and running	string	body	Yes	Valid options are pass, fail, and running

Parameter	Description	Type	Parameter Type	Required	Restriction
<code>messages</code>	An array of one or more messages for the testrun. They should be formatted in JSON in the body of the POST request. You can pass a maximum 10 messages, if you provide more than 10 messages only the first 10 are saved. Each message can contain a maximum of 80 characters, any messages with more than 80 characters will be automatically truncated.	array	body	Yes	
<code>url</code>	The CI system url for the testrun, this is the link to the CI that enables you to view the current testrun results.	string	body	Yes	1 to 1024 characters

Parameter	Description	Type	Parameter Type	Required	Restriction
uuid	The uuid for the testrun. Used to enable non- <i>admin</i> users to update the testrun. If the requested uuid does not match the review an error message is returned.	string	body	Yes	Minimum 32 max 64 characters

Example usage

Update testrun 706 for review 12345

```
curl -X PUT -H "Content-Type: application/json" -u "username:ticket" -d "@mybodyfilename.txt" "https://myswarm.url/api/v10/reviews/12345/testruns/706"
```

The "mybodyfilename.txt" file must contain all of the testrun details:

```
{
  "id": 706,
  "change": 12345,
  "version": 2,
  "test": "mytest",
  "startTime": 1567895432,
  "completedTime": 1567895562,
  "status": "fail",
  "messages": [
    "Test has failed",
    "and another message"
  ],
  "url": "http://my.jenkins.com/projectx/main/1224",
  "uuid": "FAE4501C-E4BC-73E4-A11A-FF710601BC3F"
}
```

Swarm responds with the testrun details:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data" : {
    "testruns" : [
      {
        "id": 706,
        "change": 12345,
        "version": 2,
        "test": "mytest",
        "startTime": 1567895432,
        "completedTime": 1567895562,
        "status": "fail",
        "messages": [
          "Test has failed",
          "and another message"
        ],
        "url": "http://my.jenkins.com/projectx/main/1224"
      }
    ]
  }
}
```

If the request fails

```
HTTP/1.1 <response error code>

{
  "error": <high level description>,
  "messages": [{"code" : "<code string>", "msg" : "<error message>"}],
  "data" : null
}
```

Users: Swarm users

Get a list of users

Summary

Get a list of users.

GET /api/v10/users

Description

Returns a list of users in Swarm.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>ids</code>	An optional single string or an array of users to display. Omitting this parameter or passing an empty value shows all users.	string	query	No
<code>ignoreExcludeList</code>	Determines if the list of users has the <code>user_exclude_list</code> filter applied or not. Add the parameter to ignore the <code>user_exclude_list</code> filter.	boolean	query	No

Example usage

Get a list of users

```
curl -u "username:ticket" "https://myswarm-url/api/v10/users"
```

Swarm responds with all users in Swarm:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "users": [
      {
        "id": "bruno",
        "type": "standard",
        "email": "bruno@perforce.com",
        "update": "2013/10/11 21:05:15",
        "access": "2021/01/21 16:55:38",
        "fullName": "Bruno First",
        "jobView": null,
        "password": null,
        "authMethod": "perforce",
        "reviews": []
      },
      ...
      <other users formatted as above>
      ...
    ]
  }
}
```

Fetch a specific user

To fetch details for bruno:

```
curl -u "username:ticket" "https://myswarm-url/api/v10/users?ids=bruno"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "users": [
      {
        "id": "bruno",
        "type": "standard",
        "email": "bruno@perforce.com",
        "update": "2013/10/11 21:05:15",
        "access": "2021/01/21 16:55:38",
        "fullName": "Bruno First",
        "jobView": null,
        "password": null,
        "authMethod": "perforce",
        "reviews": []
      }
    ]
  }
}
```

Fetch a number of specific users

To fetch details for bruno and alice:

```
curl -u "username:ticket" "https://myswarm-url/api/v10/users?ids[]=bruno&ids[]=alice"
```

Swarm responds with:

```
HTTP/1.1 200 OK

{
  "error": null,
```

```
"messages": [],
"data": {
  "users": [
    {
      "id": "bruno",
      "type": "standard",
      "email": "bruno@perforce.com",
      "update": "2013/10/11 21:05:15",
      "access": "2021/01/21 16:55:38",
      "fullName": "Bruno First",
      "jobView": null,
      "password": null,
      "authMethod": "perforce",
      "reviews": []
    },
    {
      "user": "alice",
      "type": "standard",
      "email": "asmith@perforce.com",
      "update": "2018/09/11 18:17:14",
      "access": "2020/12/24 16:55:38",
      "fullName": "Alice Smith",
      "jobView": null,
      "password": null,
      "authMethod": "perforce",
      "reviews": []
    }
  ]
}
```

If a request fails

<error code>:

401 Incorrect or missing credentials

```
HTTP/1.1 <response error code>
```

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Get information for a specific user

Summary

Get information for a specific user.

GET /api/v10/users/{id}

Description

Returns information for a specific user in Swarm.

Parameters

Parameter	Description	Type	Parameter Type	Required
id	User ID	string	path	Yes

Example usage

Get information for a specific user

To fetch information for bruno:

```
curl -u "username:ticket" "https://myswarm-url/api/v10/users/bruno"
```

Swarm responds with:

```
HTTP/1.1 200 OK
```

```
{
```

```

"error": null,
"messages": [],
"data": {
  "users": [
    {
      "id": "bruno",
      "type": "standard",
      "email": "bruno@perforce.com",
      "update": "2013/10/11 21:05:15",
      "access": "2021/01/21 16:55:38",
      "fullName": "Bruno First",
      "jobView": null,
      "password": null,
      "authMethod": "perforce",
      "reviews": []
    }
  ]
}

```

If a request fails

<error code>:

- **401** Incorrect or missing credentials
- **404** User does not exist

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```


Workflows : Swarm workflows

Get workflows

Summary

Gets workflows

GET /api/v10/workflows/

Description

Get all workflows

Parameters

Parameter	Description	Type	Parameter Type	Required
fields	An optional comma-separated list (or array) of fields to show for each workflow. Omitting this parameter or passing an empty value shows all fields.	string	query	No
noCache	If provided and has a value of 'true' a query will always be performed and the cache of workflows is ignored. Otherwise the cache will be used if it exists.	boolean	query	No

Usage example

Get a list of workflows

```
curl -u "username:ticket" "https://myswarm-url/api/v10/workflows"
```

JSON Response:

Tip

The global workflow has an `id` of 0.

HTTP/1.1 200 OK

```
{
  "error": null,
  "messages": [],
  "data": {
    "workflows": [
      {
        "on_submit": {
          "with_review": {
            "rule": "no_checking",
            "mode": "default"
          },
          "without_review": {
            "rule": "no_checking",
            "mode": "policy"
          }
        },
        "name": "Global Workflow",
        "description": "This is the Global workflow",
        "shared": "false",
        "owners": [
          "swarm"
        ],
        "end_rules": {
          "update": {
            "rule": "no_checking",
            "mode": "default"
          }
        },
        "auto_approve": {
          "rule": "never",
          "mode": "default"
        }
      }
    ]
  }
}
```

```
    },
    "counted_votes":{
      "rule": "anyone",
      "mode": "default"
    },
    "group_exclusions":{
      "rule": [ ],
      "mode": "policy"
    },
    "user_exclusions":{
      "rule": [ ],
      "mode": "policy"
    },
    "tests": [ ],
    "id": "0"
  },
  {
    "on_submit":{
      "with_review":{
        "rule": "no_checking",
        "mode": "inherit"
      },
      "without_review":{
        "rule": "no_checking",
        "mode": "inherit"
      }
    },
    "name": "myWorkflow 1",
    "description": "Another description",
    "shared": "true",
    "owners": [
      "user3",
      "user4"
    ],
    "end_rules":{
```

```

        "update":{
            "rule":"no_revision",
            "mode": "inherit"
        },
        "auto_approve":{
            "rule": "votes",
            "mode": "inherit"
        },
        "counted_votes": {
            "rule": "members",
            "mode": "inherit"
        },
        "tests": [],
        "id": "1"
    },
]
}

```

If a request fails

<error code>:

501 Workflows are not enabled. To enable workflows, see ["workflow " on page 752](#)

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}

```

Get a workflow by id

Summary

Gets a workflow by id

GET /api/v10/workflows/{id}

Description

Gets a workflow by id. The global workflow `id` is `0`.

Parameters

Parameter	Description	Type	Parameter Type	Required
<code>fields</code>	An optional comma-separated list (or array) of fields to show for each workflow. Omitting this parameter or passing an empty value shows all fields.	string	query	No

Example usage

Get a workflows by id

```
curl -u "username:ticket" "https://my-swarm-host/api/v10/workflows/1"
```

JSON Response:

```
HTTP/1.1 200 OK

{
  "error": null,
  "messages": [],
  "data": {
    "workflows": [
      {
        "on_submit": {
          "with_review": {
            "rule": "no_checking",
```

```

        "mode": "inherit"
      },
      "without_review":{
        "rule": "no_checking",
        "mode": "inherit"
      }
    },
    "name": "myWorkflow 1",
    "description": "Another description",
    "shared": "true",
    "owners": [
      "user3",
      "user4"
    ],
    "end_rules":{
      "update":{
        "rule":"no_revision",
        "mode": "inherit"
      }
    },
    "auto_approve":{
      "rule": "votes",
      "mode": "inherit"
    },
    "counted_votes": {
      "rule": "members",
      "mode": "inherit"
    },
    "tests": [],
    "id": "1"
  },
]
}

```

If a request fails

<error code>:

- **401** Insufficient permissions to view the workflow
- **404** Workflow does not exist
- **501** Workflows are not enabled. To enable workflows, see ["workflow " on page 752](#)

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Get workflow details by workflow name or test id

Summary

Gets workflow details by workflow name or for workflows associated with a test id

GET /api/v10/workflows

Description

Gets workflow details by workflow name or for workflows associated with a test id. You can enter multiple workflow names or test ids in the request if required.

Tip

If the **name** and **testdefinitions** parameters are both in the request, only **name** is used for the request and the **testdefinitions** parameter is ignored.

Parameters

Parameter	Description	Type	Parameter Type	Required
name	Specify the workflow you want to return details for with the workflow name parameter. To request details for multiple workflows, enter the workflow names as a comma separated list.	string	query	No
testdefinitions	Get details of workflows associated with a test id with the testdefinitions parameter. To request multiple test ids, enter the test ids as a comma separated list.	string	query	No

Example usage

Get details for myWorkflow 1

```
curl -u "username:ticket" "https://my-swarm-host/api/v10/workflows?name=myWorkflow 1"
```

JSON Response:

```
HTTP/1.1 200 OK
```

```
{
  "error": null,
  "messages": [],
  "data": {
    "workflows": [
      {
        "on_submit": {
          "with_review": {
            "rule": "no_checking",
```



```
        "mode": "inherit"
      },
      "without_review":{
        "rule": "no_checking",
        "mode": "inherit"
      }
    },
    "name": "myWorkflow 1",
    "description": "Another description",
    "shared": "true",
    "owners": [
      "user3",
      "user4"
    ],
    "end_rules":{
      "update":{
        "rule":"no_revision",
        "mode": "inherit"
      }
    },
    "auto_approve":{
      "rule": "votes",
      "mode": "inherit"
    },
    "counted_votes": {
      "rule": "members",
      "mode": "inherit"
    },
    "tests": [],
    "id": "1"
  },
]
}
```

Get details for workflows associated with test id 1 or test id 2

```
curl -u "username:ticket" "https://my-swarm-host/api/v10/workflows?testdefinitions=1,2"
```

JSON Response:

```
HTTP/1.1 200 OK
```

```
{
  "error": null,
  "messages": [],
  "data": {
    "workflows": [
      {
        "on_submit": {
          "with_review": {
            "rule": "no_checking",
            "mode": "default"
          },
          "without_review": {
            "rule": "no_checking",
            "mode": "policy"
          }
        },
        "name": "Our Workflow",
        "description": "This is the our workflow and it is good",
        "shared": "true",
        "owners": [
          "swarm"
        ],
        "end_rules": {
          "update": {
            "rule": "no_checking",
            "mode": "default"
          }
        },
        "auto_approve": {
          "rule": "never",
```

```
        "mode": "default"
    },
    "counted_votes":{
        "rule": "anyone",
        "mode": "default"
    },
    "group_exclusions":{
        "rule": [ ],
        "mode": "policy"
    },
    "user_exclusions":{
        "rule": [ ],
        "mode": "policy"
    },
    "tests": [
        {
            "id": 1,
            "event": "onUpdate"
            "blocks": "approved"
        },
        {
            "id": 2,
            "event": "onSubmit"
            "blocks": "nothing"
        }
    ],
    "id": "3"
},
{
    "on_submit":{
        "with_review":{
            "rule": "no_checking",
            "mode": "inherit"
        },
        "without_review":{
```

```
        "rule": "no_checking",
        "mode": "inherit"
    },
    {
        "name": "Another Workflow",
        "description": "Yet another workflow description",
        "shared": "false",
        "owners": [
            "user3",
            "user4"
        ],
        "end_rules": {
            "update": {
                "rule": "no_revision",
                "mode": "inherit"
            }
        },
        "auto_approve": {
            "rule": "votes",
            "mode": "inherit"
        },
        "counted_votes": {
            "rule": "members",
            "mode": "inherit"
        },
        "tests": [
            {
                "id": 2,
                "event": "onSubmit"
                "blocks": "approved"
            },
            {
                "id": 5,
                "event": "onSubmit"
                "blocks": "approved"
            }
        ]
    }
]
```

```
        },
      ],
      "id": "6"
    },
  ]
}
```

If a request fails

<error code>:

- **401** Insufficient permissions to view the workflow
- **404** Workflow or id does not exist
- **501** Workflows are not enabled. To enable workflows, see ["workflow " on page 752](#)

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Create a workflow

Summary

Create a workflow

POST /api/v10/workflows/

Description

Create a new workflow.

Tip

The global workflow is present by default so it can only be edited.

Parameters

Parameter	Description	Type	Parameter Type	Required
name	The workflow name. Will be compared against other workflows and rejected if not unique.	string	body	Yes
description	Description for the new workflow.	string	body	Yes
shared	Whether this workflow is shared for other users that do not own it. Defaults to not shared.	boolean	body	Yes
owners	A list owners for the workflow. Can be users or group names (prefixed with swarm-group-). Users and group names must exist or the workflow will be rejected.	array (of strings)	body	Yes
mode	Each workflow rule must have a mode parameter. The only valid value for mode is inherit .	string	body	Yes
on_submit	Data for rules when changes are submitted. Valid values for with_review are no_checking , approved , strict . Valid values for without_review are no_checking , auto_create , reject .	array	body	Yes

Parameter	Description	Type	Parameter Type	Required
end_rules	Data for rules when changes are submitted. Valid values are no_checking , no_revision .	array	body	Yes
auto_approve	Data for rules when changes are submitted. Valid values are votes , never .	array	body	Yes
counted_votes	Data for rules when counting votes up. Valid values are anyone , members .	array	body	Yes
tests	<p>A list of test ids for the workflow. Each test id has an event that determines when the test is run and a blocks value that determines whether a test failure blocks review approval.</p> <ul style="list-style-type: none"> Valid values for event are onUpdate, onSubmit. Valid values for blocks are nothing and approved. 	array	body	No

Example usage

Create a workflow

```
curl -X POST -H "Content-Type: application/json" -u "username:ticket" -d "@mybodyfilename.txt" "https://myswarm-url/api/v10/workflows/"
```

The "mybodyfilename.txt" file contains the information for the new workflow:

```
{
  "name": "My workflow"
  "description": "This is my workflow."
  "shared": "true"
  "owners": "Francois_Piccard,Anna_Schmidt"
  "on_submit":{
    "with_review":{
      "rule":"no_checking",
      "mode": "inherit"
    },
    "without_review":{
      "rule":"no_checking",
      "mode": "inherit"
    }
  },
  "end_rules": {
    "update":{
      "rule":"no_revision",
      "mode": "inherit"
    }
  },
  "auto_approve":{
    "rule": "never",
    "mode": "inherit"
  },
  "counted_votes": {
    "rule": "members",
    "mode": "inherit"
  },
  "tests": [
    {
      "id": 2,
      "event": "onUpdate"
      "blocks": "nothing"
    },
  ],
}
```



```
{
  "id": 5,
  "event": "onSubmit"
  "blocks": "approved"
}
]
```

Swarm responds with:

HTTP/1.1 200 OK

```
{
  "error": null,
  "messages": [],
  "data": {
    "workflows": [
      {
        "on_submit":{
          "with_review":{
            "rule": "no_checking",
            "mode": "inherit"
          },
          "without_review":{
            "rule": "no_checking",
            "mode": "inherit"
          }
        },
        "name": "My workflow",
        "description": "This is my workflow.",
        "shared": "true",
        "owners": [
          "Francois_Piccard",
          "Anna_Schmidt"
        ],
        "end_rules":{
          "update":{
```

```

        "rule": "no_revision",
        "mode": "inherit"
    },
    {
        "auto_approve": {
            "rule": "votes",
            "mode": "inherit"
        },
        "counted_votes": {
            "rule": "members",
            "mode": "inherit"
        },
        "tests": [
            {
                "id": 2,
                "event": "onUpdate"
                "blocks": "nothing"
            },
            {
                "id": 5,
                "event": "onSubmit"
                "blocks": "approved"
            }
        ],
        "id": "2"
    },
]
}

```

If a request fails

<error code>:

- **400** Invalid parameter data specified
- **401** Insufficient permissions to create a workflow
- **501** Workflows are not enabled. To enable workflows, see ["workflow " on page 752](#)

```
HTTP/1.1 <response error code>

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Update a workflow

Summary

Update a workflow

PUT /api/v10/workflows/{id}

Description

Update a workflow. All values should be provided in the request. If not provided any missing values are reverted to default.

Parameters

Parameter	Description	Type	Parameter Type	Required
id	The id of the workflow being updated	integer	path	Yes
name	The workflow name. Will be compared against other workflows and rejected if not unique	string	body	Yes

Parameter	Description	Type	Parameter Type	Required
owners	A list owners for the workflow. Can be users or group names (prefixed with swarm-group-). Users and group names must exist or the workflow will be rejected	array (of strings)	body	Yes
description	Description for the new workflow	string	body	Yes
shared	Whether this workflow is shared for other users that do not own it. Defaults to not shared	boolean	body	Yes

Parameter	Description	Type	Parameter Type	Required
<code>mode</code>	<p>Workflow rules:</p> <p>Each workflow rule must have a mode parameter. The only valid value for mode is inherit.</p> <p>Global workflow rules only:</p> <ul style="list-style-type: none"> ▪ default applies the workflow rule setting to projects and project branches that don't have an associated workflow. If a project or project branch has an associated Swarm workflow, the global rule is ignored. ▪ policy applies a minimum workflow rule setting to all projects and project branches. If a project or project branch has an 	string	body	Yes

Parameter	Description	Type	Parameter Type	Required
	associated workflow, the global rule is merged with the workflow rule and the most restrictive setting is used. Displayed as Enforce on in the Swarm UI, see " Global workflow " on page 753.			
<code>on_submit</code>	Data for rules when changes are submitted. Valid values for <code>with_review</code> are <code>no_checking</code> , <code>approved</code> , <code>strict</code> . Valid values for <code>without_review</code> are <code>no_checking</code> , <code>auto_create</code> , <code>reject</code>	array	body	Yes
<code>end_rules</code>	Data for rules when changes are submitted. Valid values are <code>no_checking</code> , <code>no_revision</code> .	array	body	Yes
<code>auto_approve</code>	Data for rules when changes are submitted. Valid values are <code>votes</code> , <code>never</code> .	array	body	Yes

Parameter	Description	Type	Parameter Type	Required
<code>counted_votes</code>	Data for rules when counting votes up. Valid values are <code>anyone</code> , <code>members</code> .	array	body	Yes
<code>tests</code>	<p>A list of test <code>ids</code> for the workflow. Each test <code>id</code> has an <code>event</code> that determines when the test is run and a <code>blocks</code> value that determines whether a test failure blocks review approval.</p> <ul style="list-style-type: none"> Valid values for <code>event</code> are <code>onUpdate</code>, <code>onSubmit</code>. Valid values for <code>blocks</code> are <code>nothing</code> and <code>approved</code>. 	array	body	No

Example usage

Update a workflow

```
curl -X PUT -H "Content-Type: application/json" -u "username:ticket" -d "@mybodyfilename.txt" "https://myswarm-url/api/v10/workflows/2"
```

The "mybodyfilename.txt" file contains the updated workflow information:

```
{
  "name": "My workflow",
  "description": "This is my workflow and I have changed the
description of it.",
  "shared": "true",
  "owners": ["Francois_Piccard,Anna_Schmidt"],
  "on_submit":{
    "with_review":{
```

```

        "rule": "no_checking",
        "mode": "inherit"
    },
    "without_review": {
        "rule": "no_checking",
        "mode": "inherit"
    }
},
"end_rules": {
    "update": {
        "rule": "no_revision",
        "mode": "inherit"
    }
},
"auto_approve": {
    "rule": "never",
    "mode": "inherit"
},
"counted_votes": {
    "rule": "members",
    "mode": "inherit"
},
"tests": [{
    "id": 2,
    "event": "onUpdate",
    "blocks": "nothing"
},
{
    "id": 5,
    "event": "onSubmit",
    "blocks": "approved"
}
]
}

```

JSON Response:

HTTP/1.1 200 OK

```
{
  "error": null,
  "messages": [],
  "data": {
    "workflows": [{
      "on_submit": {
        "with_review": {
          "rule": "no_checking",
          "mode": "inherit"
        },
        "without_review": {
          "rule": "no_checking",
          "mode": "inherit"
        }
      },
      "name": "My workflow",
      "description": "This is my workflow and I have changed the
description of it.",
      "shared": "true",
      "owners": [
        "Francois_Piccard",
        "Anna_Schmidt"
      ],
      "end_rules": {
        "update": {
          "rule": "no_revision",
          "mode": "inherit"
        }
      },
      "auto_approve": {
        "rule": "votes",
        "mode": "inherit"
      }
    }
  ]
}
```

```

        "counted_votes": {
            "rule": "members",
            "mode": "inherit"
        },
        "tests": [
            {
                "id": 2,
                "event": "onUpdate",
                "blocks": "nothing"
            },
            {
                "id": 5,
                "event": "onSubmit",
                "blocks": "approved"
            }
        ],
        "id": "2"
    }
}

```

If a request fails

<error code>:

- **400** Invalid parameter data specified
- **401** Insufficient permissions to update the workflow
- **404** Workflow does not exist or you do not have permission to update the workflow
- **501** Workflows are not enabled. To enable workflows, see ["workflow " on page 752](#)

HTTP/1.1 <response error code>

```

{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
}

```

```
"data" : null
}
```

Delete a workflow

Summary

Delete a workflow

DELETE /api/v10/workflows/{id}

Description

Delete a workflow for the provided id. This call must be authenticated and the user must have permission to edit the workflow. If the workflow is in use it cannot be deleted and an error message will be returned. The global workflow cannot be deleted.

Parameters

Parameter	Description	Type	Parameter Type	Required
id	The id of the workflow being deleted	string	form	Yes

Example usage

Delete a workflow that is not in use by a project

```
curl -u "username:ticket" -X DELETE "https://my-swarm-host/api/v10/workflows/1"
```

JSON Response

```
HTTP/1.1 200
```

```
{
  "error": null,
  "messages": [
    "Workflow [1] was deleted"
  ],
  "data": null
}
```

If a request fails

<error code>:

- **401** Insufficient permissions to delete the workflow
- **403** Cannot delete the global workflow
- **404** Workflow does not exist or you do not have permission to view this workflow
- **501** Workflows are not enabled. To enable workflows, see ["workflow " on page 752](#)

HTTP/1.1 <response error code>

```
{
  "error": <error code>,
  "messages": [{
    "code" : "<code string>",
    "text" : "<error message>"
  }],
  "data" : null
}
```

Glossary

A

access level

A permission assigned to a user to control which commands the user can execute. See also the 'protections' entry in this glossary and the 'p4 protect' command in the P4 Command Reference.

admin access

An access level that gives the user permission to privileged commands, usually super privileges.

APC

The Alternative PHP Cache, a free, open, and robust framework for caching and optimizing PHP intermediate code.

archive

1. For replication, versioned files (as opposed to database metadata). 2. For the 'p4 archive' command, a special depot in which to copy the server data (versioned files and metadata).

atomic change transaction

Grouping operations affecting a number of files in a single transaction. If all operations in the transaction succeed, all the files are updated. If any operation in the transaction fails, none of the files are updated.

avatar

A visual representation of a user or group. Avatars are used in Swarm to show involvement in (or ownership of) projects, groups, changelists, reviews, comments.

B

base

For files: The file revision that contains the most common edits or changes among the file revisions in the source file and target file paths. For checked out streams: The public have version from which the checked out version is derived.

binary file type

A Helix Core Server file type assigned to a non-text file. By default, the contents of each revision are stored in full, and file revision is stored in compressed format.

branch

(noun) A set of related files that exist at a specific location in the Helix Core depot as a result of being copied to that location, as opposed to being added to that location. A group of related files is often referred to as a codeline. To associate code reviews in Helix Swarm with the projects they are part of, add the "branch" paths in the Swarm project. (verb) To create a codeline by copying another codeline with the 'p4 integrate', 'p4 copy', or 'p4 populate' command.

branch form

The form that appears when you use the 'p4 branch' command to create or modify a branch specification.

branch mapping

Specifies how a branch is to be created or integrated by defining the location, the files, and the exclusions of the original codeline and the target codeline. The branch mapping is used by the integration process to create and update branches.

branch view

A specification of the branching relationship between two codelines in the depot. Each branch view has a unique name and defines how files are mapped from the originating codeline to the target codeline. This is the same as branch mapping.

broker

Helix Broker, a server process that intercepts commands to the Helix Core Server and is able to run scripts on the commands before sending them to the Helix Core Server.

C

change review

The process of sending email to users who have registered their interest in changelists that include specified files in the depot.

changelist

A list of files, their version numbers, the changes made to the files, and a description of the changes made. A changelist is the basic unit of versioned work in Helix Core Server. The changes specified in the changelist are not stored in the depot until the changelist is submitted to the depot. See also atomic change transaction and changelist number.

changelist form

The form that appears when you modify a changelist using the 'p4 change' command.

changelist number

An integer that identifies a changelist. Submitted changelist numbers are ordinal (increasing), but not necessarily consecutive. For example, 103, 105, 108, 109. A pending changelist number might be assigned a different value upon submission.

check in

To submit a file to the Helix Core Server depot.

check out

To designate one or more files, or a stream, for edit.

checkpoint

A backup copy of the underlying metadata at a particular moment in time. A checkpoint can recreate db.user, db.protect, and other db.* files. See also metadata.

classic depot

A repository of Helix Core files that is not streams-based. Uses the Perforce file revision model, not the graph model. The default depot name is depot. See also default depot, stream depot, and graph depot.

client form

The form you use to define a client workspace, such as with the 'p4 client' or 'p4 workspace' commands.

client name

A name that uniquely identifies the current client workspace. Client workspaces, labels, and branch specifications cannot share the same name.

client root

The topmost (root) directory of a client workspace. If two or more client workspaces are located on one machine, they should not share a client root directory.

client side

The right-hand side of a mapping within a client view, specifying where the corresponding depot files are located in the client workspace.

client workspace

Directories on your machine where you work on file revisions that are managed by Helix Core Server. By default, this name is set to the name of the machine on which your client workspace is located, but it can be overridden. Client workspaces, labels, and branch specifications cannot share the same name.

code review

A process in Helix Swarm by which other developers can see your code, provide feedback, and approve or reject your changes.

codeline

A set of files that evolve collectively. One codeline can be branched from another, allowing each set of files to evolve separately.

comment

Feedback provided in Helix Swarm on a changelist, review, job, or a file within a changelist or review.

commit server

The innermost Helix Core Server in a multi-server topology that includes edge servers. The commit server processes submitted files (changelists), global workspaces, and promoted shelves.

conflict

1. A situation where two users open the same file for edit. One user submits the file, after which the other user cannot submit unless the file is resolved. 2. A resolve where the same line is changed when merging one file into another. This type of conflict occurs when the comparison of two files to a base yields different results, indicating that the files have been changed in different ways. In this case, the merge cannot be done automatically and must be resolved manually. See file conflict.

copy up

A Helix Core Server best practice to copy (and not merge) changes from less stable lines to more stable lines. See also merge.

counter

A numeric variable used to track variables such as changelists, checkpoints, and reviews.

D

database

The database in the P4ROOT directory contains db.* files with metadata that the Helix Core Server uses to operate on versioned files, users, protections, streams, changelists, and more.

default changelist

The changelist used by a file add, edit, or delete, unless a numbered changelist is specified. A default pending changelist is created automatically when a file is opened for edit.

deleted file

In Helix Core Server, a file with its head revision marked as deleted. Older revisions of the file are still available. See also "obliterate".

delta

The differences between two files.

depot

A file repository hosted on the Helix Core Server. A depot is the top-level unit of storage for versioned files, which are also known as depot files, archive files, or source files. It contains all versions of all files ever submitted to the depot, including deleted files (but not obliterated files). There can be multiple depots on a single installation.

depot root

The topmost (root) directory for a depot.

depot side

The left side of any client view mapping, specifying the location of files in a depot.

depot syntax

Helix Core Server syntax for specifying the location of files in the depot. Depot syntax begins with:
`//depot/`

diff

(noun) A set of lines that do not match when two files, or stream versions, are compared. A conflict is a pair of unequal diffs between each of two files and a base, or between two versions of a stream. (verb) To compare the contents of files or file revisions, or of stream versions. See also conflict.

E

edge server

A replica server that is part of an edge/commit system that is able to process most read/write commands, including 'p4 integrate', and also deliver versioned files (depot files).

exclusionary access

A permission that denies access to the specified files.

exclusionary mapping

A view mapping that excludes specific files or directories.

extension

Custom logic that runs in a Lua engine embedded into the Helix Server. Helix Core Extensions are a recent alternative to triggers, which require a runtime that is external to the server. See the Helix Core Extensions Developer Guide.

F

file conflict

In a three-way file merge, a situation in which two revisions of a file differ from each other and from their base file. Also, an attempt to submit a file that is not an edit of the head revision of the file in the depot, which typically occurs when another user opens the file for edit after you have opened the file for edit.

file pattern

Helix Core Server command line syntax that enables you to specify files using wildcards.

file repository

The server's copy of all files, which is shared by all users. In Helix Core Server, this is called the depot.

file revision

A specific version of a file within the depot. Each revision is assigned a number, in sequence. Any revision can be accessed in the depot by its revision number, preceded by a pound sign (#), for example testfile#3.

file tree

All the subdirectories and files under a given root directory.

file type

An attribute that determines how Helix Core Server stores and diffs a particular file. Examples of file types are text and binary.

fix

A job that has been closed in a changelist.

form

A screen displayed by certain Helix Core Server commands. For example, you use the change form to enter comments about a particular changelist to verify the affected files.

forwarding replica

A replica server that can process read-only commands and deliver versioned files (depot files). One or more replicate servers can significantly improve performance by offloading some of the master server load. In many cases, a forwarding replica can become a disaster recovery server.

G

Git Connector

The Git Connector component of Git Connector interacts with your native Git clients. It serves Git commands and performs all Git repo operations. See "hybrid workspace".

Git Fusion

(No longer offered to new customers) A Perforce product that integrates Git with Helix, offering enterprise-ready Git repository management, and workflows that allow Git and Helix Core Server users to collaborate on the same projects using their preferred tools.

graph depot

A depot of type graph that is used to store Git repos in the Helix Core Server. See also Git Connector and classic depot.

group

A feature in Helix Core Server that makes it easier to manage permissions for multiple users.

H

have list

The list of file revisions currently in the client workspace.

head revision

The most recent revision of a file within the depot. Because file revisions are numbered sequentially, this revision is the highest-numbered revision of that file.

heartbeat

A process that allows one server to monitor another server, such as a standby server monitoring the master server (see the p4 heartbeat command).

Helix Core Server

The Helix Core Server depot and metadata; also, the program that manages the depot and metadata, also called Helix Core Server.

Helix TeamHub

A Perforce management platform for code and artifact repository. TeamHub offers built-in support for Git, SVN, Mercurial, Maven, and more.

hybrid workspace

A workspace that supports both repos of type graph (see "Git Connector"), and the Helix Core file revision model.

I

integrate

To compare two sets of files (for example, two codeline branches) and determine which changes in one set apply to the other, determine if the changes have already been propagated, and propagate any outstanding changes from one set to another.

J

job

A user-defined unit of work tracked by Helix Core Server. The job template determines what information is tracked. The template can be modified by the Helix Core Server system administrator. A job describes work to be done, such as a bug fix. Associating a job with a changelist records which changes fixed the bug.

job specification

A form describing the fields and possible values for each job stored in the Helix Core Server machine.

journal

A file containing a record of every change made to the Helix Core Server's metadata since the time of the last checkpoint. This file grows as each Helix Core Server transaction is logged. The file should be automatically truncated and renamed into a numbered journal when a checkpoint is taken.

journal rotation

The process of renaming the current journal to a numbered journal file.

journaling

The process of recording changes made to the Helix Core Server's metadata.

L

label

A named list of user-specified file revisions.

label view

The view that specifies which filenames in the depot can be stored in a particular label.

lazy copy

A method used by Helix Core Server to make internal copies of files without duplicating file content in the depot. A lazy copy points to the original versioned file (depot file). Lazy copies minimize the consumption of disk space by storing references to the original file instead of copies of the file.

librarian

The librarian subsystem of the server stores, manages, and provides the archive files to other subsystems of the Helix Core server.

license file

A file that ensures that the number of Helix Core Server users on your site does not exceed the number for which you have paid.

list access

A protection level that enables you to run reporting commands but prevents access to the contents of files.

local depot

Any depot located on the currently specified Helix Core Server.

local syntax

The syntax for specifying a filename that is specific to an operating system.

lock

1. A file lock that prevents other clients from submitting the locked file. Files are unlocked with the 'p4 unlock' command or by submitting the changelist that contains the locked file. 2. A database lock that prevents another process from modifying the database db.* file.

log

Error output from the Helix Core Server. To specify a log file, set the P4LOG environment variable or use the p4d -L flag when starting the service.

M

mapping

A single line in a view, consisting of a left side and a right side that specify the correspondences between files in the depot and files in a client, label, or branch. See also workspace view, branch view, and label view.

master server

The innermost Helix Core server in a multi-server topology. In the server spec, the Services field must be set to 'commit-server' for edge-commit architecture, and is typically set to 'standard' for master-replica architecture.

MDS checksum

The method used by Helix Core Server to verify the integrity of versioned files (depot files).

merge

1. To create new files from existing files, preserving their ancestry (branching). 2. To propagate changes from one set of files to another. 3. The process of combining the contents of two conflicting file revisions into a single file, typically using a merge tool like P4Merge.

merge file

A file generated by the Helix Core Server from two conflicting file revisions.

metadata

The data stored by the Helix Core Server that describes the files in the depot, the current state of client workspaces, protections, users, labels, and branches. Metadata is stored in the Perforce database and is separate from the archive files that users submit.

modification time or modtime

The time a file was last changed.

MPM

Multi-Processing Module, a component of the Apache web server that is responsible for binding to network ports, accepting requests, and dispatch operations to handle the request.

N

nonexistent revision

A completely empty revision of any file. Syncing to a nonexistent revision of a file removes it from your workspace. An empty file revision created by deleting a file and the `#none` revision specifier are examples of nonexistent file revisions.

numbered changelist

A pending changelist to which Helix Core Server has assigned a number.

O

obliterate

The `p4 obliterate` command permanently removes files and their history from the depot. See also "deleted file".

offline database

A copy of a Helix Core Server database that is kept up to date by manual or scripted processes that replay journals from a Helix Core Server.

opened file

A file you have checked out in your client workspace as a result of a Helix Core Server operation (such as an edit, add, delete, integrate). Opening a file from your operating system file browser is not tracked by Helix Core Server.

owner

The Helix Core Server user who created a particular client, branch, or label.

P

p4

1. The Helix Core Server command line program. 2. The command you issue to execute commands from the operating system command line.

p4d

The program that runs the Helix Core Server; p4d manages depot files and metadata.

P4PHP

The PHP interface to the Helix API, which enables you to write PHP code that interacts with a Helix Core Server machine.

PECL

PHP Extension Community Library, a library of extensions that can be added to PHP to improve and extend its functionality.

pending changelist

A changelist that has not been submitted.

project

In Helix Swarm, a group of Helix Core Server users who are working together on a specific codebase, defined by one or more branches of code, along with options for a job filter, automated test integration, and automated deployment.

protections

The permissions stored in the Helix Core Server's protections table.

proxy server

A Helix Core Server that stores versioned files. A proxy server does not perform any commands. It serves versioned files to Helix Core Server clients.

R

RCS format

Revision Control System format. Used for storing revisions of text files in versioned files (depot files). RCS format uses reverse delta encoding for file storage. Helix Core Server uses RCS format to store text files. See also reverse delta storage.

read access

A protection level that enables you to read the contents of files managed by Helix Core Server but not make any changes.

remote depot

A depot located on another Helix Core Server accessed by the current Helix Core Server.

replica

A Helix Core Server that automatically maintains a full or partial copy of metadata and none, some, or all of the related file archives by copying data from a master Helix Core Server using 'p4 pull' or 'p4 journalcopy'.

repo

A graph depot contains one or more repos, and each repo contains files from Git users.

reresolve

The process of resolving a file after the file is resolved and before it is submitted.

resolve

The process you use to manage the differences between two revisions of a file, or two versions of a stream. You can choose to resolve file conflicts by selecting the source or target file to be submitted, by merging the contents of conflicting files, or by making additional changes. To resolve stream conflicts, you can choose to accept the public source, accept the checked out target, manually accept changes, or combine path fields of the public and checked out versions while accepting all other changes made in the checked out version.

reverse delta storage

The method that Helix Core Server uses to store revisions of text files. Helix Core Server stores the changes between each revision and its previous revision, plus the full text of the head revision.

revert

To discard the changes you have made to a file in the client workspace before a submit.

review access

A special protections level that includes read and list accesses and grants permission to run the p4 review command.

review daemon

A program that periodically checks the Helix Core Server machine to determine if any changelists have been submitted. If so, the daemon sends an email message to users who have subscribed to any of the files included in those changelists, informing them of changes in files they are interested in.

revision number

A number indicating which revision of the file is being referred to, typically designated with a pound sign (#).

revision range

A range of revision numbers for a specified file, specified as the low and high end of the range. For example, myfile#5,7 specifies revisions 5 through 7 of myfile.

revision specification

A suffix to a filename that specifies a particular revision of that file. Revision specifiers can be revision numbers, a revision range, change numbers, label names, date/time specifications, or client names.

RPM

RPM Package Manager. A tool, and package format, for managing the installation, updates, and removal of software packages for Linux distributions such as Red Hat Enterprise Linux, the Fedora Project, and the CentOS Project.

S

server data

The combination of server metadata (the Helix Core Server database) and the depot files (your organization's versioned source code and binary assets).

server root

The topmost directory in which p4d stores its metadata (db.* files) and all versioned files (depot files or source files). To specify the server root, set the P4ROOT environment variable or use the p4d -r flag.

service

In the Helix Core Server, the shared versioning service that responds to requests from Helix Core Server client applications. The Helix Core Server (p4d) maintains depot files and metadata describing the files and also tracks the state of client workspaces.

shelve

The process of temporarily storing files in the Helix Core Server without checking in a changelist.

status

For a changelist, a value that indicates whether the changelist is new, pending, or submitted. For a job, a value that indicates whether the job is open, closed, or suspended. You can customize job statuses. For the 'p4 status' command, by default the files opened and the files that need to be reconciled.

storage record

An entry within the db.storage table to track references to an archive file.

stream

A "branch" with built-in rules that determine what changes should be propagated to files in a stream depot, and in what order they should be propagated. A stream spec defines a stream. A user creates a stream spec by using either the "p4 stream" command or in P4V with File > New Stream. In P4V, stream specs are visible in the Streams Graph and the Streams tab.

stream depot

A depot used with streams and stream clients. Has structured branching, unlike the free-form branching of a "classic" depot. Uses the file revision model, not the graph model. See also classic depot and graph depot.

stream hierarchy

The set of parent-to-child relationships between streams in a stream depot.

stream view

A stream view is defined by the Paths, Remapped, and Ignored fields of the stream specification. (See Form Fields in the p4 stream command)

submit

To send a pending changelist into the Helix Core Server depot for processing.

super access

An access level that gives the user permission to run every Helix Core Server command, including commands that set protections, install triggers, or shut down the service for maintenance.

symlink file type

A Helix Core Server file type assigned to symbolic links. On platforms that do not support symbolic links, symlink files appear as small text files.

sync

To copy a file revision (or set of file revisions) from the Helix Core Server depot to a client workspace.

T

target file

The file that receives the changes from the donor file when you integrate changes between two codelines.

text file type

Helix Core Server file type assigned to a file that contains only ASCII text, including Unicode text. See also binary file type.

theirs

The revision in the depot with which the client file (your file) is merged when you resolve a file conflict. When you are working with branched files, theirs is the donor file.

three-way merge

The process of combining three file revisions. During a three-way merge, you can identify where conflicting changes have occurred and specify how you want to resolve the conflicts.

topology

The set of Helix Core services deployed in a multi-server installation, which might include commit-server, edge servers, standby servers, proxies, brokers, and more.

trigger

A script that is automatically invoked by Helix Core Server when various conditions are met. (See "Helix Core Server Administrator Guide" on "Triggers".)

two-way merge

The process of combining two file revisions. In a two-way merge, you can see differences between the files.

typemap

A table in Helix Core Server in which you assign file types to files.

U**user**

The identifier that Helix Core Server uses to determine who is performing an operation. The three types of users are standard, service, and operator.

V**versioned files**

Source files stored in the Helix Core Server depot, including one or more revisions to each file. Also known as archive files, archives, and depot files. Versioned files typically use the naming convention 'filename,v' or '1.changelist.gz'.

view

A description of the relationship between two sets of files. See workspace view, label view, branch view.

W

wildcard

A special character used to match other characters in strings. The following wildcards are available in Helix Core Server: * matches anything except a slash; ... matches anything including slashes; %%0 through %%9 is used for parameter substitution in views.

workspace

See client workspace.

workspace view

A set of mappings that specifies the correspondence between file locations in the depot and the client workspace.

write access

A protection level that enables you to run commands that alter the contents of files in the depot. Write access includes read and list accesses.

X

XSS

Cross-Site Scripting, a form of web-based attack that injects malicious code into a user's web browser.

Y

yours

The edited version of a file in your client workspace when you resolve a file. Also, the target file when you integrate a branched file.

Getting help

Tip

Perforce Support cannot help you with user account creation or with resetting your password, please contact your company's Perforce administrator for help with this.

We look forward to hearing about your experiences with Swarm, positive or negative, including *must-haves* or *it would be great if Swarm....* Please feel free to contact us.

Please visit the [Perforce Support portal](#) for product support and contact information:

- Search our knowledge base
- Post on the Perforce forum
- Submit a Support request
- Call us for support
- View our video tutorials
- View our training options
- View the documentation and release notes
- Download Perforce software

Complete contact information is available on the [Perforce web site](#).

License statements

For the licensing information of the third-party software included in this Perforce product, see the license file on the Swarm server in `<swarm_root>/readme/LICENSE.txt`.