

PERFORCE

Introducing Helix

2015.2

August 2015

Introducing Helix 2015.2

August 2015

Copyright © 2015 Perforce Software.

All rights reserved.

Perforce software and documentation is available from <http://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 Software.

Perforce Software 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 Software.

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

Any additional software included within Perforce software is listed in [License Statements on page 25](#).

Table of Contents

Chapter 1	Basic Concepts	1
	The basics of version control	1
	Helix as a version control implementation	2
	Multiple user access to a set of files	3
	Balancing stability and innovation: the mainline model	4
	Streams	6
	Organizing your work: jobs and labels	7
	Working together and working apart: centralized and distributed development	8
	Collaborating within a Git Ecosystem	9
	Performance, scaling, and high availability	10
	Using proxies to improve performance	10
	Using a replica for disaster recovery	11
	Commit-edge architecture	11
	Data center scaling and high availability	12
	Securing the system	13
Chapter 2	Clients, Collaboration, and Security	15
	Client applications	15
	Analytics tools	15
	Helix Insights	15
	Helix Threat Detection	16
	Collaboration	17
	Helix Swarm	17
	Helix GitSwarm	17
	Integration	18
	Interactive development environment integrations	18
	Build and reporting integrations	18
Chapter 3	Customizing and Extending Helix	19
	APIs	19
	Triggers	19
	Helix Broker	20
Chapter 4	Use Cases	21
	Software development	21
	Digital asset management	21
	Hybrid product development	21
	Ease of use for Helix administration	22
	Growing into the future	22

Chapter 5	Documentation and Other Resources	23
	To learn more about Helix	23
	Documentation links	23
	Download URLs	24
	Videos	24
	License Statements	25

This document introduces Helix, a secure, scalable, and highly available version control system that supports parallel development. You should read this document before you start working with Helix.

This document:

- introduces the basic concepts and tasks of version control
- explains how you can configure Helix to improve performance, to scale the system, and to provide automated failover
- suggests some ways you can extend and customize Helix
- explains how you can use Helix with other products to get additional functionality
- discusses use cases for Helix
- introduces the many support resources that can help you use Helix.

If you are familiar with version control systems, you can skip the first section and begin by reading “Helix as a version control implementation.”

The basics of version control

When you work alone on a document, the latest is usually the greatest: you successively open the document, make changes, and save the document. Each time you save, you overwrite the existing copy. The situation is different when you are working with a large, globally distributed team on a project consisting of hundreds or even thousands of files. In this case, it is important to track authorship and changes, and to resolve conflicts when several users make contending changes to the same file. Version control systems allow you to do this; you can track and manage changes to any large collection of digital assets: documents, source code, web sites, audio files, and so on.

Version control systems address these issues using a variety of techniques; one of the most basic is versioning. Rather than overwriting earlier versions of a file when it is saved, under version control, each saved copy of the file is versioned and assigned a number or letter that reflects the order in which it was saved.

In addition to identifying a file version within a sequence of versions, a version control system automatically associates certain information with each version: it records who made the change, when the change was made, and why the change was made. Taken together this information provides an audit trail that you can always consult to understand how a project developed and when specific changes were made. Because no version of a file is overwritten, when bugs arise, it is possible to look back in time to identify the point at which the bug was introduced. This can be critical in fixing bugs that cannot be reproduced. Equally, looking at file history and understanding why certain decisions were made can help project participants stay on track or find appropriate options for future directions.

Sharing data under version control requires a certain amount of gatekeeping to determine who can access the data and how conflicts are resolved when two users make changes to the same file. We have explained so far about editing and saving files; but to support this gatekeeping function, version control systems introduce the additional step of checking out a file and checking in or submitting a file.

The basic version control workflow looks like this:

1. Assets under version control are placed in a specified repository.
2. Assets are associated with specific permissions that users must have in order to read or modify them.
3. A user checks out a working copy of an asset and makes changes.
4. Another user checks out a working copy of the same asset and makes changes.
5. The first user saves changes to their working copy and checks in that copy.
6. The second user saves changes to their working copy and checks in that copy.
7. The version control system detects the fact that the same asset was changed in parallel, and it asks the second user to merge changes with those of the first user before the second user's changes can be checked in. The work of comparing and merging changes is called *resolving*.

In this way, the version control system makes sure that changes are predictable, manageable, and auditable.

Version control systems are traditionally either centralized or distributed:

- *Centralized version control systems* use a single repository from which users check out one or more files to work on locally.
- *Distributed version control systems* allow users to host repositories locally, check out entire repositories with all history, work independently of one another, and combine their work through merging when necessary.

Helix supports either model, as well as a hybrid of the two.

Version control systems can be used as stand-alone applications or they can be incorporated into development or authoring tools as a means of managing the assets produced by these tools.

Helix as a version control implementation

Helix uses a client-server architecture to implement version control management.

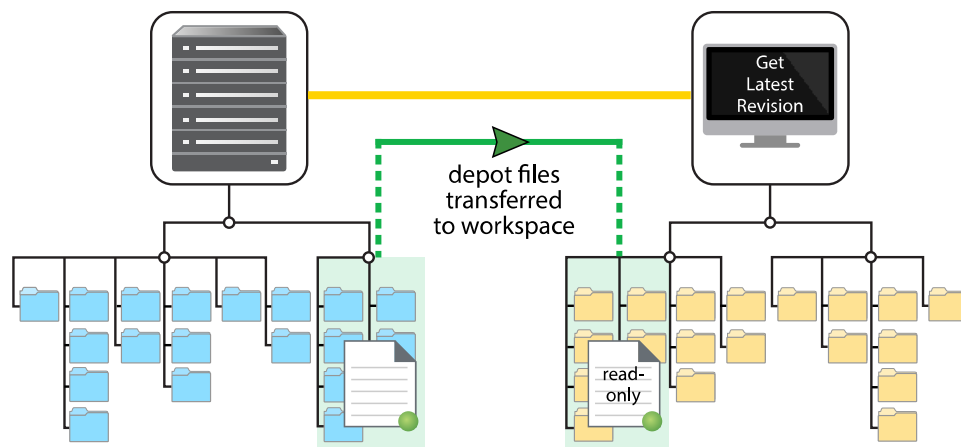
- The Helix *server* (also known as the Helix Versioning Engine or **p4d**) manages shared file repositories, or depots, that contain every revision of every file under version management. Files are organized into directory trees. The server also maintains a database to track data associated with files and client activity: logs, user permissions, metadata, configuration values, and so on.
- Helix *clients* provide an interface that allows you to check files in and out of the depot, resolve conflicts, detect threats, track change requests, and more.

Helix clients include a command-line client, a graphical user interface client, and various plugins that work with commercial IDEs and productivity software. A Helix server can provide services to a mix of Helix clients.

You also use Helix clients to manage a special area of your workstation called a workspace. Directories in the depot are mapped to directories in your workspace, which contain local copies of managed files. You always work on managed files in your workspace:

1. You check the files out of the depot (and into your workspace).
2. You modify the files.
3. You check them back into the depot.
4. If the changes you try to submit conflict with changes that other users, working in parallel with you, have already submitted, you must resolve conflicts as needed.

The next figure shows the mapping between depot files (shown on the left) and workspace files (shown on the right). Until files are checked out from the depot, they remain as read-only in the workspace. To have Helix update your workspace so that it reflects current work on the depot, synchronize your workspace to the depot by getting the latest revision of the files.



We have explained about checking files in and out of the depot, suggesting that single files may be checked in and out. In fact, the means we use to check files in and out of the depot is the *changelist*. A changelist must contain at least one file and may contain tens of thousands. A changelist is numbered and allows you to track all changes with respect to the contents of the depot: file modifications, the addition of a file, or the deletion of a file.

The changelist is the simplest way to organize your work. A changelist also represents the atomic unit of work in Helix: if a changelist includes several files, changes for all the files are committed to the depot or none of the changes are. For example, if a network connection between the client and the server fails during changelist submission, the entire submit fails.

Multiple user access to a set of files

Version control systems must address the fundamental need for multiple users to work on the same project simultaneously. Helix offers two ways to do this:

- **File locking:** Helix locks a file while someone is working in it. This means that no one else can alter the file until the work is done and the file is unlocked.
- **Branching and merging:** By branching streams and then merging them later, multiple users can work on the same files simultaneously. See [“Balancing stability and innovation: the mainline model” on page 4](#).

Balancing stability and innovation: the mainline model


So far, we have explained how version control systems handle the problem of different users working on the same file at the same time: when the different copies of the file are checked in, all but the first user must resolve changes by deciding which changes will be preserved in the latest version of the file.

This is the simplest use case for parallel development. More complicated cases arise when large development projects require many people to work in parallel both because they must support multiple releases and because they involve multiple functional teams working together to create the desired product.

For example, a game development company depends on the combined efforts of artists, musicians, programmers, testers, and build engineers to create a release. One way to organize this effort is to split the set of files that make up the project into multiple parallel branches, allowing development and testing to occur along each branch. Integration then occurs across branches to promote everyone's work into a releasable product.


Consider the following cases:

- Extremely short release cycles, such as occur with fast changing web content, require overlapping cycles of testing and release.



To handle this case, we move code through the branches shown above. Development occurs along the development line. When a milestone is reached, code is copied up to the QA line where it is thoroughly tested. After it clears all tests, it is copied up to the Beta test line where it is subjected to real-world use. Having satisfied Beta users, it can then be copied up to a release line. Note the purple dashed lines, which indicate the flow of code as it is copied up through the branch hierarchy.

- Unforeseen delays cause some features under development to miss deadlines. Not wanting to imperil the project as a whole, the project is released while development continues for the laggard components.

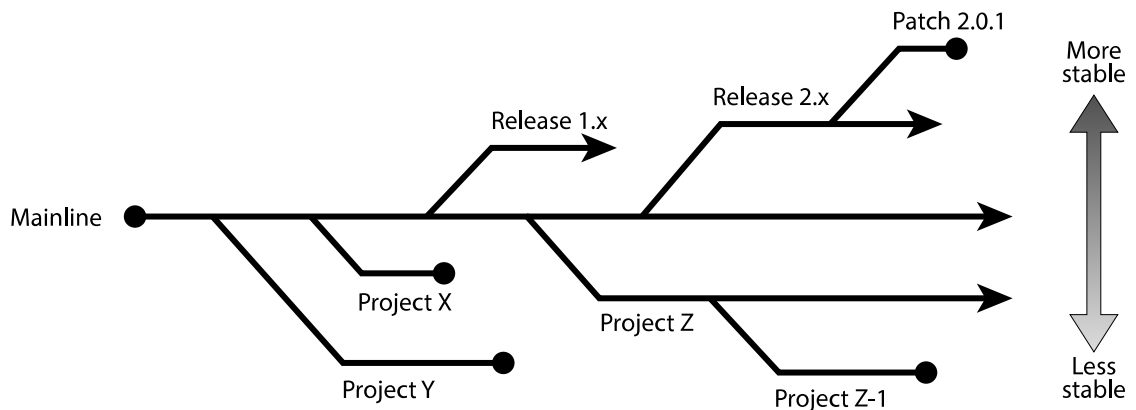


In the model shown above, projects X, Y, and Z have been branched off from the development mainline and worked on independently. When work has completed and projects X and Y have


passed development tests, they are copied up to the mainline. Project Z however could not be completed and work continues on that project without putting the larger release at risk.

What does branching have to offer? It allows us to balance the need for stability with the need for innovation. On the one hand, we have release branches that hold the most tested and stable code; on the other hand, we have development branches that allow for experimentation and exploration without putting the release at risk.

Branches can be organized in a variety of ways: you can create branches for different platforms, you can create branches along organizational lines, and so on. One common model used for product development is the Mainline model shown below:



The Mainline forms the trunk from which release branches and development branches are created. Each branch normally contains a subset of the Mainline files. Release branches might contain fewer files because files needed for testing are excluded; development branches might contain a different subset of files because the projects they represent focus on discrete product features. In the Mainline model, the “up” direction indicates increased stability or confidence.

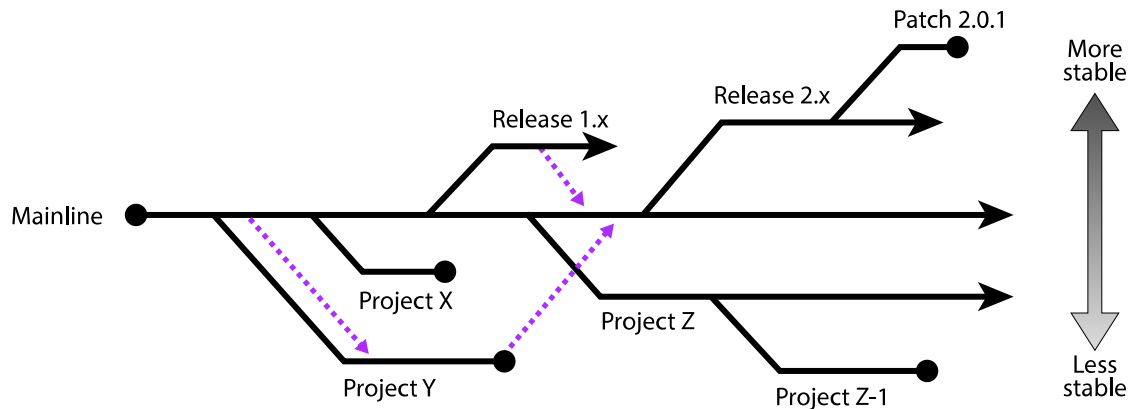


When you create branches, you are free to integrate changes in any direction you like. Unfortunately, this can lead to big problems if you inadvertently integrate untested changes into an otherwise stable branch. For this reason, in addition to defining branches that isolate changes, the Mainline model is most useful when it can implement some protocols that limit what changes can be made and in what direction.

Streams

Helix streams implement the Mainline model, adding intelligence that determines what changes can be made and in what order they must be made.

Let's look at the Mainline example again and add some information to indicate flow of change:

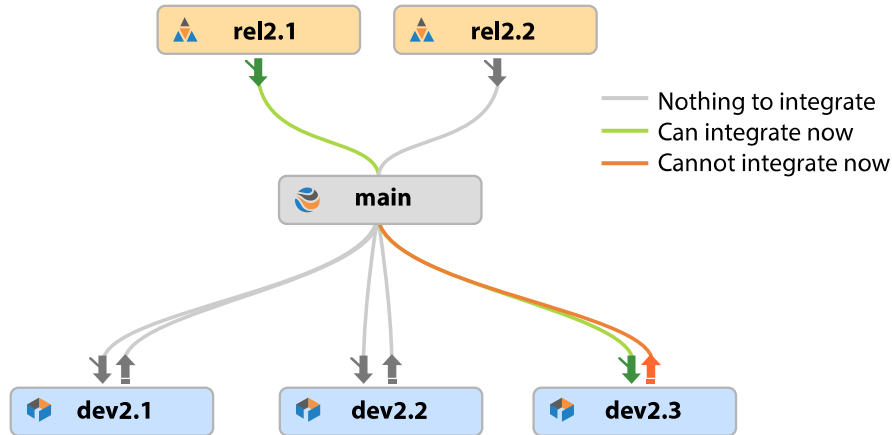


- Project Y has been branched from Mainline; work and testing continues until it is complete. It is then ready to be copied back up to the mainline. However, while development has taken place in Y, Mainline has continued to change. Before we can copy the contents up to the Mainline, we need to make sure that Project Y files reflect changes that have been made in Mainline; we merge those changes into Project Y before we copy Project Y files to Mainline.
- A bug is found in Release 1.x. The bug is fixed and tested. We now want the bug fix in Mainline, so we merge files from Release 1.x down to Mainline. We do not copy anything up because Release 1.x should not include any features added after it has been branched.

The Mainline model arranges branches in terms of stability: the most stable branches are at the top; the least stable branches are at the bottom. The flow of change needs to support this model by merging changes down and copying up.

Typically, when you work with streams, you define and populate the mainline first. You then create development streams and release streams as children of the Mainline stream. The type of a stream and its relationship to other streams determines what sort of changes can be made and in what order they are made.

Rather than using a timeline, the streams GUI—found in the Perforce Visual Client (P4V)—represents related streams as shown below:



The children of Main are shown both above and below Main. Release type streams are at the top; development and task streams are at the bottom. Stability grows as streams near the top of the diagram. The direction and color of arrows linking streams indicate both the direction of flow and the order of flow.

When you create a stream, you specify its type, its relationship to other streams, and how files are to be treated for merging and branching. The information you provide is then used by the streams application to encourage good behavior.

Streams provide visual clues for where and how to branch and merge. They guide behavior that supports stability and innovation. Using streams eliminates much of the work needed to define branches, to create workspaces, and to manage integrations.

An additional advantage of using streams is that when you switch from one stream to another, the contents of your workspace are updated automatically to reflect the contents of the current stream.

Streams automate branching, but you do not have to use them. You can create your own branches and manage them as you see fit. Custom branching gives you finer grained control but you lose the convenience of built-in flow control and workspace updating.

For information on streams, see the [P4 User Guide](#).

Organizing your work: jobs and labels

In addition to using changelists and streams to organize your work, you can use two other methods: jobs and labels.

- *Jobs* provide lightweight issue tracking that integrates well with third party defect tracking and workflow systems. They allow you to track the status of a bug or an enhancement request. Jobs have a status and a creator and are associated with changelists that implement the bug fix or the enhancement. An administrator can customize the type of information tracked by jobs add more fine grained status values, or define additional fields for information to be tracked: which customer the enhancement is for; what was done to test the fix, and so on.

You can integrate the jobs function with third-party defect tracking and workflow systems. For more information, see the [Defect Tracking Gateway](#) page.

- *Labels* are sets of tagged file revisions that allow you to handle a heterogeneous group of files as one unit. While a changelist refers only to the contents of a given set of files at the time they were submitted, a label can refer to a group of file revisions from different points in time. You might want to use labels to define the group of files contained in a particular release, to sync a set of files, to populate a workspace, or to specify a set of file revisions to be branched. You can also use a label as an alias for a changelist number, which makes it easier to remember the changelist and easier to refer to it in issuing commands.

For information about jobs and labels from a user's perspective see the [P4 User Guide](#).

For information about managing jobs and labels, see the [Perforce Server Administrator's Guide: Fundamentals](#).


Working together and working apart: centralized and distributed development

We mentioned earlier that version control systems can implement either a centralized model or a distributed model:

- *Centralized version control systems* use a single repository from which users check out one or more files to work on in their local directories.
- *Distributed version control systems* allow users to host repositories locally, check out entire repositories with history, work independently of one another, and combine their work through merging when necessary.

Helix supports either model, as well as a hybrid of the two.


In the centralized model, clients work with a depot on a shared server. A mapping of files from the depot to their workspace determines which files they are able to work with in their workspace. The following figure illustrates the centralized model:



Users check files out of the same depot, work on them, and check in their changed files. If multiple users work on the same file, they use merging and conflict resolution to make sure the resulting version is satisfactory to all authors. Although users can disconnect from the shared server and continue to work on the files in their workspace, some manual work is required to sync back to the server and to check in files when the user reconnects. For information on working with Helix using this model, see the [P4 User Guide](#).

In the distributed model, users work with a depot on personal servers that are then connected to a shared server. The depot on their personal server might contain a subset or the entire set of files on the shared server. Each user can work disconnected from the shared server but still be able to access all the files in their workspace and place these under version control using their personal server. Each user can access the entire history of a file locally, can rewrite and revise history, and can manage the files and streams on their local machine without interacting with the shared server.

The following figure illustrates this model:



When users decide to share their code or digital assets with other users, they connect to and then push their content to the shared server. This allows other users to fetch that content from the shared server and work with it on their own personal server. Users might need to merge content before pushing if their changes conflict with changes made by others.

In addition to supporting these two models, Helix also allows for a hybrid architecture in which some users connect directly to the shared server while others connect to personal servers that are connected to the shared server.

For more information about distributed development and file management, see [Using Helix for Distributed Versioning](#).

Collaborating within a Git Ecosystem

Git users who want to take advantage of Helix enterprise-level features can work seamlessly with Helix Versioning Engines using Helix Git Fusion.

Helix Git Fusion is a Git remote repository service that uses the Helix Versioning Engine as its backend. Git users interact with Git Fusion as they do with any other Git remote repository (repo), issuing standard Git commands to clone repos and transfer data. When a Git user pushes changes, Git Fusion translates those Git changes into Helix changes and submits them to the Helix depot. When a Git user pulls changes, Git Fusion translates the pull request into p4d commands to download changes from the Helix depot.



Performance, scaling, and high availability

Version control systems are key to managing large projects: with Helix, “large” can be large indeed. With enterprise-level features that you can use to fine tune and improve performance, Helix lets you scale your system to accommodate a global workforce, and to automate failover for a highly available system. For example, Helix can accommodate the needs of a gaming development company whose files might take up hundreds of terabytes or even petabytes of data; or it can support the work of a software company, whose activity level includes massive automated testing as well as focused, analytic bug fixing and tracking work.

To support these tasks, Helix uses the following additional server types:


- *Proxies* are used where bandwidth to remote sites is limited; they mediate between remote clients and the versioning service. By caching frequently used files, the proxy reduces demand on the server and keeps network traffic to a minimum.
- *Brokers* mediate between clients and servers to implement policies that solve routing or security problems. For an example of how brokers can be used, see [“Data center scaling and high availability” on page 12](#). The cluster router is a broker used to filter and load-balance incoming client requests.
- *Replicas* duplicate server data. They can be used to provide a warm standby server or to reduce load on a primary server.

The following sections explain how these servers are used singly or how they are combined to provide enterprise-level performance. For complete information about using proxies, brokers, and replicas, see [Perforce Server Administrator’s Guide: Multi-site Deployment](#).

Using proxies to improve performance

To improve performance for users accessing a shared Helix repository across a WAN, you can configure a proxy on the side of the network close to the users and configure the users to access the service through the proxy; then configure the proxy to access the master Helix service.


The following diagram illustrates a typical proxy configuration:



In this configuration, file revisions requested by users at a remote development site are fetched first from a shared server (**p4d** running on Central) and transferred over a relatively slow WAN. Subsequent requests for that same revision, however, are delivered from the Helix Proxy, (**p4p** running on Outpost), over the remote development site's LAN; this architecture reduces both network traffic across the WAN and CPU load on the shared server.

Using a replica for disaster recovery

Replication is the duplication of server data from one Helix server to another. The replicated server is called the master server; its replica is called a replica server. You can use replication to provide a warm standby server, or to reduce load and downtime on a primary server when performing builds. The following figure shows how you set up a replica to provide a warm standby to aid in disaster recovery.




When you create the replica, you specify which server it should get its data from. The replica then periodically updates itself by copying files and metadata from the master. If the master fails, all you need do is reconfigure the replica to be the new master and then reconnect clients to communicate with it.

Edge servers and workspace servers, described in the following sections, are special examples of replica servers.

Commit-edge architecture

This architecture supports optimal performance for geographically distributed work groups. At a minimum it is made up of the following kinds of servers:

- A *commit* server that stores archives and metadata.
- An *edge* server that contains a replicated copy of the commit server data and a unique, local copy of some workspace and work-in-progress information. This server can handle read-only operations and operations that only write to the local data. You can connect multiple edge servers to a commit server as shown in the next figure.



Since an edge server can handle most routine operations locally, the edge-commit architecture offloads a significant amount of processing work from the commit server and reduces data transmission between commit and edge servers. This greatly improves performance.

For more information about these options, see [Perforce Server Administrator's Guide: Multi-site Deployment](#).

Data center scaling and high availability

Cluster architecture offers superior performance and automated failover.

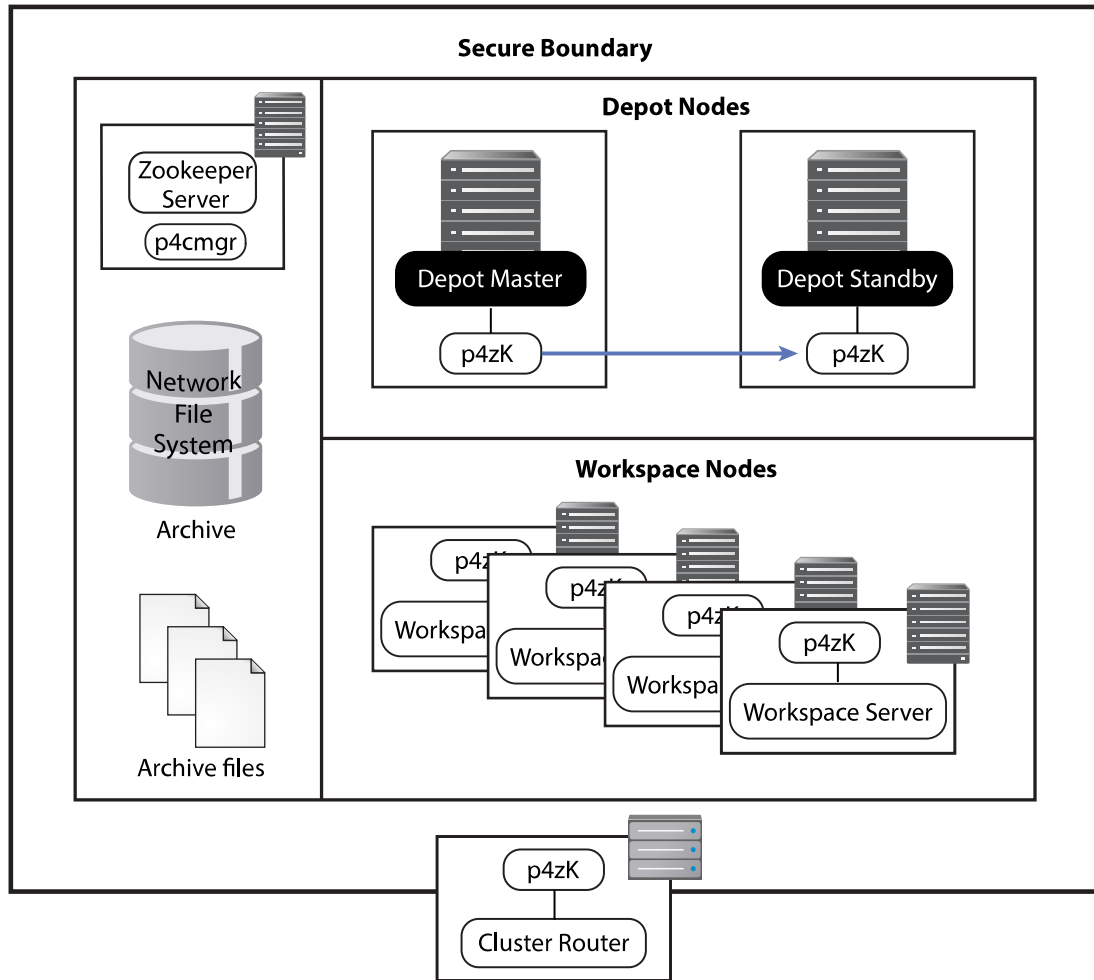
A Helix cluster includes the following elements:

- A depot master that records changes to revision data.
- A depot standby that is automatically synchronized to the depot master and that assumes the role of the depot master if the depot master fails.
- A fault tolerant shared file system that is accessible to all cluster nodes.
- Optional workspace servers that manage client data and help distribute the load of clients.
- A router that dispatches user requests to the workspace servers or directly to the depot master if no workspace servers are used.

A Helix cluster implements a highly available, horizontally scalable commit-edge architecture.

A server cluster uses an Apache ZooKeeper service to automate the process of failure detection and of role migration from the failing depot master to the depot standby. The fault-tolerant ZooKeeper service maintains cluster state information and monitors the health of the cluster nodes. The cluster uses the **p4zk** process, which supports each server in the cluster, to register the existence of the server and to respond to cluster state changes.

Server cluster topology is shown in the following figure:



For more information about Helix clusters, [Perforce Server Administrator Guide: Cluster Management](#).

Securing the system

Independent of the Helix architecture you use, secure communication is guaranteed both with respect to communication between clients and servers as well as communication between servers.

- User authentications can be done using passwords or tickets, and the strength of the password can be defined by an administrator. Users can be authenticated against an Active Directory or LDAP server, or against an internal Helix user database.
- Communication between clients and servers can be secured using the SSL protocol, which you specify when connecting to the server.
- Communication between servers in a distributed environment can be secured using a trust file and by setting permissions for the service users that own the different servers in the environment.

In addition to user authentication, digital assets are further secured by a protection scheme to prevent unauthorized access. Protections determine which Helix commands can be run, on which files, by whom, and from which host. This scheme provides the finest grained control possible.

Finally, Perforce offers Helix Threat Detection, which quickly and accurately identifies internal and external security risks to intellectual property managed in Helix. See [Chapter 2, “Clients, Collaboration, and Security” on page 15](#) for details.

For more information on security see the [Perforce Server Administrator’s Guide: Fundamentals](#) and the [Perforce Server Administrator’s Guide: Multi-site Deployment](#).

In addition to its core offering of products—the Helix Versioning Engine and the Helix client applications—Perforce offers a number of tools and software packages that allow you to analyze your data, enhance collaboration on your coding projects, and seamlessly integrate Helix into interactive development environments.

Client applications

Helix client applications include the following:

- A command line interface available for all platforms.
- A GUI interface for Mac OS X, UNIX, Linux, and Windows.
- Integrations, or plug-ins, that work with commercial IDEs and productivity software.

Analytics tools

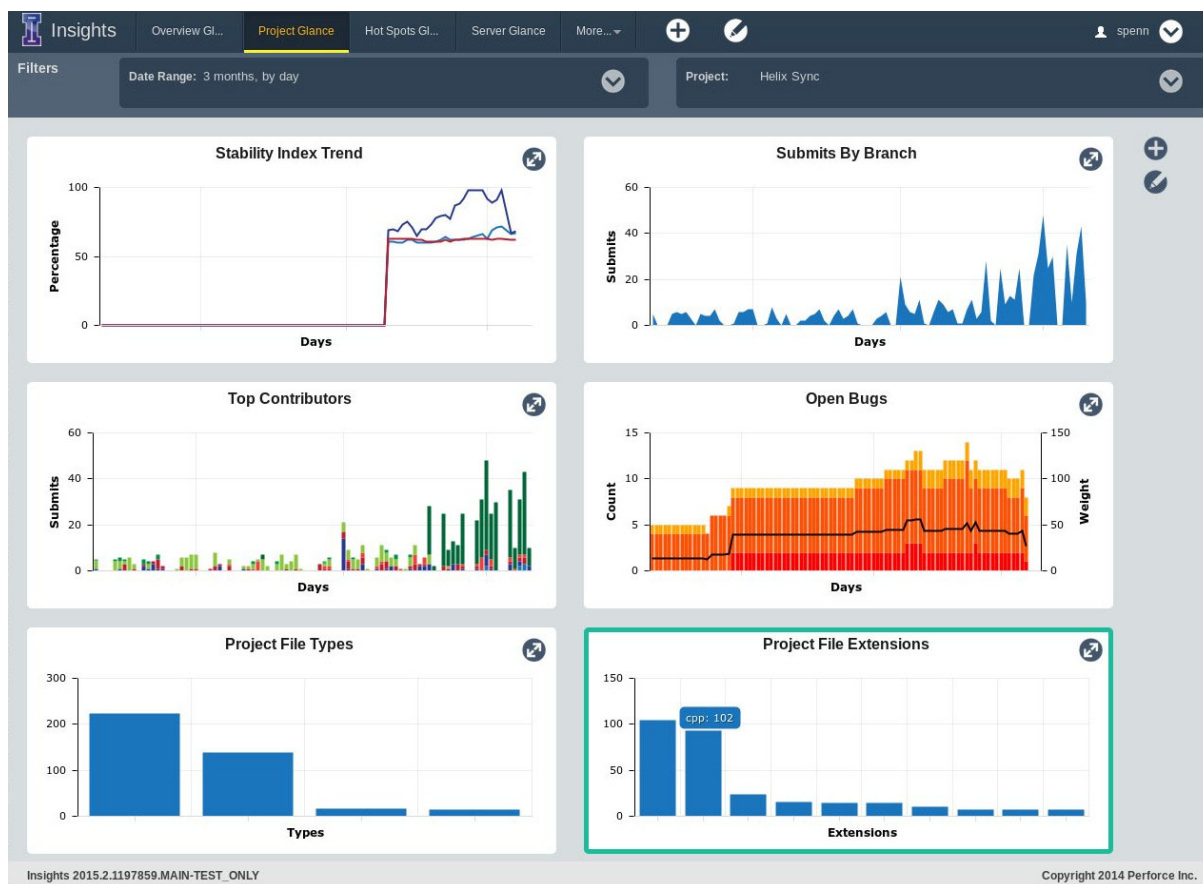
Perforce offers two tools for analysis: Helix Insights and Helix Threat Detection.

Helix Insights

Helix Insights is an analytics tool that presents vital information about how your projects, codelines, and teams are performing. Insights draws on data from the Helix Versioning Engine and logs, displaying it in graphical form so that you can quickly understand what's going on.

Insights reports on project stability by examining bug and submit activity. It also lets you visualize metrics for change, contributions, and issue tracking.

The following figure shows some of the information and analysis provided by Insights.



Using Insights, you can customize analytics dashboards to get a view into your system and analyze any information from Helix or third-party sources. You can also employ the Insights Analytics API to set custom metrics.

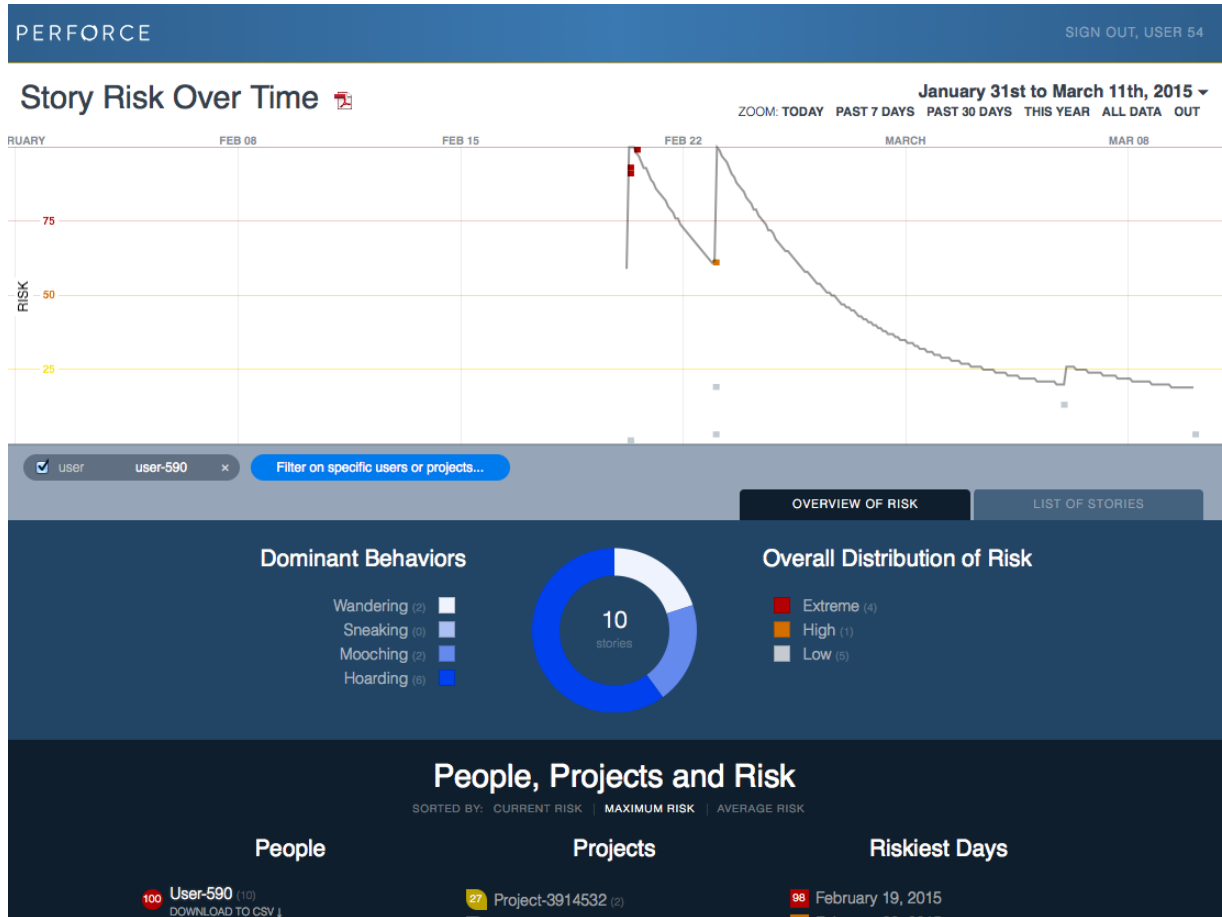
Finally, you can export Insights data for use in third-party analytics tools (for example, Splunk or Tableau).

Helix Threat Detection

Helix Threat Detection quickly and accurately identifies internal and external security risks to intellectual property managed in Helix. Applying advanced behavioral analytics to interactions with source code, product designs and related assets stored in Helix, Helix threat detection proactively surfaces threats in real time and assesses the risk level for each threat event.

Helix Threat Detection effectively distinguishes actual threats from “noise.” It features an intuitive graphical interface and plain-language reporting, linked to analytics and machine learning, that enables you to successfully protect valuable intellectual property assets.

The following figure illustrates how Helix Threat Detection displays possible thefts over time.



Collaboration

Using the Helix Swarm and Helix GitSwarm collaboration tools, your team can work together effectively on complex and extensive sets of data and code.

Helix Swarm

Helix Swarm is a web-based collaborative tool for small to medium teams sharing code and assets. It provides functionality to review, comment upon, and approve changes to code and assets. Using Swarm, you can customize projects, memberships, and reviewers.

Helix Swarm helps you review code before or after committing it, bring continuous integration into the review, and merge work that passed the review. It integrates cleanly with JIRA and Helix Jobs, and provides an API for the development of custom extensions.

Helix GitSwarm

Helix GitSwarm is Perforce's code review solution for Git users. Built on the open source product GitLab, it includes Helix-specific enhancements to support both Git and Helix code reviews. It allows users to push their code, review others' code, comment, and track issues.

Helix GitSwarm offers a pure Git-based workflow and an intuitive web-based UI. In addition to keeping your Git-based projects in Helix, GitSwarm allows you to mirror content from repositories in other popular tools such as GitHub and GitLab. It also allows you to restrict who can see what projects, as well as providing other security mechanisms. In addition, it provides issue tracking and project Wikis.

Integration

Helix includes two types of integration plugins: plugins for integration into interactive development environments and plugins for build and reporting integrations.

Interactive development environment integrations

- The P4VS plugin embeds the power of Perforce Helix features into the Microsoft Visual Studio IDE.
- The P4Eclipse plugin integrates the strengths of Perforce Helix with Eclipse's powerful IDE.
- P4Connect, the Perforce Helix Plugin for Unity, enables you to perform Perforce Helix operations directly from within Unity.

You also have access to many community-built integration plugins, which are summarized on the Perforce website, on the [Helix Clients & Integrations page](#).

Build and reporting integrations

Helix offers three main build and reporting integrations: one for the Jenkins build system, one to query from Helix into a SQL database, and one providing an FTP interface to the Helix Versioning Engine:

- The Jenkins integration allows Jenkins users to use Helix repositories directly. The plugin makes it simple to pull code, automate reviews, apply labels, and so on.
- The P42DB integration replicates Helix metadata to open source and commercial SQL databases.
- The P4FTP integration allows FTP-based HTML authoring tools such as Dreamweaver to access files in Perforce Helix depots.

Helix provides a number of features to customize and extend Helix to build applications.

APIs

Helix includes eight APIs for customizing and extending Helix to build applications. They are summarized in the following table:

This API	Serves this function
Helix C/C++ API	Enables you to create applications that interact with end users, send commands to the Helix Versioning Engine and process data returned from the versioning service
Helix P4API.NET API	Allows you to access Helix version management services from within a .NET program
P4Java	Enables Java applications to access the Helix version control system
P4Perl	Provides an object-oriented Perl API to the Helix version management system
P4Ruby	An extension to the Ruby programming language, it allows you to run Helix commands from within Ruby scripts, and get the results in a Ruby-friendly format
P4Python	The Python interface to the Helix Versioning Engine, it enables you to write Python code that interacts with a Helix Versioning Engine
P4PHP	The PHP interface to the Helix Versioning Engine, it enables you to write PHP code that interacts with a Helix Versioning Engine
P4ObjC	Helix's API for Objective-C enables you to write Objective-C code that interacts with a Helix Versioning Engine

See [Chapter 5, “Documentation and Other Resources” on page 23](#) for pointers to further information about these APIs.

Triggers

Triggers allow you to extend or customize Helix functionality, changing existing behavior or creating new behavior in the Helix system. They are user-written programs or scripts that are called by a Helix Versioning Engine whenever certain operations (such as changelist submissions, changes to forms,

attempts by users to log in or change passwords) are performed. If the script returns a value of **0**, the operation continues; if the script returns any other value, the operation fails.

Consider the following examples:

- You can use a pre-submit trigger to ensure that whenever **file1** is submitted in a changelist, **file2** is also submitted.
- You can use form triggers to generate a customized default workspace view when users run the **p4 client** command, or to ensure that users always enter a meaningful workspace description.

See [Perforce Server Administrator's Guide: Fundamentals](#) for further information about triggers.

Helix Broker

The Helix Broker enables you to implement local policies in your Helix environment by allowing you to restrict the commands that can be executed, or redirect specific commands to an alternate (replica or edge) Helix Versioning Engine.

The Helix Broker is a server process that mediates between Helix client applications and Helix services, including proxy servers. For example, Helix client applications can connect to a proxy server that connects to the broker, which then connects to a Helix Versioning Engine. Or, Helix client applications can connect to a broker configured to redirect reporting-related commands to a read-only replica server, while passing other commands through to a master server.

See [Perforce Server Administrator's Guide: Multi-site Deployment](#) for further information about the Helix Broker.

The Ace Rocket company develops software to manage the production and distribution of their rockets and rocket launchers. They have a variety of technical needs, all of which can be met by members of the Helix product family. The following sections describe these needs and the Helix products that address them.

Software development

Ace Rocket uses a variety of programming languages to develop the components and products that they build. Because they use open source projects, they have numerous developers that use Git. For their internal development teams, they prefer the flexibility of Helix Versioning Engine. Using Git Fusion, Ace Rocket can track their Git commits in the Helix Versioning Engine. This makes the Helix Versioning Engine a single source of truth for all of their assets and provides an unbroken history going all the way back in time.

Ace's software developers use both Swarm and GitSwarm to collaborate and communicate about the code they write. Ace Rocket's Windows developers take advantage of the native Visual Studio integration (P4VS), which makes it easy for them to kick off code reviews from within their IDE. Ace's Java developers get the full power of Helix inside of their IDE of choice: Eclipse.

Using the Jenkins plugin, the Ace Rocket testing team continually tests code as it is checked into a single place by both P4 and Git developers. The Ace Rocket management team uses Insights to obtain analytics for code activity and quality.

Digital asset management

The Ace Rocket Company tracks design versions of their products using Helix. They track versions of their CAD/CAM files as well as large test datasets. Ace takes advantage of Helix's file locking to ensure that multiple people do not attempt to change a model at the same time.

Using configurable storage, Ace Rocket chooses to limit its storage consumption to only 20 versions of each file. Using Helix's multi-site deployment architecture, Ace seamlessly replicates these large files across the globe to improve access times for their engineers.

Using P4GT, Ace Rocket's modelers interact with Helix easily from within their favorite modeling tool.

Hybrid product development

With both source code and CAD/CAM files tracked in Helix, Ace Rocket can easily track dependencies between their code and their models. Developers take advantage of Helix's fluid distributed workflows, while Ace's model engineers use locking to obtain the tight control they need to collaborate with each other.

Having both the source code and modeling assets in one repository lets Ace Rocket use Helix's theft detection feature to track all of the company's intellectual property in a unified fashion.

Ace Rocket's IT team takes advantage of the numerous Helix APIs to build custom tools that meet their unique needs and to integrate Helix into all of their systems. In addition, they enjoy the benefits of Helix's native Active Directory support to make it easy to manage all of their users.

Ease of use for Helix administration

Ace Rocket employs an outside agency for contractors. They want to give these contractors access to only a certain subset of the depot, while they want to give in-house developers full access. Because Helix Versioning Engine permissions are highly configurable, it's easy for an administrator to use finely grained access control to manage access for different types of users.

Growing into the future

As Ace Rocket grows and develops offices in Shanghai, Buenos Aires, and Johannesburg, they deploy edge servers and replicas to support thousands of developers in these remote offices. Both developers and digital asset managers in remote offices benefit from having local access to their software or assets.

Helix is keeping one unified view, worldwide, of digital assets content and history so the distributed teams can collectively work as one.

This chapter describe the resources available for learning more about Helix.

To learn more about Helix

To obtain online help from within all Helix applications:

- Use the help menu from within graphical Perforce applications
- Type **p4 help** from the command line for help with the Command-Line Client

Documentation for Helix is available on the web at: <http://www.perforce.com/documentation>

The Helix Knowledge Base. A complete list of articles is available: <http://answers.perforce.com>

The Helix Forums are a place for users to ask questions and to hear from other users: <http://forums.perforce.com/>

The **perforce-user** mailing list is mirrored on the Helix Forums: <http://maillist.perforce.com/mailman/listinfo/perforce-user>

Perforce support personnel are available for email and telephone support. Perforce also offers training, consulting, and other professional services. For details, see: <http://www.perforce.com/support-services>

Documentation links

The following table lists documentation available for all Helix products:

For specific information about...	See this documentation
Using the command-line interface to perform software version management, working with Helix streams, jobs, reporting, scripting, and more.	Helix Versioning Engine User Guide
Basic workflows using P4V, the cross-platform Helix desktop client.	P4V User Guide
Working with personal and shared servers and understanding the distributed versioning features of the Helix Versioning engine.	Using Helix for Distributed Versioning
Installing and working with Helix Swarm, the code review and collaboration tool.	http://www.perforce.com/perforce/doc.current/manuals/swarm/index.html
p4 command line flags and options (reference)	P4 Command Reference , p4 help
Installing and administering the Helix versioning engine, including user management, and security settings.	Helix Server Administrator Guide: Fundamentals

For specific information about...	See this documentation
Installing and configuring Helix servers (proxies, replicas, and edge servers) in a distributed environment.	Helix Server Administrator Guide: Multi-site Deployment
Installing and administering a Helix server cluster for high performance and automated failover.	Helix Server Administrator Guide: Cluster Management
Helix plug-ins and integrations	IDEs: Using IDE Plug-ins Defect trackers: Defect Tracking Gateway Guide Others: online help from the Helix menu or web site
Developing custom Helix applications using the Helix C/C++ API	C/C++ API User's Guide
Working with Helix in Ruby, Perl, Python, and PHP	APIs for Scripting

Download URLs

To download Helix products, visit: <http://www.perforce.com/downloads>

Videos

To view tutorials explaining how Helix products work, visit: <http://www.perforce.com/resources/tutorials>

Perforce software includes software developed by the University of California, Berkeley and its contributors. This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>).

