



WHITE PAPER

Helix Core for Subversion (SVN) Users

Introduction

This guide is designed to help users familiar with SVN more quickly adopt Perforce Helix Core version control. Use this guide to:

- Learn the similarities and differences between SVN and Helix Core.
- Discover new workflow capabilities with Helix Core that are not present in SVN.
- Adapt development processes to use Helix Core to its full potential.

While this guide is quite comprehensive, it shouldn't be considered technical documentation for the functionality described.

This information in this guide is based on SVN 1.9 and Helix Core 2018.1

Contents

Conceptual Differences Between SVN and Helix Core	3
Repos vs. Depots	3
Authentication	3
Workspaces	3
Why Workspaces?.....	4
Files Under Source Control.....	4
Pending Changes and Shelves.....	5
Storing a Change	5
Labels	5
Identifying File Revisions.....	6
Branching and Merging	7
Access Permissions	8
Replication	9
Code Reviews with Helix Swarm	9
Adapting Development Processes	9
Choosing Client Software	9
Setting up the Environment.....	10
Getting Help	10
Creating and Using a Client Workspace	10
Why Explicit Checkouts.....	12
Changelists	13
Migrating SVN Data	14
Helix Core to SVN Conversion Tool	14

Conceptual Differences Between SVN and Helix Core

REPOS VS. DEPOTS

In many situations, an SVN server contains a set of isolated repositories. A client connects to one repository at a time, but it can reference other repositories through `svn-externals`.

The Helix Core server does not have isolated storage. There is no technical restriction in accessing files across depots. Files can be branched and integrated (or merged) from one depot to another, and users can access files from several depots at the same time. Depots are all managed as a collection in the Helix Core server.

With Helix Core, there are several depot types. We won't discuss them all in detail here. By default, there is a *primary depot* named *depot* on every Helix Core installation. Other important types of depots include Stream, remote, spec, unload, archive, tangent, and graph. Each has a specific purpose in building workflows.

Assets being used in active development are managed in three kinds of depots: classic depots, Stream depots, and graph depots. Stream depots are the most commonly used type of depot today. More about Streams later in this document.

Graph depots are depots in which Git files and whole Git repos are managed inside the Helix Core server.

Classic depots provide compatibility with older (pre-2011) versions of Helix Core, and they are a useful store for assets that require minimal workflow.

AUTHENTICATION

Both SVN and Helix Core require named user accounts for all operations.

In Helix Core, users first log in to the server and are issued tickets that are valid for a certain amount of time – 12 hours by default.

Passwords in Helix Core are either administered and stored within the server itself or via an authentication agent such as LDAP or Active Directory (AD). Helix Core also supports multi-factor authentication (MFA).

WORKSPACES

To make a change to a file under source control, SVN users check out a working copy of a repo, or part of a repo, onto their local machine. In SVN, this working copy contains an administrative directory called `.svn` that holds the connection information and state and a pristine copy of all workspace files. The SVN working copy has the same structure of folders and files as the repository.

Helix Core uses a concept called the “client workspace.” From a user perspective, it is used to map files in the server depot to files on the workstation. A client workspace is sometimes interchangeably called a client or a workspace in Helix Core. In fact, the command-line client, P4, allows both names. Both “p4 client” and “p4 workspace” are aliases.

Helix Core doesn't use hidden files to manage local assets. Instead, the state of each workspace is kept on the server. A client workspace consists of a:

- Unique name
- Root directory
- Set of options
- View

The workspace view specifies the set of files in the depot that should be mapped to the local machine, because no one wants all the files that are available on the server. This selects just the desired set of files. Note that content can be mapped from multiple depots into a workspace.

In addition, the structure of the workspace can be quite different to the structure of the repository if required. For example, files or folders can be mapped to different levels in the workspace. This flexibility can be very useful in practice. For example, this allows artists and software developers to work on their own separate content for a game, and a build workspace can have everything.

Many teams use a branching feature of Helix Core called Streams (discussed later in this document) for development. This can automatically generate a workspace view, or you can generate the view using scripts or template workspaces. Others let their users generate their own workspaces.

Workspaces are a powerful tool for automation, in addition to being used by developers creating code. There are two special types of workspaces:

- Partitioned workspaces
- Read-only workspaces

Both are designed to enhance performance of CI/CD automation processes by streamlining the handling of files in the DevOps pipeline. This also

helps maintain optimal performance by preventing database table fragmentation.

WHY WORKSPACES?

Workspaces are a powerful concept. They are flexible enough to allow different types of users to define the layout of files from depots to suit their working needs. For example, developers, designers, QA automation engineers, and DevOps pros can define their own views for their own unique needs.

Workspaces in Helix Core are not only used to map the set of files a user wants to work with; the server can also track exactly which revisions of each file the user has synced. This approach allows the system to send the correct set of files to the user when syncing without having to scan the file system first to see which files need to be updated. With a large number of files, this can be a huge performance win. This is also very popular in industries that have very strict auditing rules. Helix Core admins can easily track and log who has synced which files.

A powerful productivity advantage of being able to map an assortment of modules to one workspace is the ability to easily modify multiple code modules in one check-in, guaranteeing that anyone with a similar client view who syncs to a check-in will have all the code in the correct state.

FILES UNDER SOURCE CONTROL

SVN stores text and binary files, as well as empty directories.

Helix Core also stores text and binary files, but the server does not version empty directories

directly. If coding practices require empty directories in the repository, place an empty hidden file into it (e.g., “.dir” or “.p4ignore”).

Helix Core distinguishes among text, binary, and Unicode-encoded files as well as symbolic links. Additionally, each version of a file can have extra attributes such as *executable* or *exclusive checkout*.

PENDING CHANGES AND SHELVES

SVN (since version 1.5) has the concept of a local changelist attached to a working copy that allows users to group changes under an arbitrarily named change.

Helix Core stores changelists on the server. It has the concept of a default pending change to which all changes in a workspace are attached. Each file operation such as adding, editing, or deleting a file is automatically attached to the pending changelist.

A pending change can also be explicitly saved; it is then given a unique change number. Each work-space can have unlimited number of pending changes.

Similar to SVN, a file can be in only one changelist at a time on the client.

But Helix Core also has a feature called “shelve.” This transfers the changed content to a central server where it can be inspected, compared, and reviewed by all users. For example, when a developer believes they have completed a change, but is submitting it to a code review. She doesn’t want to submit it to the server (i.e., commit it) prior to the reviewers approving her changes.

Another great use for shelving is to keep several states of a file in the same workspace at the same time, since each file can only be opened once in a pending change for a workspace.

STORING A CHANGE

A change in SVN is committed with a change comment and is then visible to all other users. It is possible to only commit a single named changelist. Once committed, the change is given a repository-unique, strictly increasing number called a *revision*.

The Helix Core operation to commit a change is called a *submit*. In most cases, users will submit all changes in their current pending changelist. As in SVN, each submitted change is assigned a unique, strictly increasing number, but it is called a *change* or *changelist*.

In Helix Core, a changelist can include a description and, optionally, a list of *jobs* that the change fixes. Helix Core has a built-in system for handling work items and tracking status milestones. Jobs are often used to allow developers to get tasks from Jira, Helix ALM, and other issue management systems, and then to report the status of developer’s progress back to such systems, automatically.

SVN’s revision and Helix Core’s change both identify the state of all files in the repository at a given time, not just a particular change.

LABELS

Labels in SVN are called *tags*. They are represented as a file path, and by convention only, stored in a subdirectory called “tags.”

Helix Core's label feature assigns a name (label) to a set of files using the file specifiers described below. Labels are commonly used for automation tasks, supporting CI/CD build, and release workflows. They are also popular with development teams sharing libraries and artifacts. teams sharing libraries and artifacts.

Helix Core stores labels as metadata on the server. There are three kinds of labels:

- *Static* labels resemble SVN's tags in flexibility
- *Automatic* labels are aliases for changes and fixes and have minimal impact on the repository metadata
- *Unloaded* labels act like static labels but store their data in a single file

For example, to access files at a label, a Helix Core user syncs "files/...@labelname".

There are several commands to create and act upon labels, that are discussed extensively in the Helix Core documentation.

IDENTIFYING FILE REVISIONS

SVN identifies file revisions only through its global revision number, or HEAD. A particular file might have been modified only in revisions 1, 15, 31, and 73.

Files in SVN are only identified through their local path. Recursion can be defined through a specified depth.

Helix Core provides users with a comprehensive, cross-platform-capable approach to specifying

files, whether employed at the command line, or via scripting.

Helix Core users specify a file in three ways:

1. Depot syntax — an absolute path including the depot root: //depot/dir/file
2. Client syntax — an absolute path including the client root: //client/dir/file
3. Local syntax — a relative or absolute path of the operating system

In Helix Core, any file can be specified within any command in client syntax, depot syntax, or local syntax with the latter two options being most commonly used.

Depot names and client workspace names share the same namespace; there is no way for the Helix Core server to confuse a client name with a depot name.

Helix Core server's own method of file specification, using Depot syntax, remains unchanged across operating systems.

If a file is specified relative to a client root, it is said to be in client syntax. If it is specified relative to the top of the depot, it is said to be in depot syntax. A file specified in either manner can be said to have been specified in Perforce syntax.

Local syntax refers to filenames as specified by the local shell or operating system. Filenames referred to in local syntax can be specified by their absolute paths or relative to the current working directory. (Relative path components can only appear at the beginning of a file specifier.)

SYNTAX	EXAMPLE
Local syntax	<pre>/staff/user/usercws/file.c usercws/file.c (if current directory is /staff/user) ../usercws/file.c (if current directory is /staff/user/project)</pre>
Depot syntax	<pre>//depot/source/module/file.c</pre>
Client syntax	<pre>//usercws/file.c</pre>

Helix Core file specifiers always begin with two slashes (`//`), followed by the client or depot name, followed by the full pathname of the file relative to the client or depot root directory.

Path components in client and depot syntax are always separated by slashes (`/`), regardless of the component separator used by the local operating system or shell.

Wildcards can also be used to specify files:

WILDCARD	MEANING
*	Matches all characters except slashes within one directory.
...	Matches all files under the current working directory and all subdirectories. (matches anything, including slashes, and does so across subdirectories)
%%1 - %%9	Positional specifiers for substring rearrangement in filenames, when used in views.

Revision specifiers can be used to operate on many files at once, as the name implies: `p4 sync //myclient/...#4` copies the fourth revision of all non-open files into the client workspace.

Table: File Specifier Modifiers

IDENTIFIER	SYMBOL USED	EXAMPLES
Change	@	README.txt@rel2.2
Revision	#	README.txt#14
Label	@	README.txt@rel2.2
Head Revisions	#head	README.txt#head

BRANCHING AND MERGING

SVN by convention stores its branches in a directory called “branches” parallel to the “trunk” directory that contains the mainline. SVN does not keep track of the relationships among branches; it’s up to the team to do so, with a naming scheme or with external documentation.

The real difference between SVN and Helix Core is the way merges are treated. Helix Core uses a separate database table to keep track of every merge and the choice a user made when resolving a conflict on the server. This feature allows Helix Core to make an accurate choice of which file revision still requires merging and the common base of a merge. This minimizes any merge conflicts, even if the merge across branches is only indirectly related.

There are references to “Classic” depots and branching in Helix Core documentation, but one of the most exciting innovations developed by Perforce is the concept/technology of *Streams*.

SVN has no concept of a relationship between branches beyond a naming convention.

In contrast, Perforce Streams have workflows and relationships that makes complex tasks simple, and it minimizes the need for developers to hand process management tasks such as integrating bug fixes from a dev or QA branch to multiple production codelines.

Perforce Streams have several benefits to the user:

- Streams defines the purpose of each branch: mainline, development, task, or release. Streams follow the mainline model — all changes flow toward the mainline, similar to an SVN trunk.
- Each Stream (except the mainline, which is the root) has a parent Stream that defines a clear hierarchy along which changes flow.
- The Stream graph identifies changes that still need to be propagated.
- Client workspaces are locked to a Stream, eliminating the need to set up the view manually. Client workspaces can be switched from one related Stream to another; only the files that differ between these Streams will be updated.

Mainline: A Stream with no parent. Expects merging down from more stable child Streams. Expects copying up from less stable child Streams. Used as the stable trunk of a Stream system.

Release: A Stream that is more stable than its parent. Expects merging down from more stable child Streams. Does not expect copying up from its parent Stream. Useful for ongoing stabilization, bug fixing, and release maintenance.

Development: A Stream that is less stable than its parent. Expects merging down from its parent

Stream. Expects copying up from its less stable child Streams. Does not expect to have more stable child Streams. Useful for long-term projects and major new features.

There are two additional Stream types with special characteristics:

Task Streams are lightweight, short-term branches that can be used for work that affects a small portion of a full project branch. Task Streams enable work to be done privately, let developers switch contexts quickly, and reduce the amount of metadata managed by the Helix Core server.

Virtual Streams provide users with the ability to restrict the workspace view of a real Stream. Virtual Streams act as a filter. They are used to sync a particular set of files, rather than all the files in the Stream view to a workspace.

ACCESS PERMISSIONS

SVN sets its access permissions through an Apache module at the repository or directory level with either full or no access.

Helix Core has a more fine-grained approach. Access permissions are stored in the server's protections table and define list, read, open, and write access. Access permissions are typically defined for a group of users and can be restricted to an individual.

Helix Core has several ways to identify a file revision. A concept similar to an SVN revision is a change, an integer number that is unique to the entire server. Each individual file has its own revision number that starts with 1 when the file is added to the repository. Additionally, a label

may also be used to identify a file revision. also be used to identify a file revision.

REPLICATION

SVN has the `svnsync` command, which creates and maintains read-only mirrors (copies) of its repositories. It works by replaying commits that occurred in one repository and committing it into another. The primary use case for this command is to enable backups for the SVN server.

Helix Core replication is different from other version control solutions, and it's been steadily evolving since 2012. With Perforce federated architecture, also known as "Commit/Edge" servers, each location can have its own server. It features lightweight, intelligent replication, which offers a sharp contrast to copying.

A commit server stores the canonical archives and permanent metadata. This goes in a data center in, for example, a corporate headquarters or in a private cloud. Then, an edge server contains replicated copies of the commit server data and a unique, local copy of some workspace and work-in-progress information. To achieve high performance, edge servers process read-only operations and operations that only write to the local data. Multiple edge servers may be connected to a commit server.

The edge server offloads a significant amount of processing work from the commit server. It also reduces data transmission between commit and edge servers. As workloads grow, additional CPUs and memory can be added, and performance continues to improve in a linear fashion. There's virtually no ceiling to performance improvements.

From a developer perspective, most typical operations (until the point of submit) are handled by the edge server. Read operations, such as obtaining a list of files or viewing file history, are performed locally. In addition, with an edge server, syncing, checking out, merging, resolving, and reverting files are also local operations. Developers don't even know there are multiple servers. To them, it's all transparent so they can focus on creating great code.

CODE REVIEWS WITH HELIX SWARM

SVN does not provide an integrated code review tool, but instead works with tools like Atlassian Crucible, SmartBear Collaborator, and others.

Helix Swarm, which is included with Helix Core, is a scalable code review and collaboration tool for all types of intellectual property. Swarm seamlessly handles 100s or 1000s of reviews at once, regardless of file type or size.

Contributors share files, comment, suggest tasks, vote up or down, and submit final work directly within its web-based interface. Swarm automates the entire process via notifications and makes it easy to monitor progress.

Large, distributed teams use Swarm to support Agile methodologies and CI/CD workflows requiring multiple iterations, incremental delivery, and build automation.

Adapting Development Processes

CHOOSING CLIENT SOFTWARE

SVN itself does not provide an official GUI, but TortoiseSVN, which adds SVN commands

to Windows File Explorer and is maintained by the community, is a popular choice. There are also add-ons for IDEs such as Visual Studio and Eclipse. also add-ons for IDEs such as Visual Studio and Eclipse.

Perforce supplies an official GUI client for Helix Core called P4V that is available on Windows, Mac OS X, and Linux. Perforce also has a Helix Core plugin for Windows 10 File Explorer, that lets users access version control functionality from “right-click” menus in Windows.

Helix Core plugins can be used by designers from within applications on Windows such as Photoshop CC, Autodesk 3DS, and Maya. Additionally, numerous third-party integrations and clients exist.

Both SVN and Helix Core provide rich command-line interfaces (“svn” and “p4,” respectively) that offer access to all functions.

SETTING UP THE ENVIRONMENT

SVN stores its connection information in a hidden “.svn” file. This directory is created when a user checks out files from an SVN server for the first time, and it contains the connection parameters needed to commit changes back.

In Helix Core, users specify the connection parameters directly in any tool such as P4V, the desktop client, or an IDE. From the command line, the connection information is typically defined in the environment or, on Windows, in the registry.

The important variables are:

- P4PORT=my-Helix-server:1666
- P4USER=my-user-name
- P4CLIENT=my-client-workspace-name

If the client machine has only one client workspace, then these three variables are all a user needs.

If a workstation has more than one workspace, the environment variable P4CONFIG can be defined to point to a file usually named “.p4” or “p4config.txt”. Place a text file with this name containing the connection parameters into the workspace root. This technique is closest to how SVN works.

GETTING HELP

Need help understanding a command? Helix Core provides a comprehensive help system, invoked via “p4 help.”

CREATING AND USING A CLIENT WORKSPACE

After connecting to the Helix Core server, the first step is to create a client workspace. This must be done before uploading and submitting any files.

A client workspace has a name, a root directory, a view, and a set of options.

Client Workspace Name

The client workspace name is unique to the Helix Core server. By default, the name matches the hostname of the user’s workstation because each workspace is locked to a host.

Most organizations have an established naming convention for client workspaces (which can be enforced through automation, i.e., triggers). A useful convention could be user.host.project (e.g., myname.macbook.p4python).

Set the P4CLIENT value to whatever name is chosen.

Creating a Client Workspace

To create a client, run the command:

```
p4 client
```

...from the command line after P4CLIENT is set. An editor will open that presents the client workspace information to be edited. Make changes (and save and close the editor). The server will create or update the client workspace.

Alternatively, create the client workspace through a GUI tool such as P4V.

Root

The root of the client workspace is the directory under which all files under Helix Core control will be placed. This is similar to the root of an SVN working copy.

View

The view maps files from Helix Core server to the local drive (and the other way around for new files). Usually, the layout of the files on the server and the client workspace match, but this does not always have to be the case.

In general, a workspace view maps a depot path to a client workspace path. When the server is first created, the first workspace path will most likely look like this:

```
//depot/... //workspace_name/...
```

Here “workspace_name” is replaced with the client workspace name. The three dots are a

wildcard to match all files and directories below the specified path.

Because Helix Core often stores many projects, this wide open view that maps every single file on the server can download many more files than desired, unless it is changed to something like the following (depending on the depot layout):

```
//depot/project/branch/... //workspace_name/...
```

...where “project” is the project name and “branch” is the current branch where work is in progress.

The client workspace mapping is somewhat similar to specifying a subdirectory in the URL when checking out a working copy from an SVN server.

Options

There are several options for configuring a client workspace. For those most familiar with the functionality of SVN, the significant option is noallwrite. More details on workspace options can be found in the [P4 User Guide](#).

WORKING WITH FILES

Active development with SVN has similarities to Helix Core in basic operations working with files. Retrieve a local copy of the files, then edit, build, and test in this environment, and ultimately commit the changes back to the server.

In SVN, all files are always writable in the workspace. Edit the files, and when the changes are committed, SVN determines which files have changed and offers to include them in the commit.

In Helix Core, there is a choice. By default, the files in the workspace are read-only to begin with. The command “p4 edit” explicitly checks out the files. This has several effects:

- The file is made writable in the workspace.
- The file is marked for edit on the Helix Core. This step makes it clear to team members that the file is being editing by another developer.
- If the file has the file type “exclusive check-out,” no other user can edit the file until it is committed or reverted. This is useful for files that cannot be merged (e.g., most binary files).

When the changes are submitted, Helix Core does not have to search the entire workspace for changed files. Instead, it simply looks up changes in the database.

If the SVN mode of working is preferable, a workspace can be set to allwrite by modifying its options. With this setting, all files synced to the workspace are now writable by default and can be modified without checking them out explicitly first, just like with SVN. If an existing workspace is switched to allwrite, keep in mind that previously synced files need to be resynced or use OS methods to make them writable.

Helix Core offers two commands to identify which files have changed: “p4 status” shows which files have been edited, added, or deleted, and “p4 reconcile” adds the changed files to the pending changelist (“p4 status” is an alias for “p4 reconcile n”).

The typical work flow is then:

```
p4 sync          # update a workspace
vi hello.c       # change a file locally
p4 reconcile     # update the pending changelist
p4 submit        # submit the changes
```

“p4 reconcile” can also discover renamed and moved files even if their content has slightly changed. This is particularly important when refactoring Java code because the class name inside the file and the file name itself are linked. In SVN, the lack of this functionality can lead to frustration and loss and productivity tracking down files and changes.

To prevent “p4 reconcile” from adding unwanted or unnecessary files to the server (such as generated object or class files), add these files to an ignore list in a file specified by the environment variable P4IGNORE (e.g., “p4ignore”).

Before switching workspace options, keep in mind that most IDEs and editors integrated with Helix Core will check out the file automatically when the developer starts typing.

WHY EXPLICIT CHECKOUTS?

One reason for using explicit checkouts is that it removes the need to scan files for content changes.

With smaller projects, calculating hashes for each file is fairly cheap. However, many Helix Core teams have millions of files in a workspace and/or have individual files larger than 100MB. Calculating all the hashes in those cases is extremely time consuming.

Explicit checkouts let Helix Core know exactly which files it needs to work with. This behavior is one of the reasons Perforce is so popular in industries that use large files like game studios, movie producers, and hardware.

Another benefit: Explicit checkouts provide a form of asynchronous communication that lets everyone know which files are being worked on, and by whom. It can let team members avoid working in a certain area to prevent a needless conflict, or it can alert an admin or manager to the fact that a new developer on the team has wandered into code that perhaps doesn't need to be edited.

Explicit checkouts also play nicely with the Helix Core concept of pending changelists. Pending changelists are buckets that hold open files, to organize work. Branches are great, but sometimes it is nice to be able to organize work into multiple named changes before actually submitting to the server.

With the Helix Core model of potentially mapping multiple branches or multiple projects into one workspace, pending changelists make it easy to keep separate changes organized.

CHANGELISTS

In SVN, changes can be grouped into local changelists and committed individually. However, in most cases, developers will have simply committed all changed files in the working copy together.

In Helix Core, a changelist is a global object on the server that has several states described next.

Default Pending Changelist

Each workspace always has a default changelist associated with it; there is no need to create it. When a file is checked out without specifying an explicit change, it is automatically assigned to that associated default changelist.

Default changelists are not numbered, but they can be accessed with the name "default."

Numbered Pending Changelist

A numbered pending changelist can be created and can have zero or one or more open files associated with it. The server sets the changelist number when it is created.

Helix Core uses a single atomic counter for all changes: submitted, pending, and shelved. When a new numbered pending change is created, it receives the next available number, just like a submitted change.

In Helix Core, submitted changes are ordered by time — a higher change number implies the change was submitted later. It is likely that other users will have submitted changes after the creation of a pending change and before this change is submitted, so the pending change will probably be renumbered when it is submitted. This is expected behavior. The pending change number is simply there for convenience to be able to update the pending change.

There can be an arbitrary number of pending changes associated with a workspace, and open files can be moved between numbered pending changes and from and to the default changelist.

Note that each file can only be opened once in a pending change for a workspace because the file content is still stored on the local disk. If there is a need to keep several states of a file in the same workspace at the same time, the file can be shelved.

Shelved Changes

Numbered pending changes can be shelved, that is, the content of the open files can be saved on the Helix Core server for safekeeping and sharing. This feature can be used to:

- Backup a change if it is not ready for a submit.
- Stash (as in the Git command) changes temporarily to work on something else.
- Share changes for review or transport between workspaces.

Shelved changes are visible to other team members and can be unshelved in different workspaces, by different users, and even different branches or Streams. Shelves are temporary — they need to be deleted when the change they are associated with is submitted.

Submitted Changes

When submitting a default pending change, a description should be provided. For numbered pending changes, the description already stored will be suggested. Once a change is submitted, its contents cannot be modified further. It is possible, however, to update its description and any associated fixes for jobs.

Migrating SVN Data

All migrations from one version control system to another offer two basic choices:

1. Keep the old repository in read-only mode for reference and only import the head revision into the new tool. This is certainly the fastest and easiest way to migrate data, but it is not applicable if the migration is happening in the middle of an existing project or if there are long-running releases that need to be supported.
2. Migrate some or all history into the new tool. There will always be a mismatch between different version control systems because storage and usage are often very different. SVN and Helix Core have enough basic similarities to make a full migration possible and the outcome acceptable.

HELIX CORE TO SVN CONVERSION TOOL

Perforce provides a conversion tool for SVN repositories, which can be found here:

<https://swarm.workshop.perforce.com/projects/perforce-software-p4convert/files/main/release>

This tool supports two modes:

- Full import into a new Helix Core Server.
- Incremental import into an existing Helix Core Server.

Please refer to the [documentation](#) for more information.

For assistance with migration requirements, please [contact Perforce Professional Services](#).

Learn More

Perforce has technical documentation, video tutorials, and many other resources available to help teams get more familiar with Helix Core and use this powerful tool more effectively.

SELF-SERVICE

[A home base for all support resources](https://www.perforce.com/support)

<https://www.perforce.com/support>

INTRODUCTION TO PERFORCE

[Get step-by-step instructions for everything](http://www.perforce.com/perforce/doc.current/manuals/intro/index.html)

<http://www.perforce.com/perforce/doc.current/manuals/intro/index.html>

P4 USER'S GUIDE

[Learn to use Helix Core](http://www.perforce.com/perforce/doc.current/manuals/p4guide/index.html)

<http://www.perforce.com/perforce/doc.current/manuals/p4guide/index.html>

PERFORCE SYSTEM ADMINISTRATOR'S GUIDE

[Your role as a system administrator](http://www.perforce.com/perforce/doc.current/manuals/p4sag/index.html)

<http://www.perforce.com/perforce/doc.current/manuals/p4sag/index.html>

PERFORCE DIRECTORY STANDARD

[Establish a formal, documented Perforce Directory Standard](http://info.perforce.com/PDS.html)

<http://info.perforce.com/PDS.html>

About Perforce

Enterprises across the globe rely on Perforce to build and deliver complex digital products faster and with higher quality. Perforce is best known for its highly scalable version management and collaboration platform that securely manages change across all digital content - source code, art files, video files, images, libraries - while supporting the developer and build tools your teams need to be productive, such as Git, Visual Studio, Jenkins, Adobe, Maya and many others. Perforce also offers complete project lifecycle management tools to accelerate a project's delivery cycle by linking the requirements, test plans, source code, and helpdesk in an integrated platform. Perforce is trusted by the world's most innovative brands, including Pixar, NVIDIA, Scania, Ubisoft, and VMware. The company has offices in the US, the United Kingdom, Germany, Canada and Australia, and sales partners around the globe. For more information, please visit www.perforce.com