

Jenkins User Handbook

jenkinsci-docs@googlegroups.com

Table of Contents

Getting Started with Jenkins	1
Installing Jenkins	2
Prerequisites	3
Installation platforms	4
Docker	4
WAR file	7
macOS	8
Linux	8
Windows	9
Other operating systems	9
Post-installation setup wizard	12
Unlocking Jenkins	12
Customizing Jenkins with plugins	12
Creating the first administrator user	13
Using Jenkins	14
Pipeline	15
What is Jenkins Pipeline?	16
Why Pipeline?	18
Pipeline Terms	19
Getting Started with Pipeline	20
Prerequisites	21
Defining a Pipeline	22
Defining a Pipeline in the Web UI	22
Defining a Pipeline in SCM	25
Built-in Documentation	26
Snippet Generator	26
Global Variable Reference	27
Further Reading	28
Additional Resources	28
Using a Jenkinsfile	29
Creating a Jenkinsfile	30
Build	31
Test	32
Deploy	33
Advanced Syntax for Pipeline	35
String Interpolation	35
Working with the Environment	35
Parameters	37

Handling Failures	37
Using multiple agents	38
Optional step arguments	40
Advanced Scripted Pipeline	41
Branches and Pull Requests	43
Creating a Multibranch Pipeline	44
Additional Environment Variables	47
Supporting Pull Requests	47
Using Organization Folders	48
Using Docker with Pipeline	49
Customizing the execution environment	50
Caching data for containers	50
Using multiple containers	51
Using a Dockerfile	52
Specifying a Docker Label	54
Advanced Usage with Scripted Pipeline	55
Running "sidecar" containers	55
Building containers	56
Using a remote Docker server	57
Using a custom registry	58
Extending with Shared Libraries	59
Defining Shared Libraries	60
Directory structure	60
Global Shared Libraries	61
Folder-level Shared Libraries	61
Automatic Shared Libraries	61
Using libraries	62
Loading libraries dynamically	63
Library versions	64
Retrieval Method	64
Writing libraries	67
Accessing steps	67
Defining global variables	68
Defining custom steps	69
Defining a more structured DSL	70
Using third-party libraries	71
Loading resources	71
Pretesting library changes	72
Defining Declarative Pipelines	72
Pipeline Development Tools	74
Blue Ocean Editor	75

Command-line Pipeline Linter	76
Examples	76
"Replay" Pipeline Runs with Modifications	78
Usage	78
Features	79
Limitations	79
Pipeline Unit Testing Framework	80
Pipeline Syntax	81
Declarative Pipeline	82
Sections	82
Directives	88
Parallel	98
Steps	99
Scripted Pipeline	101
Flow Control	101
Steps	102
Differences from plain Groovy	102
Syntax Comparison	103
Blue Ocean	104
What is Blue Ocean?	105
Join the community	106
Frequently Asked Questions	107
Why does Blue Ocean exist?	107
Where is the name from?	107
What does this mean for the classic Jenkins UI?	108
What does this mean for my plugins?	108
What technologies are currently in use?	108
Where can I find the source code?	108
Getting Started with Blue Ocean	109
Installing	110
In an Existing Jenkins Environment	110
With Docker	111
Starting Blue Ocean	112
Navigation bar	113
Switching to the "Classic" UI	114
Creating Pipelines	115
Starting Pipeline Creation	116
Creating a Pipeline for a Git Repository	117
Creating Pipelines for GitHub Repositories	118
Provide a GitHub Access Token	118
Select a GitHub Account or Organization	118

Dashboard	120
Navigation Bar	121
Favorites	122
Pipelines	123
Health Icons	123
Pipeline Run Status	124
Activity View	125
Navigation Bar	126
Activity	127
Branches	128
Pull Requests	129
Pipeline Run Details View	130
Pipeline Run Status	131
Special cases	132
Pipelines outside of Souce Control	132
Freestyle Projects	132
Matrix projects	132
Tabs	133
Pipeline	133
Changes	133
Tests	133
Artifacts	134
Pipeline Editor	136
Starting the editor	137
Limitations	138
Navigation bar	139
Pipeline settings	140
Agent	140
Environment	140
Stage editor	141
Stage configuration	142
Step configuration	143
Save Pipeline dialog	144
Managing Jenkins	145
Configuring the System	146
Managing Security	147
Enabling Security	148
JNLP TCP Port	148
Access Control	149
Markup Formatter	151
Cross Site Request Forgery	152

Caveats	152
Agent/Master Access Control	153
Customizing Access	153
Disabling	155
Managing Tools	157
Built-in tool providers	158
Ant	158
Git	158
JDK	158
Maven	158
Managing Plugins	159
Installing a plugin	160
From the web UI	160
Using the Jenkins CLI	161
Advanced installation	161
Updating a plugin	163
Removing a plugin	164
Uninstalling a plugin	164
Disabling a plugin	165
Pinned plugins	166
Jenkins CLI	167
Using the CLI over SSH	168
Authentication	168
Common Commands	169
Using the CLI client	172
Downloading the client	172
Using the client	172
Client connection modes	172
Script Console	177
Managing Nodes	178
In-process Script Approval	179
Getting Started	180
Groovy Sandbox	181
Script Approval	183
Approve assuming permissions check	184
Managing Users	185
System Administration	186
Backing-up/Restoring Jenkins	187
Monitoring Jenkins	188
Securing Jenkins	189
Access Control	190

Protect users of Jenkins from other threats	191
Disabling Security	192
Managing Jenkins with Chef	193
Managing Jenkins with Puppet	194
Scaling Jenkins	195
Appendix	196
Glossary	197
General Terms	198

Getting Started with Jenkins

This chapter is intended for new users unfamiliar with Jenkins or those without experience with recent versions of Jenkins.

This chapter will lead you through installing an instance of Jenkins on a system for learning purposes and understanding basic Jenkins concepts. It will provide simple step-by-step tutorials on how to do a number common tasks. Each section is intended to be completed in order, with each building on knowledge from the previous section. When you are done you should have enough experience with the core of Jenkins to continue exploring on your own.

If you are already familiar with Jenkins basics and would like to delve deeper into how to use specific features, see [Using Jenkins](#).

If you are a Jenkins administrator and want to know more about managing Jenkins nodes and instances, see [Managing Jenkins](#).

If you are a system administrator and want learn how to back-up, restore, maintain as Jenkins servers and nodes, see [Jenkins System Administration](#).

Installing Jenkins

The procedures on this page are for new installations of Jenkins on a single/local machine.

Jenkins is typically run as a standalone application in its own process with the built-in [Java servlet container/application server \(Jetty\)](#).

Jenkins can also be run as a servlet in different Java servlet containers such as [Apache Tomcat](#) or [GlassFish](#). However, instructions for setting up these types of installations are beyond the scope of this page.

Note: Although this page focusses on local installations of Jenkins, this content can also be used to help set up Jenkins in production environments.

Prerequisites

Minimum hardware requirements:

- 256 MB of RAM
- 1 GB of drive space (although 10 GB is a recommended minimum if running Jenkins as a [Docker](#) container)

Recommended hardware configuration for a small team:

- 1 GB+ of RAM
- 50 GB+ of drive space

Software requirements:

- Java 8 - either a Java Runtime Environment (JRE) or a Java Development Kit (JDK) is fine
Note: This is not a requirement if running Jenkins as a [Docker](#) container.

Installation platforms

This section describes how to install/run Jenkins on different platforms and operating systems.

Docker

[Docker](#) is a platform for running applications in an isolated environment called a "container" (or Docker container). Applications like Jenkins can be downloaded as read-only "images" (or Docker images), each of which is run in Docker as a container. A Docker container is in effect a "running instance" of a Docker image. From this perspective, an image is stored permanently more or less (i.e. insofar as image updates are published), whereas containers are stored temporarily. Read more about these concepts in the Docker documentation's [Getting Started, Part 1: Orientation and setup](#) page.

Docker's fundamental platform and container design means that a single Docker image (for any given application like Jenkins) can be run on any operating system (e.g. macOS, Linux and Windows) which is also running Docker.

Installing Docker

To install Docker on your operating system, visit the [Docker](#) website and navigate to **Get Docker** at the top of their site to begin accessing the Docker download that is suitable for your operating system. Follow the installation instructions on their website.

CAUTION

If you are installing Docker on a Linux-based operating system, ensure you configure Docker so it can be managed as a non-root user. Read more about this in Docker's [Post-installation steps for Linux](#) page of their documentation. This page also contains information about how to configure Docker to start on boot.

Downloading and running Jenkins in Docker

There are several Docker images of Jenkins available.

The recommended Docker image to use is the [jenkinsci/blueocean image](#) (from the [Docker Hub repository](#)). This image contains the current [Long-Term Support \(LTS\) release of Jenkins](#) (which is production-ready) bundled with all Blue Ocean plugins and features. This means that you do not need to install the Blue Ocean plugins separately.

NOTE

A new [jenkinsci/blueocean](#) image is published each time a new release of Blue Ocean is published. You can see a list of previously published versions of the [jenkinsci/blueocean](#) image on the [tags](#) page.

There are also other Jenkins Docker images you can use (accessible through [jenkins/jenkins](#) on Docker Hub). However, these do not come with Blue Ocean, which would need to be installed via the [Manage Jenkins > Manage Plugins](#) page in Jenkins. Read more about this in [Getting started with Blue Ocean](#).

On macOS and Linux

1. Open up a terminal window.
2. Download the `jenkinsci/blueocean` image and run it as a container in Docker using the following `docker run` command:

```
docker run \  
  -u root \  
  --rm \ ①  
  -d \ ②  
  -p 8080:8080 \ ③  
  -v jenkins-data:/var/jenkins_home \ ④  
  -v /var/run/docker.sock:/var/run/docker.sock \  
  jenkinsci/blueocean ⑤
```

- ① (*Optional*) Automatically removes the Docker container (which is the instantiation of the `jenkinsci/blueocean` image below) when it is shut down. This keeps things tidy if you need to quit Jenkins.
- ② (*Optional*) Runs the `jenkinsci/blueocean` container in the background (i.e. "detached" mode) and outputs the container ID. If you do not specify this option, then the running Docker log for this container is output in the terminal window.
- ③ Makes (i.e. "publishes") port 8080 of the `jenkinsci/blueocean` container to port 8080 on the host machine. The 1st number represents the port on the host while the 2nd represents the container's port. Therefore, if you specified `-p 49001:8080` for this option, you would be accessing Jenkins on your host machine through port 49001.
- ④ Maps the `/var/jenkins_home` directory in the container to the Docker `volume` with the name `jenkins-data`. If this volume does not exist, then this `docker run` command will automatically create the volume for you.
Notes: This volume could also be created independently using the `docker volume create` command:
`docker volume create jenkins-data`
Instead of mapping the `/var/jenkins_home` directory to a Docker volume, you could also map this directory to one on your machine's local file system. For example, specifying the option `-v $HOME/jenkins:/var/jenkins_home` would map the container's `/var/jenkins_home` directory to the `jenkins` subdirectory within the `$HOME` directory on your local machine, which would typically be `/Users/<your-username>/jenkins` or `/home/<your-username>/jenkins`.
- ⑤ The `jenkinsci/blueocean` Docker image itself. If this image has not already been downloaded, then this `docker run` command will automatically download the image for you. Furthermore, if any updates to this image were published since you last ran this command, then running this command again will automatically download these published image updates for you.
Note: This Docker image could also be downloaded (or updated) independently using the `docker pull` command:
`docker pull jenkinsci/blueocean`

Note: If copying and pasting the command snippet above does not work, try copying and pasting this annotation-free version here:

```
docker run \  
-u root \  
--rm \  
-d \  
-p 8080:8080 \  
-v jenkins-data:/var/jenkins_home \  
-v /var/run/docker.sock:/var/run/docker.sock \  
jenkinsci/blueocean
```

3. Proceed to the [Post-installation setup wizard](#).

On Windows

1. Open up a command prompt window.
2. Download the `jenkinsci/blueocean` image and run it as a container in Docker using the following `docker run` command:

```
docker run ^  
-u root ^  
--rm ^  
-d ^  
-p 8080:8080 ^  
-v jenkins-data:/var/jenkins_home ^  
-v /var/run/docker.sock:/var/run/docker.sock ^  
jenkinsci/blueocean
```

For an explanation of each of these options, refer to the [macOS and Linux](#) instructions above.

3. Proceed to the [Post-installation setup wizard](#).

Accessing the Jenkins/Blue Ocean Docker container

If you have some experience with Docker and you wish or need to access the `jenkinsci/blueocean` container through a terminal/command prompt using the `docker exec` command, you can add an option like `--name jenkins-blueocean` (with the `docker run` above), which would give the `jenkinsci/blueocean` container the name "jenkins-blueocean".

This means you could access the container (through a separate terminal/command prompt window) with a `docker exec` command like:

```
docker exec -it jenkins-blueocean bash
```

Accessing the Jenkins console log through Docker logs

There is a possibility you may need to access the Jenkins console log, for instance, when [Unlocking Jenkins](#) as part of the [Post-installation setup wizard](#).

If you did not specify the detached mode option `-d` with the `docker run` ... command [above](#), then the

Jenkins console log is easily accessible through the terminal/command prompt window from which you ran this Docker command.

Otherwise, you can access the Jenkins console log through the [Docker logs](#) of the `jenkinsci/blueocean` container using the following command:

```
docker logs <docker-container-name>
```

Your `<docker-container-name>` can be obtained using the `docker ps` command. If you specified the `--name jenkins-blueocean` option in the `docker run ...` command above (see also [Accessing the Jenkins/Blue Ocean Docker container](#)), you can simply use the `docker logs` command:

```
docker logs jenkins-blueocean
```

Accessing the Jenkins home directory

There is a possibility you may need to access the Jenkins home directory, for instance, to check the details of a Jenkins build in the `workspace` subdirectory.

If you mapped the Jenkins home directory (`/var/jenkins_home`) to one on your machine's local file system (i.e. in the `docker run ...` command [above](#)), then you can access the contents of this directory through your machine's usual terminal/command prompt.

Otherwise, if you specified the `-v jenkins-data:/var/jenkins_home` option in the `docker run ...` command, you can access the contents of the Jenkins home directory through the `jenkinsci/blueocean` container's terminal/command prompt using the `docker exec` command:

```
docker exec -it <docker-container-name> bash
```

As mentioned [above](#), your `<docker-container-name>` can be obtained using the `docker ps` command. If you specified the

`--name jenkins-blueocean` option in the `docker run ...` command above (see also [Accessing the Jenkins/Blue Ocean Docker container](#)), you can simply use the `docker exec` command:

```
docker exec -it jenkins-blueocean bash
```

WAR file

The Web application ARchive (WAR) file version of Jenkins can be installed on any operating system or platform that supports Java.

To download and run the WAR file version of Jenkins:

1. Download the [latest stable Jenkins WAR file](#) to an appropriate directory on your machine.
2. Open up a terminal/command prompt window to the download directory.
3. Run the command `java -jar jenkins.war`.
4. Browse to `localhost:8080` and wait until the **Unlock Jenkins** page appears.
5. Continue on with the [Post-installation setup wizard](#) below.

Notes:

- Unlike downloading and running Jenkins with Blue Ocean in Docker ([above](#)), this process does not automatically install the Blue Ocean features, which would need to be installed separately via the [Manage Jenkins > Manage Plugins](#) page in Jenkins. Read more about the specifics for installing Blue Ocean on the [Getting Started with Blue Ocean](#) page.
- You can change the port by specifying the `--httpPort` option when you run the `java -jar jenkins.war` command. For example, to make Jenkins accessible through port 49001, then run Jenkins using the command:
`java -jar jenkins.war --httpPort=49001`

macOS

To install from the website, using a package:

- [Download the latest package](#)
- Open the package and follow the instructions

Jenkins can also be installed using `brew`:

- Install the latest release version

```
brew install jenkins
```

- Install the LTS version

```
brew install jenkins-lts
```

Linux

Debian/Ubuntu

On Debian-based distributions, such as Ubuntu, you can install Jenkins through `apt`.

Recent versions are available in [an apt repository](#). Older but stable LTS versions are in [this apt repository](#).

```
wget -q -O - https://pkg.jenkins.io/debian/jenkins.io.key | sudo apt-key add -
sudo sh -c 'echo deb http://pkg.jenkins.io/debian-stable binary/ >
/etc/apt/sources.list.d/jenkins.list'
sudo apt-get update
sudo apt-get install jenkins
```

This package installation will:

- Setup Jenkins as a daemon launched on start. See `/etc/init.d/jenkins` for more details.

- Create a 'jenkins' user to run this service.
- Direct console log output to the file `/var/log/jenkins/jenkins.log`. Check this file if you are troubleshooting Jenkins.
- Populate `/etc/default/jenkins` with configuration parameters for the launch, e.g `JENKINS_HOME`
- Set Jenkins to listen on port 8080. Access this port with your browser to start configuration.

NOTE

If your `/etc/init.d/jenkins` file fails to start Jenkins, edit the `/etc/default/jenkins` to replace the line `----HTTP_PORT=8080----` with `----HTTP_PORT=8081----` Here, "8081" was chosen but you can put another port available.

Windows

To install from the website, using the installer:

- [Download the latest package](#)
- Open the package and follow the instructions

Other operating systems

OpenIndiana Hipster

On a system running [OpenIndiana Hipster](#) Jenkins can be installed in either the local or global zone using the [Image Packaging System](#) (IPS).

IMPORTANT

Disclaimer: This platform is NOT officially supported by the Jenkins team, use it at your own risk. Packaging and integration described in this section is maintained by the OpenIndiana Hipster team, bundling the generic `jenkins.war` to work in that operating environment.

For the common case of running the newest packaged weekly build as a standalone (Jetty) server, simply execute:

```
pkg install jenkins
svcadm enable jenkins
```

The common packaging integration for a standalone service will:

- Create a `jenkins` user to run the service and to own the directory structures under `/var/lib/jenkins`.
- Pull the OpenJDK8 and other packages required to execute Jenkins, including the `jenkins-core-weekly` package with the latest `jenkins.war`.

CAUTION

Long-Term Support (LTS) Jenkins releases currently do not support OpenZFS-based systems, so no packaging is provided at this time.

- Set up Jenkins as an SMF service instance (`svc:/network/http:jenkins`) which can then be enabled with the `svcadm` command demonstrated above.
- Set up Jenkins to listen on port 8080.
- Configure the log output to be managed by SMF at `/var/svc/log/network-http:jenkins.log`.

Once Jenkins is running, consult the log (`/var/svc/log/network-http:jenkins.log`) to retrieve the generated administrator password for the initial set up of Jenkins, usually it will be found at `/var/lib/jenkins/home/secrets/initialAdminPassword`. Then navigate to `localhost:8080` to [complete configuration of the Jenkins instance](#).

To change attributes of the service, such as environment variables like `JENKINS_HOME` or the port number used for the Jetty web server, use the `svccfg` utility:

```
svccfg -s svc:/network/http:jenkins editprop
svcadm refresh svc:/network/http:jenkins
```

You can also refer to `/lib/svc/manifest/network/jenkins-standalone.xml` for more details and comments about currently supported tunables of the SMF service. Note that the `jenkins` user account created by the packaging is specially privileged to allow binding to port numbers under 1024.

The current status of Jenkins-related packages available for the given release of OpenIndiana can be queried with:

```
pkg info -r '*jenkins*'
```

Upgrades to the package can be performed by updating the entire operating environment with `pkg update`, or specifically for Jenkins core software with:

```
pkg update jenkins-core-weekly
```

CAUTION

Procedure for updating the package will restart the currently running Jenkins process. Make sure to prepare it for shutdown and finish all running jobs before updating, if needed.

Solaris, OmniOS, SmartOS, and other siblings

Generally it should suffice to install Java 8 and [download](#) the `jenkins.war` and run it as a standalone process or under an application server such as [Apache Tomcat](#).

Some caveats apply:

- Headless JVM and fonts: For OpenJDK builds on minimized-footprint systems, there may be [issues running the headless JVM](#), because Jenkins needs some fonts to render certain pages.
- ZFS-related JVM crashes: When Jenkins runs on a system detected as a `SunOS`, it tries to load

integration for advanced ZFS features using the bundled `libzfs.jar` which maps calls from Java to native `libzfs.so` routines provided by the host OS. Unfortunately, that library was made for binary utilities built and bundled by the OS along with it at the same time, and was never intended as a stable interface exposed to consumers. As the forks of Solaris legacy, including ZFS and later the OpenZFS initiative evolved, many different binary function signatures were provided by different host operating systems - and when Jenkins `libzfs.jar` invoked the wrong signature, the whole JVM process crashed. A solution was proposed and integrated in `jenkins.war` since weekly release 2.55 (and not yet in any LTS to date) which enables the administrator to configure which function signatures should be used for each function known to have different variants, apply it to their application server initialization options and then run and update the generic `jenkins.war` without further workarounds. See [the libzfs4j Git repository](#) for more details, including a script to try and "lock-pick" the configuration needed for your particular distribution (in particular if your kernel updates bring a new incompatible `libzfs.so`).

Also note that forks of the OpenZFS initiative may provide ZFS on various BSD, Linux, and macOS distributions. Once Jenkins supports detecting ZFS capabilities, rather than relying on the `SunOS` check, the above caveats for ZFS integration with Jenkins should be considered.

Post-installation setup wizard

After downloading, installing and running Jenkins using one of the procedures above, the post-installation setup wizard begins.

This setup wizard takes you through a few quick "one-off" steps to unlock Jenkins, customize it with plugins and create the first administrator user through which you can continue accessing Jenkins.


Unlocking Jenkins

When you first access a new Jenkins instance, you are asked to unlock it using an automatically-generated password.

1. Browse to `localhost:8080` (or whichever port you configured for Jenkins when installing it) and wait until the **Unlock Jenkins** page appears.



2. From the Jenkins console log output, copy the automatically-generated alphanumeric password (between the 2 sets of asterisks).



3. On the **Unlock Jenkins** page, paste this password into the **Administrator password** field and click **Continue**.

Notes:

- If you ran Jenkins in Docker in detached mode, you can access the Jenkins console log from the Docker logs ([above](#)).
- The Jenkins console log indicates the location (in the Jenkins home directory) where this password can also be obtained. This password must be entered in the setup wizard on new Jenkins installations before you can access Jenkins's main UI. This password also serves as the default administrator account's password (with username "admin") if you happen to skip the subsequent user-creation step in the setup wizard.

Customizing Jenkins with plugins

After [unlocking Jenkins](#), the **Customize Jenkins** page appears. Here you can install any number of useful plugins as part of your initial setup.

Click one of the two options shown:

- **Install suggested plugins** - to install the recommended set of plugins, which are based on most common use cases.

- **Select plugins to install** - to choose which set of plugins to initially install. When you first access the plugin selection page, the suggested plugins are selected by default.

NOTE If you are not sure what plugins you need, choose **Install suggested plugins**. You can install (or remove) additional Jenkins plugins at a later point in time via the [Manage Jenkins > Manage Plugins](#) page in Jenkins.

The setup wizard shows the progression of Jenkins being configured and your chosen set of Jenkins plugins being installed. This process may take a few minutes.

Creating the first administrator user

Finally, after [customizing Jenkins with plugins](#), Jenkins asks you to create your first administrator user.

1. When the **Create First Admin User** page appears, specify the details for your administrator user in the respective fields and click **Save and Finish**.
2. When the **Jenkins is ready** page appears, click **Start using Jenkins**.

Notes:

- This page may indicate **Jenkins is almost ready!** instead and if so, click **Restart**.
 - If the page does not automatically refresh after a minute, use your web browser to refresh the page manually.
3. If required, log in to Jenkins with the credentials of the user you just created and you are ready to start using Jenkins!

IMPORTANT From this point on, the Jenkins UI is only accessible by providing valid username and password credentials.

Using Jenkins

This chapter will describe how to work with Jenkins as a non-administrator user. It will cover topics applicable to anyone using Jenkins on a day-to-day basis. This includes basic topics such as selecting, running, and monitoring existing jobs, and how to find and review jobs results. It will continue on to discussing a number of topics around designing and creating projects.

This chapter is intended to be used by Jenkins users of all skill levels. The sections are structured in a feature-centric way for easier searching and reference by experienced users. At the same time, to help beginners, we've attempted to order sections in the chapter from simpler to progressively more complex feature areas. Also, topics within each section will progress from basic to advanced, with expert-level considerations and corner-cases at the end or in a separate section later in the chapter.

If you are not yet familiar with basic Jenkins terminology and features, start with [Getting Started with Jenkins](#).

If you are a Jenkins administrator and want to know more about managing Jenkins nodes and instances, see [Managing Jenkins](#).

If you are a system administrator and want learn how to back-up, restore, maintain as Jenkins servers and nodes, see [Jenkins System Administration](#).

WARNING

To Contributors: This chapter functions as a continuation of "[Getting Started with Jenkins](#)", but the format will be slightly different - see the description above. We need to balance between providing a feature reference for experienced users with providing a continuing on-ramp for beginners. Sections should be written and ordered to only assume knowledge from "Getting Started" or from previous sections in this chapter.

Pipeline

This chapter will cover all aspects of Jenkins Pipeline, from running Pipelines to writing Pipeline code, and even extending Pipeline itself.

This chapter is intended to be used by Jenkins users of all skill levels, but beginners may need to refer to some sections of "[Using Jenkins](#)" to understand some topics covered in this chapter.

If you are not yet familiar with basic Jenkins terminology and features, start with [Getting Started with Jenkins](#).

What is Jenkins Pipeline?

Jenkins Pipeline (or simply "Pipeline" with a capital "P") is a suite of plugins which supports implementing and integrating *continuous delivery pipelines* into Jenkins.

A *continuous delivery pipeline* is an automated expression of your process for getting software from version control right through to your users and customers. Every change to your software (committed in source control) goes through a complex process on its way to being released. This process involves building the software in a reliable and repeatable manner, as well as the progression of the built software (called a "build") through multiple stages of testing and deployment.

Pipeline provides an extensible set of tools for modeling simple-to-complex delivery pipelines "as code" via the [Pipeline Domain Specific Language \(DSL\) syntax](#). [1: [Domain-Specific Language](#)]

Typically, the definition of a Jenkins Pipeline is written into a text file (called a **Jenkinsfile**) which in turn is checked into a project's source control repository. [2: [Source Control Management](#)] This is the foundation of "Pipeline-as-Code"; treating the continuous delivery pipeline a part of the application to be versioned and reviewed like any other code. Creating a **Jenkinsfile** provides a number of immediate benefits:

- Automatically create Pipelines for all Branches and Pull Requests
- Code review/iteration on the Pipeline
- Audit trail for the Pipeline
- Single source of truth [3: en.wikipedia.org/wiki/Single_Source_of_Truth] for the Pipeline, which can be viewed and edited by multiple members of the project.

While the syntax for defining a Pipeline, either in the web UI or with a **Jenkinsfile**, is the same, it's generally considered best practice to define the Pipeline in a **Jenkinsfile** and check that in to source control.

Here's an example of a **Jenkinsfile**:

```

// Declarative //
pipeline {
  agent any ①

  stages {
    stage('Build') { ②
      steps { ③
        sh 'make' ④
      }
    }
    stage('Test'){
      steps {
        sh 'make check'
        junit 'reports/**/*.xml' ⑤
      }
    }
    stage('Deploy') {
      steps {
        sh 'make publish'
      }
    }
  }
}
// Script //
node {
  stage('Build') {
    sh 'make'
  }

  stage('Test') {
    sh 'make check'
    junit 'reports/**/*.xml'
  }

  stage('Deploy') {
    sh 'make publish'
  }
}

```

- ① **agent** indicates that Jenkins should allocate an executor and workspace for this part of the Pipeline.
- ② **stage** describes a stage of this Pipeline.
- ③ **steps** describes the steps to be run in this **stage**
- ④ **sh** executes the given shell command
- ⑤ **junit** is a Pipeline **step** provided by the plugin:junit[JUnit plugin] for aggregating test reports.

Why Pipeline?

Jenkins is, fundamentally, an automation engine which supports a number of automation patterns. Pipeline adds a powerful set of automation tools onto Jenkins, supporting use cases that span from simple continuous integration to comprehensive continuous delivery pipelines. By modeling a series of related tasks, users can take advantage of the many features of Pipeline:

- **Code:** Pipelines are implemented in code and typically checked into source control, giving teams the ability to edit, review, and iterate upon their delivery pipeline.
- **Durable:** Pipelines can survive both planned and unplanned restarts of the Jenkins master.
- **Pausable:** Pipelines can optionally stop and wait for human input or approval before continuing the Pipeline run.
- **Versatile:** Pipelines support complex real-world continuous delivery requirements, including the ability to fork/join, loop, and perform work in parallel.
- **Extensible:** The Pipeline plugin supports custom extensions to its DSL [1: [Domain-Specific Language](#)] and multiple options for integration with other plugins.

While Jenkins has always allowed rudimentary forms of chaining Freestyle Jobs together to perform sequential tasks, [4: Additional plugins have been used to implement complex behaviors utilizing Freestyle Jobs such as the Copy Artifact, Parameterized Trigger, and Promoted Builds plugins] Pipeline makes this concept a first-class citizen in Jenkins.

Building on the core Jenkins value of extensibility, Pipeline is also extensible both by users with [Pipeline Shared Libraries](#) and by plugin developers. [5: `plugin:github-organization-folder`[GitHub Organization Folder plugin]]

The flowchart below is an example of one continuous delivery scenario easily modeled in Jenkins Pipeline:

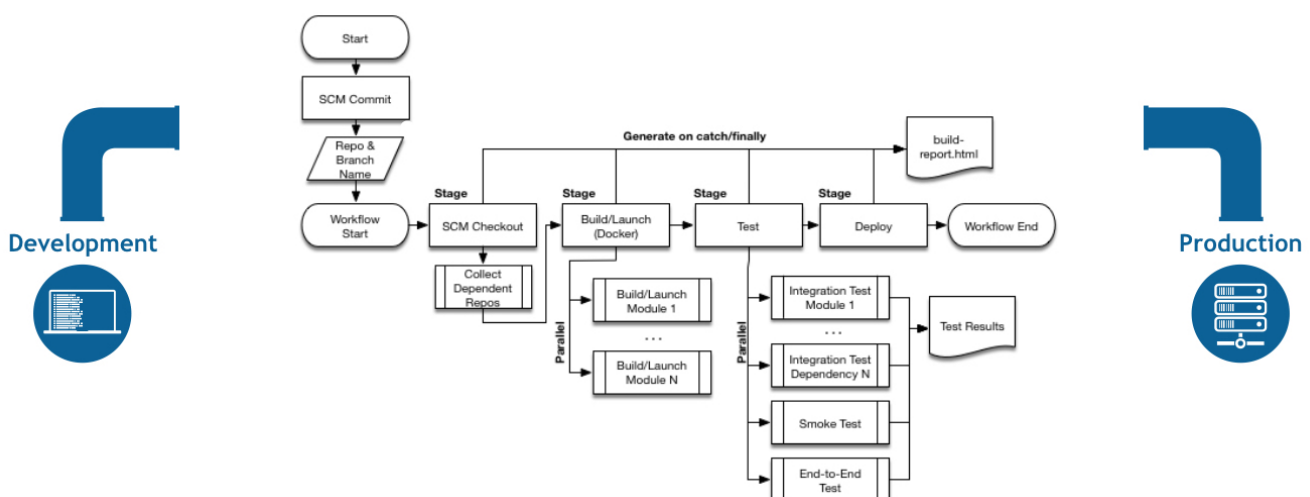


Figure 1. Pipeline Flow

Pipeline Terms

Step

A single task; fundamentally steps tell Jenkins *what* to do. For example, to execute the shell command `make` use the `sh` step: `sh 'make'`. When a plugin extends the Pipeline DSL, that typically means the plugin has implemented a new *step*.

Node

Most *work* a Pipeline performs is done in the context of one or more declared `node` steps. Confining the work inside of a node step does two things:

1. Schedules the steps contained within the block to run by adding an item to the Jenkins queue. As soon as an executor is free on a node, the steps will run.
2. Creates a workspace (a directory specific to that particular Pipeline) where work can be done on files checked out from source control.

CAUTION

Depending on your Jenkins configuration, some workspaces may not get automatically cleaned up after a period of inactivity. See tickets and discussion linked from [JENKINS-2111](#) for more information.

Stage

`stage` is a step for defining a conceptually distinct subset of the entire Pipeline, for example: "Build", "Test", and "Deploy", which is used by many plugins to visualize or present Jenkins Pipeline status/progress. [6: [Blue Ocean, Pipeline Stage View plugin](#)]

Getting Started with Pipeline

Jenkins Pipeline is a suite of plugins which supports implementing and integrating continuous delivery pipelines into Jenkins. Pipeline provides an extensible set of tools for modeling simple-to-complex delivery pipelines "as code" via the Pipeline DSL. [7: [Domain-Specific Language](#)]

This section introduces some of the key concepts to Jenkins Pipeline and help introduce the basics of defining and working with Pipelines inside of a running Jenkins instance.

Prerequisites

To use Jenkins Pipeline, you will need:

- Jenkins 2.x or later (older versions back to 1.642.3 may work but are not recommended)
- Pipeline plugin [8: [Pipeline plugin](#)]

To learn how to install and manage plugins, consult [Managing Plugins](#).

Defining a Pipeline

Scripted Pipeline is written in [Groovy](#). The relevant bits of [Groovy syntax](#) will be introduced as necessary in this document, so while an understanding of Groovy is helpful, it is not required to work with Pipeline.

A basic Pipeline can be created in either of the following ways:

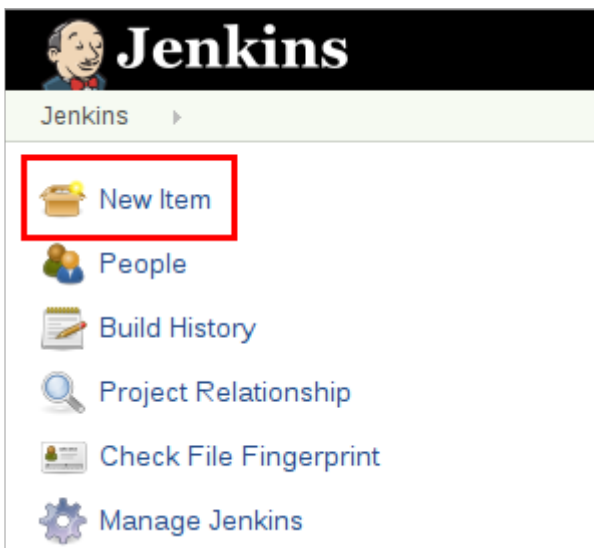
- By entering a script directly in the Jenkins web UI.
- By creating a [Jenkinsfile](#) which can be checked into a project's source control repository.

The syntax for defining a Pipeline with either approach is the same, but while Jenkins supports entering Pipeline directly into the web UI, it's generally considered best practice to define the Pipeline in a [Jenkinsfile](#) which Jenkins will then load directly from source control. [9: en.wikipedia.org/wiki/Source_control_management]

Defining a Pipeline in the Web UI

To create a basic Pipeline in the Jenkins web UI, follow these steps:

- Click **New Item** on Jenkins home page.



- Enter a name for your Pipeline, select **Pipeline** and click **OK**.

CAUTION

Jenkins uses the name of the Pipeline to create directories on disk. Pipeline names which include spaces may uncover bugs in scripts which do not expect paths to contain spaces.

Enter an item name

an-example

» Required field



Freestyle project

This is the central feature of Jenkins. Jenkins will build your project, combining any SCM with any build system, and this can be even used for something other than software build.



Pipeline

Orchestrates long-running activities that can span multiple build slaves. Suitable for building pipelines (formerly known as workflows) and/or organizing complex activities that do not easily fit in free-style job type.



External Job

This type of job allows you to record the execution of a process run outside Jenkins, even on a remote machine. This is designed so that you can use Jenkins as a dashboard of your existing automation system.



Multi-configuration project

Suitable for projects that need a large number of different configurations, such as testing on multiple environments, platform-specific builds, etc.



Folder

Creates a container that stores nested items in it. Useful for grouping things together. Unlike view, which is just a filter, a folder creates a separate namespace, so you can have multiple things of the same name as long as they are in different folders.



GitHub Organization

Scans a GitHub organization (or user account) for all repositories matching some defined markers.



Multibranch Pipeline

Creates a set of Pipeline projects according to detected branches in one SCM repository.

OK

- In the **Script** text area, enter a Pipeline and click **Save**.

Pipeline

Definition Pipeline script

Script


```
1 node {
2   echo "Hello World"
3 }
```

Use Groovy Sandbox

[Pipeline Syntax](#)










Save Apply

- Click **Build Now** to run the Pipeline.



Jenkins

Jenkins > an-example >

-  [Back to Dashboard](#)
-  [Status](#)
-  [Changes](#)
-  **Build Now**
-  [Delete Pipeline](#)
-  [Configure](#)
-  [Move](#)
-  [Full Stage View](#)
-  [Pipeline Syntax](#)

- Click **#1** under "Build History" and then click **Console Output** to see the full output from the Pipeline.

Console Output

```
Started by user admin
[Pipeline] node
Running on master in /var/jenkins_home/workspace/an-example
[Pipeline] {
[Pipeline] echo
Hello World
[Pipeline] }
[Pipeline] // node
[Pipeline] End of Pipeline
Finished: SUCCESS
```

The example above shows a successful run of a basic Pipeline created in the Jenkins web UI, using two steps.

```
// Script //
node { ①
    echo 'Hello World' ②
}
// Declarative not yet implemented //
```

① `node` allocates an executor and workspace in the Jenkins environment.

② `echo` writes simple string in the Console Output.

Defining a Pipeline in SCM

Complex Pipelines are hard to write and maintain within the text area of the Pipeline configuration page. To make this easier, Pipeline can also be written in a text editor and checked into source control as a `Jenkinsfile` which Jenkins can load via the **Pipeline Script from SCM** option.

To do this, select **Pipeline script from SCM** when defining the Pipeline.

With the **Pipeline script from SCM** option selected, you do not enter any Groovy code in the Jenkins UI; you just indicate by specifying a path where in source code you want to retrieve the pipeline from. When you update the designated repository, a new build is triggered, as long as the Pipeline is configured with an SCM polling trigger.

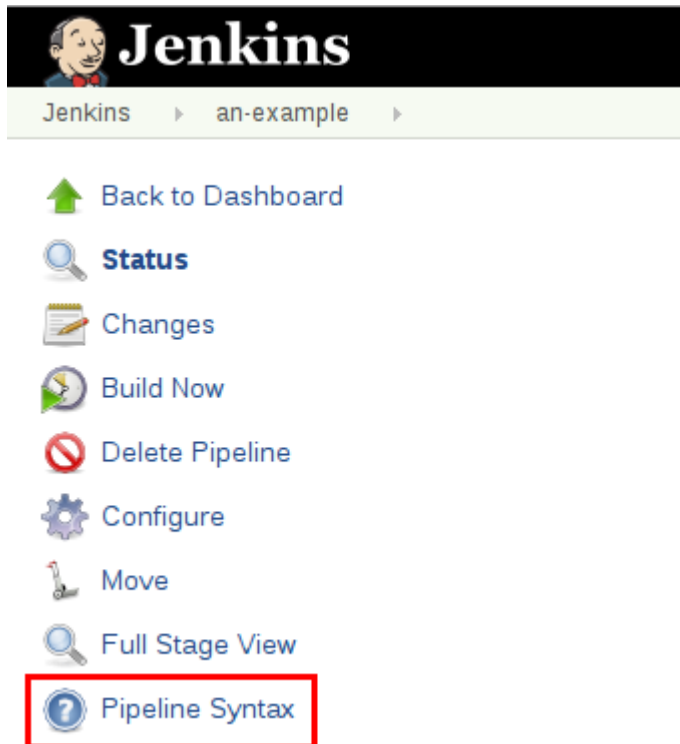
TIP

The first line of a `Jenkinsfile` should be `#!/usr/bin/env groovy` [12: en.wikipedia.org/wiki/Hashbang] [13: groovy-lang.org/syntax.html#shebang_line] which text editors, IDEs, GitHub, etc will use to syntax highlight the `Jenkinsfile` properly as Groovy code.

Built-in Documentation

Pipeline ships with built-in documentation features to make it easier to create Pipelines of varying complexities. This built-in documentation is automatically generated and updated based on the plugins installed in the Jenkins instance.

The built-in documentation can be found globally at: localhost:8080/pipeline-syntax/, assuming you have a Jenkins instance running on localhost port 8080. The same documentation is also linked as **Pipeline Syntax** in the side-bar for any configured Pipeline project.



Snippet Generator

The built-in "Snippet Generator" utility is helpful for creating bits of code for individual steps, discovering new steps provided by plugins, or experimenting with different parameters for a particular step.

The Snippet Generator is dynamically populated with a list of the steps available to the Jenkins instance. The number of steps available is dependent on the plugins installed which explicitly expose steps for use in Pipeline.

To generate a step snippet with the Snippet Generator:

1. Navigate to the **Pipeline Syntax** link (referenced above) from a configured Pipeline, or at localhost:8080/pipeline-syntax/.
2. Select the desired step in the **Sample Step** dropdown menu
3. Use the dynamically populated area below the **Sample Step** dropdown to configure the selected step.
4. Click **Generate Pipeline Script** to create a snippet of Pipeline which can be copied and pasted

into a Pipeline.

Steps

Sample Step stage: Stage

Stage Name 

[Generate Pipeline Script](#)

```
stage('Deploy') {  
  // some block  
}
```

To access additional information and/or documentation about the step selected, click on the help icon (indicated by the red arrow in the image above).

Global Variable Reference

In addition to the Snippet Generator, which only surfaces steps, Pipeline also provides a built-in "**Global Variable Reference**." Like the Snippet Generator, it is also dynamically populated by plugins. Unlike the Snippet Generator however, the Global Variable Reference only contains documentation for **variables** provided by Pipeline or plugins, which are available for Pipelines.

The variables provided by default in Pipeline are:

env

Environment variables accessible from Scripted Pipeline, for example: `env.PATH` or `env.BUILD_ID`. Consult the built-in [Global Variable Reference](#) for a complete, and up to date, list of environment variables available in Pipeline.

params

Exposes all parameters defined for the Pipeline as a read-only [Map](#), for example: `params.MY_PARAM_NAME`.

currentBuild

May be used to discover information about the currently executing Pipeline, with properties such as `currentBuild.result`, `currentBuild.displayName`, etc. Consult the built-in [Global Variable Reference](#) for a complete, and up to date, list of properties available on `currentBuild`.

Further Reading

This section merely scratches the surface of what can be done with Jenkins Pipeline, but should provide enough of a foundation for you to start experimenting with a test Jenkins instance.

In the next section, [The Jenkinsfile](#), more Pipeline steps will be discussed along with patterns for implementing successful, real-world, Jenkins Pipelines.

Additional Resources

- [Pipeline Steps Reference](#), encompassing all steps provided by plugins distributed in the Jenkins Update Center.
 - [Pipeline Examples](#), a community-curated collection of copyable Pipeline examples.
-

Using a Jenkinsfile

This section builds on the information covered in [Getting Started](#), and introduces more useful steps, common patterns, and demonstrates some non-trivial `Jenkinsfile` examples.

Creating a `Jenkinsfile`, which is checked into source control [14: en.wikipedia.org/wiki/Source_control_management], provides a number of immediate benefits:

- Code review/iteration on the Pipeline
- Audit trail for the Pipeline
- Single source of truth [15: en.wikipedia.org/wiki/Single_Source_of_Truth] for the Pipeline, which can be viewed and edited by multiple members of the project.

Pipeline supports [two syntaxes](#), Declarative (introduced in Pipeline 2.5) and Scripted Pipeline. Both of which support building continuous delivery pipelines. Both may be used to define a Pipeline in either the web UI or with a `Jenkinsfile`, though it's generally considered a best practice to create a `Jenkinsfile` and check the file into the source control repository.

Creating a Jenkinsfile

As discussed in the [Getting Started](#) section, a **Jenkinsfile** is a text file that contains the definition of a Jenkins Pipeline and is checked into source control. Consider the following Pipeline which implements a basic three-stage continuous delivery pipeline.

```
// Declarative //
pipeline {
  agent any

  stages {
    stage('Build') {
      steps {
        echo 'Building..'
      }
    }
    stage('Test') {
      steps {
        echo 'Testing..'
      }
    }
    stage('Deploy') {
      steps {
        echo 'Deploying....'
      }
    }
  }
}

// Script //
node {
  stage('Build') {
    echo 'Building....'
  }
  stage('Test') {
    echo 'Building....'
  }
  stage('Deploy') {
    echo 'Deploying....'
  }
}
```

Not all Pipelines will have these same three stages, but it is a good starting point to define them for most projects. The sections below will demonstrate the creation and execution of a simple Pipeline in a test installation of Jenkins.

NOTE

It is assumed that there is already a source control repository set up for the project and a Pipeline has been defined in Jenkins following [these instructions](#).

Using a text editor, ideally one which supports [Groovy](#) syntax highlighting, create a new [Jenkinsfile](#) in the root directory of the project.

The Declarative Pipeline example above contains the minimum necessary structure to implement a continuous delivery pipeline. The [agent directive](#), which is required, instructs Jenkins to allocate an executor and workspace for the Pipeline. Without an [agent](#) directive, not only is the Declarative Pipeline not valid, it would not be capable of doing any work! By default the [agent](#) directive ensures that the source repository is checked out and made available for steps in the subsequent stages`

The [stages directive](#), and [steps directives](#) are also required for a valid Declarative Pipeline as they instruct Jenkins what to execute and in which stage it should be executed.

For more advanced usage with Scripted Pipeline, the example above [node](#) is a crucial first step as it allocates an executor and workspace for the Pipeline. In essence, without [node](#), a Pipeline cannot do any work! From within [node](#), the first order of business will be to checkout the source code for this project. Since the [Jenkinsfile](#) is being pulled directly from source control, Pipeline provides a quick and easy way to access the right revision of the source code

```
// Script //
node {
    checkout scm ①
    /* .. snip .. */
}
// Declarative not yet implemented //
```

① The [checkout](#) step will checkout code from source control; [scm](#) is a special variable which instructs the [checkout](#) step to clone the specific revision which triggered this Pipeline run.

Build

For many projects the beginning of "work" in the Pipeline would be the "build" stage. Typically this stage of the Pipeline will be where source code is assembled, compiled, or packaged. The [Jenkinsfile](#) is **not** a replacement for an existing build tool such as GNU/Make, Maven, Gradle, etc, but rather can be viewed as a glue layer to bind the multiple phases of a project's development lifecycle (build, test, deploy, etc) together.

Jenkins has a number of plugins for invoking practically any build tool in general use, but this example will simply invoke [make](#) from a shell step ([sh](#)). The [sh](#) step assumes the system is Unix/Linux-based, for Windows-based systems the [bat](#) could be used instead.

```

// Declarative //
pipeline {
  agent any

  stages {
    stage('Build') {
      steps {
        sh 'make' ①
        archiveArtifacts artifacts: '**/target/*.jar', fingerprint: true ②
      }
    }
  }
}
// Script //
node {
  stage('Build') {
    sh 'make' ①
    archiveArtifacts artifacts: '**/target/*.jar', fingerprint: true ②
  }
}

```

① The `sh` step invokes the `make` command and will only continue if a zero exit code is returned by the command. Any non-zero exit code will fail the Pipeline.

② `archiveArtifacts` captures the files built matching the include pattern (`*/target/*.jar`) and saves them to the Jenkins master for later retrieval.

TIP

Archiving artifacts is not a substitute for using external artifact repositories such as Artifactory or Nexus and should be considered only for basic reporting and file archival.

Test

Running automated tests is a crucial component of any successful continuous delivery process. As such, Jenkins has a number of test recording, reporting, and visualization facilities provided by a [number of plugins](#). At a fundamental level, when there are test failures, it is useful to have Jenkins record the failures for reporting and visualization in the web UI. The example below uses the `junit` step, provided by the `plugin:junit[JUnit plugin]`.

In the example below, if tests fail, the Pipeline is marked "unstable", as denoted by a yellow ball in the web UI. Based on the recorded test reports, Jenkins can also provide historical trend analysis and visualization.

```

// Declarative //
pipeline {
  agent any

  stages {
    stage('Test') {
      steps {
        /* `make check` returns non-zero on test failures,
        * using `true` to allow the Pipeline to continue nonetheless
        */
        sh 'make check || true' ①
        junit '**/target/*.xml' ②
      }
    }
  }
}

// Script //
node {
  /* .. snip .. */
  stage('Test') {
    /* `make check` returns non-zero on test failures,
    * using `true` to allow the Pipeline to continue nonetheless
    */
    sh 'make check || true' ①
    junit '**/target/*.xml' ②
  }
  /* .. snip .. */
}

```

① Using an inline shell conditional (`sh 'make || true'`) ensures that the `sh` step always sees a zero exit code, giving the `junit` step the opportunity to capture and process the test reports. Alternative approaches to this are covered in more detail in the [\[handling-failures\]](#) section below.

② `junit` captures and associates the JUnit XML files matching the inclusion pattern (`*/target/*.xml`).

Deploy

Deployment can imply a variety of steps, depending on the project or organization requirements, and may be anything from publishing built artifacts to an Artifactory server, to pushing code to a production system.

At this stage of the example Pipeline, both the "Build" and "Test" stages have successfully executed. In essence, the "Deploy" stage will only execute assuming previous stages completed successfully, otherwise the Pipeline would have exited early.


```

// Declarative //
pipeline {
  agent any

  stages {
    stage('Deploy') {
      when {
        expression {
          currentBuild.result == null || currentBuild.result == 'SUCCESS' ①
        }
      }
      steps {
        sh 'make publish'
      }
    }
  }
}
// Script //
node {
  /* .. snip .. */
  stage('Deploy') {
    if (currentBuild.result == null || currentBuild.result == 'SUCCESS') { ①
      sh 'make publish'
    }
  }
  /* .. snip .. */
}

```

① Accessing the `currentBuild.result` variable allows the Pipeline to determine if there were any test failures. In which case, the value would be `UNSTABLE`.

Assuming everything has executed successfully in the example Jenkins Pipeline, each successful Pipeline run will have associated build artifacts archived, test results reported upon and the full console output all in Jenkins.

A Scripted Pipeline can include conditional tests (shown above), loops, try/catch/finally blocks and even functions. The next section will cover this advanced Scripted Pipeline syntax in more detail.

Advanced Syntax for Pipeline

String Interpolation

Jenkins Pipeline uses rules identical to [Groovy](#) for string interpolation. Groovy's String interpolation support can be confusing to many newcomers to the language. While Groovy supports declaring a string with either single quotes, or double quotes, for example:

```
def singlyQuoted = 'Hello'  
def doublyQuoted = "World"
```

Only the latter string will support the dollar-sign (\$) based string interpolation, for example:

```
def username = 'Jenkins'  
echo 'Hello Mr. ${username}'  
echo "I said, Hello Mr. ${username}"
```

Would result in:

```
Hello Mr. ${username}  
I said, Hello Mr. Jenkins
```

Understanding how to use string interpolation is vital for using some of Pipeline's more advanced features.

Working with the Environment

Jenkins Pipeline exposes environment variables via the global variable `env`, which is available from anywhere within a [Jenkinsfile](#). The full list of environment variables accessible from within Jenkins Pipeline is documented at <localhost:8080/pipeline-syntax/globals#env>, assuming a Jenkins master is running on `localhost:8080`, and includes:

BUILD_ID

The current build ID, identical to `BUILD_NUMBER` for builds created in Jenkins versions 1.597+

JOB_NAME

Name of the project of this build, such as "foo" or "foo/bar".

JENKINS_URL

Full URL of Jenkins, such as <example.com:port/jenkins/> (NOTE: only available if Jenkins URL set in "System Configuration")

Referencing or using these environment variables can be accomplished like accessing any key in a Groovy [Map](#), for example:

```
// Declarative //
pipeline {
  agent any
  stages {
    stage('Example') {
      steps {
        echo "Running ${env.BUILD_ID} on ${env.JENKINS_URL}"
      }
    }
  }
}
// Script //
node {
  echo "Running ${env.BUILD_ID} on ${env.JENKINS_URL}"
}
```

Setting environment variables

Setting an environment variable within a Jenkins Pipeline is accomplished differently depending on whether Declarative or Scripted Pipeline is used.

Declarative Pipeline supports an [environment](#) directive, whereas users of Scripted Pipeline must use the `withEnv` step.

```
// Declarative //
pipeline {
  agent any
  environment { ①
    CC = 'clang'
  }
  stages {
    stage('Example') {
      environment { ②
        DEBUG_FLAGS = '-g'
      }
      steps {
        sh 'printenv'
      }
    }
  }
}
// Script //
node {
  /* .. snip .. */
  withEnv(["PATH+MAVEN=${tool 'M3'}/bin"]) {
    sh 'mvn -B verify'
  }
}
```

- ① An `environment` directive used in the top-level `pipeline` block will apply to all steps within the Pipeline.
- ② An `environment` directive defined within a `stage` will only apply the given environment variables to steps within the `stage`.

Parameters

Declarative Pipeline supports parameters out-of-the-box, allowing the Pipeline to accept user-specified parameters at runtime via the [parameters directive](#). Configuring parameters with Scripted Pipeline is done with the `properties` step, which can be found in the Snippet Generator.

If you configured your pipeline to accept parameters using the **Build with Parameters** option, those parameters are accessible as members of the `params` variable.

Assuming that a String parameter named "Greeting" has been configuring in the `Jenkinsfile`, it can access that parameter via `${params.Greeting}`:

```
// Declarative //
pipeline {
    agent any
    parameters {
        string(name: 'Greeting', defaultValue: 'Hello', description: 'How should I
greet the world?')
    }
    stages {
        stage('Example') {
            steps {
                echo "${params.Greeting} World!"
            }
        }
    }
}

// Script //
properties([parameters([string(defaultValue: 'Hello', description: 'How should I greet
the world?', name: 'Greeting'))]))

node {
    echo "${params.Greeting} World!"
}
```

Handling Failures

Declarative Pipeline supports robust failure handling by default via its [post section](#) which allows declaring a number of different "post conditions" such as: `always`, `unstable`, `success`, `failure`, and `changed`. The [Pipeline Syntax](#) section provides more detail on how to use the various post conditions.

```

// Declarative //
pipeline {
  agent any
  stages {
    stage('Test') {
      steps {
        sh 'make check'
      }
    }
  }
  post {
    always {
      junit '**/target/*.xml'
    }
    failure {
      mail to: team@example.com, subject: 'The Pipeline failed :( '
    }
  }
}
// Script //
node {
  /* .. snip .. */
  stage('Test') {
    try {
      sh 'make check'
    }
    finally {
      junit '**/target/*.xml'
    }
  }
  /* .. snip .. */
}

```

Scripted Pipeline however relies on Groovy's built-in `try/catch/finally` semantics for handling failures during execution of the Pipeline.

In the `[test]` example above, the `sh` step was modified to never return a non-zero exit code (`sh 'make check || true'`). This approach, while valid, means the following stages need to check `currentBuild.result` to know if there has been a test failure or not.

An alternative way of handling this, which preserves the early-exit behavior of failures in Pipeline, while still giving `junit` the chance to capture test reports, is to use a series of `try/finally` blocks:

Using multiple agents

In all the previous examples, only a single agent has been used. This means Jenkins will allocate an

executor wherever one is available, regardless of how it is labeled or configured. Not only can this behavior be overridden, but Pipeline allows utilizing multiple agents in the Jenkins environment from within the *same Jenkinsfile*, which can be helpful for more advanced use-cases such as executing builds/tests across multiple platforms.

In the example below, the "Build" stage will be performed on one agent and the built results will be reused on two subsequent agents, labelled "linux" and "windows" respectively, during the "Test" stage.

```
// Declarative //
pipeline {
  agent none
  stages {
    stage('Build') {
      agent any
      steps {
        checkout scm
        sh 'make'
        stash includes: '**/target/*.jar', name: 'app' ①
      }
    }
    stage('Test on Linux') {
      agent { ②
        label 'linux'
      }
      steps {
        unstash 'app' ③
        sh 'make check'
      }
      post {
        always {
          junit '**/target/*.xml'
        }
      }
    }
    stage('Test on Windows') {
      agent {
        label 'windows'
      }
      steps {
        unstash 'app'
        bat 'make check' ④
      }
      post {
        always {
          junit '**/target/*.xml'
        }
      }
    }
  }
}
```

```

}
// Script //
stage('Build') {
    node {
        checkout scm
        sh 'make'
        stash includes: '**/target/*.jar', name: 'app' ①
    }
}

stage('Test') {
    node('linux') { ②
        checkout scm
        try {
            unstash 'app' ③
            sh 'make check'
        }
        finally {
            junit '**/target/*.xml'
        }
    }
    node('windows') {
        checkout scm
        try {
            unstash 'app'
            bat 'make check' ④
        }
        finally {
            junit '**/target/*.xml'
        }
    }
}
}

```

- ① The **stash** step allows capturing files matching an inclusion pattern (****/target/*.jar**) for reuse within the *same* Pipeline. Once the Pipeline has completed its execution, stashed files are deleted from the Jenkins master.
- ② The parameter in **agent/node** allows for any valid Jenkins label expression. Consult the [Pipeline Syntax](#) section for more details.
- ③ **unstash** will retrieve the named "stash" from the Jenkins master into the Pipeline's current workspace.
- ④ The **bat** script allows for executing batch scripts on Windows-based platforms.

Optional step arguments

Pipeline follows the Groovy language convention of allowing parentheses to be omitted around method arguments.

Many Pipeline steps also use the named-parameter syntax as a shorthand for creating a Map in Groovy, which uses the syntax **[key1: value1, key2: value2]**. Making statements like the following

functionally equivalent:

```
git url: 'git://example.com/amazing-project.git', branch: 'master'  
git([url: 'git://example.com/amazing-project.git', branch: 'master'])
```

For convenience, when calling steps taking only one parameter (or only one mandatory parameter), the parameter name may be omitted, for example:

```
sh 'echo hello' /* short form */  
sh([script: 'echo hello']) /* long form */
```

Advanced Scripted Pipeline

Scripted Pipeline is a domain-specific language [16: en.wikipedia.org/wiki/Domain-specific_language] based on Groovy, most [Groovy syntax](#) can be used in Scripted Pipeline without modification.

Executing in parallel

The example in the [section above](#) runs tests across two different platforms in a linear series. In practice, if the `make check` execution takes 30 minutes to complete, the "Test" stage would now take 60 minutes to complete!

Fortunately, Pipeline has built-in functionality for executing portions of Scripted Pipeline in parallel, implemented in the aptly named `parallel` step.

Refactoring the example above to use the `parallel` step:


```
// Script //
stage('Build') {
    /* .. snip .. */
}

stage('Test') {
    parallel linux: {
        node('linux') {
            checkout scm
            try {
                unstash 'app'
                sh 'make check'
            }
            finally {
                junit '**/target/*.xml'
            }
        }
    },
    windows: {
        node('windows') {
            /* .. snip .. */
        }
    }
}
// Declarative not yet implemented //
```

Instead of executing the tests on the "linux" and "windows" labelled nodes in series, they will now execute in parallel assuming the requisite capacity exists in the Jenkins environment.

Branches and Pull Requests

In the [previous section](#) a `Jenkinsfile` which could be checked into source control was implemented. This section covers the concept of **Multibranch** Pipelines which build on the `Jenkinsfile` foundation to provide more dynamic and automatic functionality in Jenkins.

Creating a Multibranch Pipeline

The **Multibranch Pipeline** project type enables you to implement different Jenkinsfiles for different branches of the same project. In a Multibranch Pipeline project, Jenkins automatically discovers, manages and executes Pipelines for branches which contain a **Jenkinsfile** in source control.

This eliminates the need for manual Pipeline creation and management.

To create a Multibranch Pipeline:

- Click **New Item** on Jenkins home page.



- Enter a name for your Pipeline, select **Multibranch Pipeline** and click **OK**.

CAUTION

Jenkins uses the name of the Pipeline to create directories on disk. Pipeline names which include spaces may uncover bugs in scripts which do not expect paths to contain spaces.

Enter an item name

an-example

» Required field



Freestyle project

This is the central feature of Jenkins. Jenkins will build your project, combining any SCM with any build system, and this can be even used for something other than software build.



Pipeline

Orchestrates long-running activities that can span multiple build slaves. Suitable for building pipelines (formerly known as workflows) and/or organizing complex activities that do not easily fit in free-style job type.



External Job

This type of job allows you to record the execution of a process run outside Jenkins, even on a remote machine. This is designed so that you can use Jenkins as a dashboard of your existing automation system.



Multi-configuration project

Suitable for projects that need a large number of different configurations, such as testing on multiple environments, platform-specific builds, etc.



Folder

Creates a container that stores nested items in it. Useful for grouping things together. Unlike view, which is just a filter, a folder creates a separate namespace, so you can have multiple things of the same name as long as they are in different folders.



GitHub Organization

Scans a GitHub organization (or user account) for all repositories matching some defined markers.



Multibranch Pipeline

Creates a set of Pipeline projects according to detected branches in one SCM repository.

OK

- Add a **Branch Source** (for example, Git) and enter the location of the repository.

Name

Display Name

Description

[Plain text] [Preview](#)

Branch Sources

Add source ▼

- Git**
- GitHub
- Single repository & branch
- Subversion

Trigger builds remotely (e.g., from scripts) ?

Git

Project Repository

Credentials - none - ▼ Add ▼

Ignore on push notifications

Repository browser (Auto) ▼ ?

Additional Behaviours Add ▼

Advanced...

Property strategy All branches get the same properties ▼

Add property ▼

Delete source

- **Save** the Multibranch Pipeline project.

Upon **Save**, Jenkins automatically scans the designated repository and creates appropriate items for each branch in the repository which contains a **Jenkinsfile**.

By default, Jenkins will not automatically re-index the repository for branch additions or deletions (unless using an [Organization Folder](#)), so it is often useful to configure a Multibranch Pipeline to periodically re-index in the configuration:

Build Triggers



Trigger builds remotely (e.g., from scripts) 

Build periodically 

Periodically if not otherwise run 

Interval 

Additional Environment Variables

Multibranch Pipelines expose additional information about the branch being built through the `env` global variable, such as:

BRANCH_NAME

Name of the branch for which this Pipeline is executing, for example `master`.

CHANGE_ID

An identifier corresponding to some kind of change request, such as a pull request number

Additional environment variables are listed in the [Global Variable Reference](#).

Supporting Pull Requests

With the "GitHub" or "Bitbucket" Branch Sources, Multibranch Pipelines can be used for validating pull/change requests. This functionality is provided, respectively, by the `plugin:github-branch-source`[GitHub Branch Source] and `plugin:cloudbees-bitbucket-branch-source`[Bitbucket Branch Source] plugins. Please consult their documentation for further information on how to use those plugins.

Using Organization Folders

Organization Folders enable Jenkins to monitor an entire GitHub Organization, or Bitbucket Team/Project and automatically create new Multibranch Pipelines for repositories which contain branches and pull requests containing a `Jenkinsfile`.

Currently, this functionality exists only for GitHub and Bitbucket, with functionality provided by the `plugin:github-organization-folder`[GitHub Organization Folder] and `plugin:cloudbees-bitbucket-branch-source`[Bitbucket Branch Source] plugins.

Using Docker with Pipeline

Many organizations use [Docker](#) to unify their build and test environments across machines, and to provide an efficient mechanism for deploying applications. Starting with Pipeline versions 2.5 and higher, Pipeline has built-in support for interacting with Docker from within a [Jenkinsfile](#).

While this section will cover the basics of utilizing Docker from with a [Jenkinsfile](#), it will not cover the fundamentals of Docker, which can be read about in the [Docker Getting Started Guide](#).

Customizing the execution environment

Pipeline is designed to easily use [Docker](#) images as the execution environment for a single [Stage](#) or the entire Pipeline. Meaning that a user can define the tools required for their Pipeline, without having to manually configure agents. Practically any tool which can be [packaged in a Docker container](#). can be used with ease by making only minor edits to a [Jenkinsfile](#).

```
// Declarative //
pipeline {
  agent {
    docker { image 'node:7-alpine' }
  }
  stages {
    stage('Test') {
      steps {
        sh 'node --version'
      }
    }
  }
}

// Script //
node {
  /* Requires the Docker Pipeline plugin to be installed */
  docker.image('node:7-alpine').inside {
    stage('Test') {
      sh 'node --version'
    }
  }
}
```

When the Pipeline executes, Jenkins will automatically start the specified container and execute the defined steps within it:

```
[Pipeline] stage
[Pipeline] { (Test)
[Pipeline] sh
[guided-tour] Running shell script
+ node --version
v7.4.0
[Pipeline] }
[Pipeline] // stage
[Pipeline] }
```

Caching data for containers

Many build tools will download external dependencies and cache them locally for future re-use. Since containers are initially created with "clean" file systems, this can result in slower Pipelines, as

they may not take advantage of on-disk caches between subsequent Pipeline runs.

Pipeline supports adding custom arguments which are passed to Docker, allowing users to specify custom [Docker Volumes](#) to mount, which can be used for caching data on the [agent](#) between Pipeline runs. The following example will cache `~/m2` between Pipeline runs utilizing the [maven container](#), thereby avoiding the need to re-download dependencies for subsequent runs of the Pipeline.

```
// Declarative //
pipeline {
  agent {
    docker {
      image 'maven:3-alpine'
      args '-v $HOME/.m2:/root/.m2'
    }
  }
  stages {
    stage('Build') {
      steps {
        sh 'mvn -B'
      }
    }
  }
}

// Script //
node {
  /* Requires the Docker Pipeline plugin to be installed */
  docker.image('maven:3-alpine').inside('-v $HOME/.m2:/root/.m2') {
    stage('Build') {
      sh 'mvn -B'
    }
  }
}
```

Using multiple containers

It has become increasingly common for code bases to rely on multiple, different, technologies. For example, a repository might have both a Java-based back-end API implementation *and* a JavaScript-based front-end implementation. Combining Docker and Pipeline allows a [Jenkinsfile](#) to use **multiple** types of technologies by combining the [agent {}](#) directive, with different stages.

```

// Declarative //
pipeline {
  agent none
  stages {
    stage('Back-end') {
      agent {
        docker { image 'maven:3-alpine' }
      }
      steps {
        sh 'mvn --version'
      }
    }
    stage('Front-end') {
      agent {
        docker { image 'node:7-alpine' }
      }
      steps {
        sh 'node --version'
      }
    }
  }
}
// Script //
node {
  /* Requires the Docker Pipeline plugin to be installed */

  stage('Back-end') {
    docker.image('maven:3-alpine').inside {
      sh 'mvn --version'
    }
  }

  stage('Front-end') {
    docker.image('node:7-alpine').inside {
      sh 'node --version'
    }
  }
}
}

```

Using a Dockerfile

For projects which require a more customized execution environment, Pipeline also supports building and running a container from a **Dockerfile** in the source repository. In contrast to the [previous approach](#) of using an "off-the-shelf" container, using the `agent { dockerfile true }` syntax will build a new image from a **Dockerfile** rather than pulling one from [Docker Hub](#).

Re-using an example from above, with a more custom **Dockerfile**:

Dockerfile

```
FROM node:7-alpine

RUN apk add -U subversion
```

By committing this to the root of the source repository, the **Jenkinsfile** can be changed to build a container based on this **Dockerfile** and then run the defined steps using that container:

```
// Declarative //
pipeline {
  agent { dockerfile true }
  stages {
    stage('Test') {
      steps {
        sh 'node --version'
        sh 'svn --version'
      }
    }
  }
}
// Script //
```

The `agent { dockerfile true }` syntax supports a number of other options which are described in more detail in the [Pipeline Syntax](#) section.



Example



Steps Example



✓	> Record trace of a Docker image used in FROM	<1s
✓	> docker build -t 8ccfe61fd542d23ce0abf929ebade3e62b7aedf8 -f "Dockerfile" "." -- Shell Script	<1s
✓	> Dockerfile -- Read file from workspace	<1s
✓	> General SCM	<1s
✓	> Hello World! -- Print Message	<1s
✓	✓ echo myCustomEnvVar = \$myCustomEnvVar -- Shell Script	<1s

```
1 [ne-dockerfile-sample_master-WS6YXIGKD7VZMNTXWB5A7QQYU0JTQ7ZKI6JR62T66CTEH43QKGJA] Running shell script
2 + echo myCustomEnvVar = This is a sample.
3 myCustomEnvVar = This is a sample.
```

Using a Dockerfile with Jenkins Pipeline

Specifying a Docker Label

By default, Pipeline assumes that *any* configured [agent](#) is capable of running Docker-based Pipelines. For Jenkins environments which have macOS, Windows, or other agents, which are unable to run the Docker daemon, this default setting may be problematic. Pipeline provides a global option in the **Manage Jenkins** page, and on the [Folder](#) level, for specifying which agents (by [Label](#)) to use for running Docker-based Pipelines.

Pipeline Model Definition

Docker Label	<input type="text"/>	
Docker registry URL	<input type="text"/>	
Registry credentials	<input type="text" value="- none -"/>	<input type="button" value="Add"/>

Advanced Usage with Scripted Pipeline

Running "sidecar" containers

Using Docker in Pipeline can be an effective way to run a service on which the build, or a set of tests, may rely. Similar to the [sidecar pattern](#), Docker Pipeline can run one container "in the background", while performing work in another. Utilizing this sidecar approach, a Pipeline can have a "clean" container provisioned for each Pipeline run.

Consider a hypothetical integration test suite which relies on a local MySQL database to be running. Using the `withRun` method, implemented in the `plugin:docker-workflow[Docker Pipeline]` plugin's support for Scripted Pipeline, a `Jenkinsfile` can run MySQL as a sidecar:

```
node {
  checkout scm
  /*
   * In order to communicate with the MySQL server, this Pipeline explicitly
   * maps the port (`3306`) to a known port on the host machine.
   */
  docker.image('mysql:5').withRun('-e "MYSQL_ROOT_PASSWORD=my-secret-pw" -p
3306:3306') { c ->
    /* Wait until mysql service is up */
    sh 'while ! mysqladmin ping -h0.0.0.0 --silent; do sleep 1; done'
    /* Run some tests which require MySQL */
    sh 'make check'
  }
}
```

This example can be taken further, utilizing two containers simultaneously. One "sidecar" running MySQL, and another providing the [execution environment](#), by using the Docker [container links](#).

```
node {
  checkout scm
  docker.image('mysql:5').withRun('-e "MYSQL_ROOT_PASSWORD=my-secret-pw"') { c ->
    docker.image('mysql:5').inside("--link ${c.id}:db") {
      /* Wait until mysql service is up */
      sh 'while ! mysqladmin ping -hdb --silent; do sleep 1; done'
    }
    docker.image('centos:7').inside("--link ${c.id}:db") {
      /*
       * Run some tests which require MySQL, and assume that it is
       * available on the host name `db`
       */
      sh 'make check'
    }
  }
}
```

The above example uses the object exposed by `withRun`, which has the running container's ID available via the `id` property. Using the container's ID, the Pipeline can create a link by passing custom Docker arguments to the `inside()` method.

The `id` property can also be useful for inspecting logs from a running Docker container before the Pipeline exits:

```
sh "docker logs ${c.id}"
```

Building containers

In order to create a Docker image, the `plugin:docker-workflow[Docker Pipeline]` plugin also provides a `build()` method for creating a new image, from a `Dockerfile` in the repository, during a Pipeline run.

One major benefit of using the syntax `docker.build("my-image-name")` is that a Scripted Pipeline can use the return value for subsequent Docker Pipeline calls, for example:

```
node {
  checkout scm

  def customImage = docker.build("my-image:${env.BUILD_ID}")

  customImage.inside {
    sh 'make test'
  }
}
```

The return value can also be used to publish the Docker image to [Docker Hub](#), or a [custom Registry](#), via the `push()` method, for example:

```
node {
  checkout scm
  def customImage = docker.build("my-image:${env.BUILD_ID}")
  customImage.push()
}
```

One common usage of image "tags" is to specify a `latest` tag for the most recently, validated, version of a Docker image. The `push()` method accepts an optional `tag` parameter, allowing the Pipeline to push the `customImage` with different tags, for example:

```
node {
  checkout scm
  def customImage = docker.build("my-image:${env.BUILD_ID}")
  customImage.push()

  customImage.push('latest')
}
```

Using a remote Docker server

By default, the `plugin:docker-workflow` [Docker Pipeline] plugin will communicate with a local Docker daemon, typically accessed through `/var/run/docker.sock`.

To select a non-default Docker server, such as with [Docker Swarm](#), the `withServer()` method should be used.

By passing a URI, and optionally the Credentials ID of a **Docker Server Certificate Authentication** pre-configured in Jenkins, to the method with:

```
node {
  checkout scm

  docker.withServer('tcp://swarm.example.com:2376', 'swarm-certs') {
    docker.image('mysql:5').withRun('-p 3306:3306') {
      /* do things */
    }
  }
}
```

`inside()` and `build()` will not work properly with a Docker Swarm server out of the box

For `inside()` to work, the Docker server and the Jenkins agent must use the same filesystem, so that the workspace can be mounted.

Currently neither the Jenkins plugin nor the Docker CLI will automatically detect the case that the server is running remotely; a typical symptom would be errors from nested `sh` commands such as

CAUTION

```
cannot create /...@tmp/durable-.../pid: Directory nonexistent
```

When Jenkins detects that the agent is itself running inside a Docker container, it will automatically pass the `--volumes-from` argument to the `inside` container, ensuring that it can share a workspace with the agent.

Additionally some versions of Docker Swarm do not support custom Registries.

Using a custom registry

By default the plugin:docker-workflow[Docker Pipeline] integrates assumes the default Docker Registry of [Docker Hub](#).

In order to use a custom Docker Registry, users of Scripted Pipeline can wrap steps with the `withRegistry()` method, passing in the custom Registry URL, for example:

```
node {
    checkout scm

    docker.withRegistry('https://registry.example.com') {

        docker.image('my-custom-image').inside {
            sh 'make test'
        }
    }
}
```

For a Docker Registry which requires authentication, add a "Username/Password" Credentials item from the Jenkins home page and use the Credentials ID as a second argument to `withRegistry()`:

```
node {
    checkout scm

    docker.withRegistry('https://registry.example.com', 'credentials-id') {

        def customImage = docker.build("my-image:${env.BUILD_ID}")

        /* Push the container to the custom Registry */
        customImage.push()
    }
}
```

Extending with Shared Libraries

As Pipeline is adopted for more and more projects in an organization, common patterns are likely to emerge. Oftentimes it is useful to share parts of Pipelines between various projects to reduce redundancies and keep code "DRY" [17: en.wikipedia.org/wiki/Don't_repeat_yourself].

Pipeline has support for creating "Shared Libraries" which can be defined in external source control repositories and loaded into existing Pipelines.

Defining Shared Libraries

A Shared Library is defined with a name, a source code retrieval method such as by SCM, and optionally a default version. The name should be a short identifier as it will be used in scripts.

The version could be anything understood by that SCM; for example, branches, tags, and commit hashes all work for Git. You may also declare whether scripts need to explicitly request that library (detailed below), or if it is present by default. Furthermore, if you specify a version in Jenkins configuration, you can block scripts from selecting a *different* version.

The best way to specify the SCM is using an SCM plugin which has been specifically updated to support a new API for checking out an arbitrary named version (*Modern SCM* option). As of this writing, the latest versions of the Git and Subversion plugins support this mode; others should follow.

If your SCM plugin has not been integrated, you may select *Legacy SCM* and pick anything offered. In this case, you need to include `${library.yourLibName.version}` somewhere in the configuration of the SCM, so that during checkout the plugin will expand this variable to select the desired version. For example, for Subversion, you can set the *Repository URL* to `svnserver/project/${library.yourLibName.version}` and then use versions such as `trunk` or `branches/dev` or `tags/1.0`.

Directory structure

The directory structure of a Shared Library repository is as follows:

```
(root)
+- src                # Groovy source files
|   +- org
|       +- foo
|           +- Bar.groovy # for org.foo.Bar class
+- vars
|   +- foo.groovy      # for global 'foo' variable
|   +- foo.txt        # help for 'foo' variable
+- resources          # resource files (external libraries only)
|   +- org
|       +- foo
|           +- bar.json # static helper data for org.foo.Bar
```

The `src` directory should look like standard Java source directory structure. This directory is added to the classpath when executing Pipelines.

The `vars` directory hosts scripts that define global variables accessible from Pipeline. The basename of each `.groovy` file should be a Groovy (~ Java) identifier, conventionally camelCased. The matching `.txt`, if present, can contain documentation, processed through the system's configured markup formatter (so may really be HTML, Markdown, etc., though the `txt` extension is required).

The Groovy source files in these directories get the same “CPS transformation” as in Scripted Pipeline.

A `resources` directory allows the `libraryResource` step to be used from an external library to load associated non-Groovy files. Currently this feature is not supported for internal libraries.


Other directories under the root are reserved for future enhancements.

Global Shared Libraries

There are several places where Shared Libraries can be defined, depending on the use-case. *Manage Jenkins » Configure System » Global Pipeline Libraries* as many libraries as necessary can be configured.

Global Pipeline Libraries

Sharable libraries available to any Pipeline jobs running on this system. These libraries will be trusted, meaning they run without "sandbox" restrictions and may use `@Grab`.



Since these libraries will be globally usable, any Pipeline in the system can utilize functionality implemented in these libraries.

These libraries are considered "trusted:" they can run any methods in Java, Groovy, Jenkins internal APIs, Jenkins plugins, or third-party libraries. This allows you to define libraries which encapsulate individually unsafe APIs in a higher-level wrapper safe for use from any Pipeline. Beware that **anyone able to push commits to this SCM repository could obtain unlimited access to Jenkins**. You need the *Overall/RunScripts* permission to configure these libraries (normally this will be granted to Jenkins administrators).

Folder-level Shared Libraries

Any Folder created can have Shared Libraries associated with it. This mechanism allows scoping of specific libraries to all the Pipelines inside of the folder or subfolder.

Folder-based libraries are not considered "trusted:" they run in the Groovy sandbox just like typical Pipelines.

Automatic Shared Libraries

Other plugins may add ways of defining libraries on the fly. For example, the `plugin:github-branch-source[GitHub Branch Source]` plugin provides a "GitHub Organization Folder" item which allows a script to use an untrusted library such as `github.com/someorg/somerepo` without any additional configuration. In this case, the specified GitHub repository would be loaded, from the `master` branch, using an anonymous checkout.

Using libraries

Shared Libraries marked *Load implicitly* allows Pipelines to immediately use classes or global variables defined by any such libraries. To access other shared libraries, the `Jenkinsfile` needs to use the `@Library` annotation, specifying the library's name:

Global Pipeline Libraries

Sharable libraries available to any Pipeline jobs running on this system. These libraries will be trusted, meaning they run without "sandbox" restrictions and may use `@Grab`.

The screenshot shows the configuration for a global pipeline library. The 'Name' field contains 'my-shared-library'. The 'Default version' field contains 'master'. The 'Load implicitly' checkbox is unchecked. The 'Allow default version to be overridden' checkbox is checked. Under the 'Retrieval method' section, 'Modern SCM' is selected with a radio button.

```
@Library('my-shared-library') _  
/* Using a version specifier, such as branch, tag, etc */  
@Library('my-shared-library@1.0') _  
/* Accessing multiple libraries with one statement */  
@Library(['my-shared-library', 'otherlib@abc1234']) _
```

The annotation can be anywhere in the script where an annotation is permitted by Groovy. When referring to class libraries (with `src/` directories), conventionally the annotation goes on an `import` statement:

```
@Library('somelib')  
import com.mycorp.pipeline.somelib.UsefulClass
```

TIP

For Shared Libraries which only define Global Variables (`<code>vars</code>`), or a `<code>Jenkinsfile</code>` which only needs a Global Variable, the [annotation](http://groovy-lang.org/objectorientation.html#annotation) pattern `<code>@Library('my-shared-library') _</code>` may be useful for keeping code concise. In essence, instead of annotating an unnecessary `<code>import</code>` statement, the symbol `<code></code>` is annotated.

It is not recommended to `import` a global variable/function, since this will force the compiler to interpret fields and methods as `static` even if they were intended to be instance. The Groovy compiler in this case can produce confusing error messages.

Libraries are resolved and loaded during *compilation* of the script, before it starts executing. This allows the Groovy compiler to understand the meaning of symbols used in static type checking, and

permits them to be used in type declarations in the script, for example:

```
@Library('somelib')
import com.mycorp.pipeline.somelib.Helper

int useSomeLib(Helper helper) {
    helper.prepare()
    return helper.count()
}

echo useSomeLib(new Helper('some text'))
```

Global Variables however, are resolved at runtime.

Loading libraries dynamically

As of version 2.7 of the *Pipeline: Shared Groovy Libraries* plugin, there is a new option for loading (non-implicit) libraries in a script: a **library** step that loads a library *dynamically*, at any time during the build.

If you are only interested in using global variables/functions (from the **vars/** directory), the syntax is quite simple:

```
library 'my-shared-library'
```

Thereafter, any global variables from that library will be accessible to the script.

Using classes from the **src/** directory is also possible, but trickier. Whereas the **@Library** annotation prepares the “classpath” of the script prior to compilation, by the time a **library** step is encountered the script has already been compiled. Therefore you cannot **import** or otherwise “statically” refer to types from the library.

However you may use library classes dynamically (without type checking), accessing them by fully-qualified name from the return value of the **library** step. **static** methods can be invoked using a Java-like syntax:

```
library('my-shared-library').com.mycorp.pipeline.Utils.someStaticMethod()
```

You can also access **static** fields, and call constructors as if they were **static** methods named **new**:

```

def useSomeLib(helper) { // dynamic: cannot declare as Helper
    helper.prepare()
    return helper.count()
}

def lib = library('my-shared-library').com.mycorp.pipeline // preselect the package

echo useSomeLib(lib.Helper.new(lib.Constants.SOME_TEXT))

```

Library versions

The "Default version" for a configured Shared Library is used when "Load implicitly" is checked, or if a Pipeline references the library only by name, for example `@Library('my-shared-library')` `_`. If a "Default version" is **not** defined, the Pipeline must specify a version, for example `@Library('my-shared-library@master')` `_`.

If "Allow default version to be overridden" is enabled in the Shared Library's configuration, a `@Library` annotation may also override a default version defined for the library. This also allows a library with "Load implicitly" to be loaded from a different version if necessary.

When using the `library` step you may also specify a version:

```
library 'my-shared-library@master'
```

Since this is a regular step, that version could be *computed* rather than a constant as with the annotation; for example:

```
library "my-shared-library@$BRANCH_NAME"
```


would load a library using the same SCM branch as the multibranch `Jenkinsfile`. As another example, you could pick a library by parameter:


```
properties([parameters([string(name: 'LIB_VERSION', defaultValue: 'master')]))
library "my-shared-library@${params.LIB_VERSION}"
```


Note that the `library` step may not be used to override the version of an implicitly loaded library. It is already loaded by the time the script starts, and a library of a given name may not be loaded twice.


Retrieval Method


The best way to specify the SCM is using an SCM plugin which has been specifically updated to support a new API for checking out an arbitrary named version (**Modern SCM** option). As of this writing, the latest versions of the Git and Subversion plugins support this mode.

 Library


Name 

Default version 

Load implicitly 

Allow default version to be overridden 



Retrieval method

Modern SCM 


Source Code Management

Git

Project Repository

Credentials  Add 

Ignore on push notifications

Repository browser 

Legacy SCM

SCM plugins which have not yet been updated to support the newer features required by Shared Libraries, may still be used via the **Legacy SCM** option. In this case, include `${library.yourlibrarynamehere.version}` wherever a branch/tag/ref may be configured for that particular SCM plugin. This ensures that during checkout of the library's source code, the SCM plugin will expand this variable to checkout the appropriate version of the library.

Library

Name ?

Default version ?

Load implicitly ?

Allow default version to be overridden ?

Retrieval method

Modern SCM ?

Legacy SCM ?

Source Code Management

Git

Subversion ?

Modules

Repository URL ?

⚠ This repository URL is parameterized, syntax validation skipped

Credentials Add ?

⚠ The repository URL is parameterized, connection check skipped

Local module directory ?

Repository depth ?

Ignore externals ?

Dynamic retrieval

If you only specify a library name (optionally with version after @) in the `library` step, Jenkins will look for a preconfigured library of that name. (Or in the case of a `github.com/owner/repo` automatic library it will load that.)

But you may also specify the retrieval method dynamically, in which case there is no need for the library to have been predefined in Jenkins. Here is an example:

```
library identifier: 'custom-lib@master', retriever: modernSCM(
  [$class: 'GitSCMSource',
   remote: 'git@git.mycorp.com:my-jenkins-utils.git',
   credentialsId: 'my-private-key'])
```

It is best to refer to **Pipeline Syntax** for the precise syntax for your SCM.

Note that the library version *must* be specified in these cases.

Writing libraries

At the base level, any valid [Groovy code](#) is okay for use. Different data structures, utility methods, etc, such as:

```
// src/org/foo/Point.groovy
package org.foo;

// point in 3D space
class Point {
    float x,y,z;
}
```

Accessing steps

Library classes cannot directly call steps such as `sh` or `git`. They can however implement methods, outside of the scope of an enclosing class, which in turn invoke Pipeline steps, for example:

```
// src/org/foo/Zot.groovy
package org.foo;

def checkOutFrom(repo) {
    git url: "git@github.com:jenkinsci/${repo}"
}
```

Which can then be called from a Scripted Pipeline:

```
def z = new org.foo.Zot()
z.checkOutFrom(repo)
```

This approach has limitations; for example, it prevents the declaration of a superclass.

Alternately, a set of `steps` can be passed explicitly using `this` to a library class, in a constructor, or just one method:

```
package org.foo
class Utilities implements Serializable {
    def steps
    Utilities(steps) {this.steps = steps}
    def mvn(args) {
        steps.sh "${steps.tool 'Maven'}/bin/mvn -o ${args}"
    }
}
```

When saving state on classes, such as above, the class **must** implement the `Serializable` interface.

This ensures that a Pipeline using the class, as seen in the example below, can properly suspend and resume in Jenkins.

```
@Library('utils') import org.foo.Utilities
def utils = new Utilities(this)
node {
    utils.mvn 'clean package'
}
```

If the library needs to access global variables, such as `env`, those should be explicitly passed into the library classes, or methods, in a similar manner.

Instead of passing numerous variables from the Scripted Pipeline into a library,

```
package org.foo
class Utilities {
    static def mvn(script, args) {
        script.sh "${script.tool 'Maven'}/bin/mvn -s ${script.env.HOME}/jenkins.xml -o
        ${args}"
    }
}
```

The above example shows the script being passed in to one `static` method, invoked from a Scripted Pipeline as follows:

```
@Library('utils') import static org.foo.Utilities.*
node {
    mvn this, 'clean package'
}
```

Defining global variables

Internally, scripts in the `vars` directory are instantiated on-demand as singletons. This allows multiple methods to be defined in a single `.groovy` file for convenience. For example:

var/log.groovy

```
def info(message) {
    echo "INFO: ${message}"
}

def warning(message) {
    echo "WARNING: ${message}"
}
```

Jenkinsfile

```
@Library('utils') _

log.info 'Starting'
log.warning 'Nothing to do!'
```

Declarative Pipeline does not allow global variable usage outside of a `script` directive ([JENKINS-42360](#)).

Jenkinsfile

```
@Library('utils') _

pipeline {
  agent none
  stage ('Example') {
    steps {
      script { ①
        log.info 'Starting'
        log.warning 'Nothing to do!'
      }
    }
  }
}
```

① `script` directive required to access global variables in Declarative Pipeline.

NOTE

A variable defined in a shared library will only show up in *Global Variables Reference* (under *Pipeline Syntax*) after Jenkins loads and uses that library as part of a successful Pipeline run.

WARNING

Avoid preserving state in global variables

Avoid defining global variables with methods that interact or preserve state. Use a static class or instantiate a local variable of a class instead.

Defining custom steps

Shared Libraries can also define global variables which behave similarly to built-in steps, such as `sh` or `git`. Global variables defined in Shared Libraries **must** be named with all lower-case or "camelCased" in order to be loaded properly by Pipeline. [18: gist.github.com/rtyler/e5e57f075af381fce4ed3ae57aa1f0c2]

For example, to define `sayHello`, the file `vars/sayHello.groovy` should be created and should implement a `call` method. The `call` method allows the global variable to be invoked in a manner similar to a step:

```
// vars/sayHello.groovy
def call(String name = 'human') {
    // Any valid steps can be called from this code, just like in other
    // Scripted Pipeline
    echo "Hello, ${name}."
}
```

The Pipeline would then be able to reference and invoke this variable:

```
sayHello 'Joe'
sayHello() /* invoke with default arguments */
```

If called with a block, the `call` method will receive a `Closure`. The type should be defined explicitly to clarify the intent of the step, for example:

```
// vars/windows.groovy
def call(Closure body) {
    node('windows') {
        body()
    }
}
```

The Pipeline can then use this variable like any built-in step which accepts a block:

```
windows {
    bat "cmd /?"
}
```

Defining a more structured DSL

If you have a lot of Pipelines that are mostly similar, the global variable mechanism provides a handy tool to build a higher-level DSL that captures the similarity. For example, all Jenkins plugins are built and tested in the same way, so we might write a step named `buildPlugin`:

```
// vars/buildPlugin.groovy
def call(body) {
    // evaluate the body block, and collect configuration into the object
    def config = [:]
    body.resolveStrategy = Closure.DELEGATE_FIRST
    body.delegate = config
    body()

    // now build, based on the configuration provided
    node {
        git url: "https://github.com/jenkinsci/${config.name}-plugin.git"
        sh "mvn install"
        mail to: "...", subject: "${config.name} plugin build", body: "..."
    }
}
}
```

Assuming the script has either been loaded as a [Global Shared Library](#) or as a [Folder-level Shared Library](#) the resulting `Jenkinsfile` will be dramatically simpler:

```
// Script //
buildPlugin {
    name = 'git'
}
// Declarative not yet implemented //
```

Using third-party libraries

It is possible to use third-party Java libraries, typically found in [Maven Central](#), from **trusted** library code using the `@Grab` annotation. Refer to the [Grape documentation](#) for details, but simply put:

```
@Grab('org.apache.commons:commons-math3:3.4.1')
import org.apache.commons.math3.primes.Primes
void parallelize(int count) {
    if (!Primes.isPrime(count)) {
        error "${count} was not prime"
    }
    // ...
}
```

Third-party libraries are cached by default in `~/.groovy/grapes/` on the Jenkins master.

Loading resources

External libraries may load adjunct files from a `resources/` directory using the `LibraryResource` step. The argument is a relative pathname, akin to Java resource loading:

```
def request = libraryResource 'com/mycorp/pipeline/somelib/request.json'
```

The file is loaded as a string, suitable for passing to certain APIs or saving to a workspace using `writeFile`.

It is advisable to use an unique package structure so you do not accidentally conflict with another library.

Pretesting library changes

If you notice a mistake in a build using an untrusted library, simply click the *Replay* link to try editing one or more of its source files, and see if the resulting build behaves as expected. Once you are satisfied with the result, follow the diff link from the build's status page, and apply the diff to the library repository and commit.

(Even if the version requested for the library was a branch, rather than a fixed version like a tag, replayed builds will use the exact same revision as the original build: library sources will not be checked out again.)

Replay is not currently supported for trusted libraries. Modifying resource files is also not currently supported during *Replay*.

Defining Declarative Pipelines

Starting with Declarative 1.2, released in late September, 2017, you can define Declarative Pipelines in your shared libraries as well. Here's an example, which will execute a different Declarative Pipeline depending on whether the build number is odd or even:

```
// vars/evenOrOdd.groovy
def call(int buildNumber) {
    if (buildNumber % 2 == 0) {
        pipeline {
            agent any
            stages {
                stage('Even Stage') {
                    steps {
                        echo "The build number is even"
                    }
                }
            }
        }
    } else {
        pipeline {
            agent any
            stages {
                stage('Odd Stage') {
                    steps {
                        echo "The build number is odd"
                    }
                }
            }
        }
    }
}
}
```

```
// Jenkinsfile
@Library('my-shared-library') _

evenOrOdd(currentBuild.getNumber())
```

Only entire pipeline's can be defined in shared libraries as of this time. This can only be done in `vars/*.groovy`, and only in a `call` method. Only one Declarative Pipeline can be executed in a single build, and if you attempt to execute a second one, your build will fail as a result.

Pipeline Development Tools

Jenkins Pipeline includes [built-in documentation](#) and the [Snippet Generator](#) which are key resources when developing Pipelines. They provide detailed help and information that is customized to the currently installed version of Jenkins and related plugins. In this section, we'll discuss other tools and resources that may help with development of Jenkins Pipelines.

Blue Ocean Editor

The [Blue Ocean Pipeline Editor](#) provides a [WYSIWYG](#) way to create Declarative Pipelines. The editor offers a structural view of all the stages, parallel branches, and steps in a Pipeline. The editor validates Pipeline changes as they are made, eliminating many errors before they are even committed. Behind the scenes it still generates Declarative Pipeline code.

Command-line Pipeline Linter

Jenkins can validate, or "lint", a Declarative Pipeline from the command line before actually running it. This can be done using a Jenkins CLI command or by making an HTTP POST request with appropriate parameters. We recommended using the [SSH interface](#) to run the linter. See the [Jenkins CLI documentation](#) for details on how to properly configure Jenkins for secure command-line access.

Linting via the CLI with SSH

```
# ssh (Jenkins CLI)
# JENKINS_SSHD_PORT=[sshd port on master]
# JENKINS_HOSTNAME=[Jenkins master hostname]
ssh -p $JENKINS_SSHD_PORT $JENKINS_HOSTNAME declarative-linter < Jenkinsfile
```

Linting via HTTP POST using curl

```
# curl (REST API)
# Assuming "anonymous read access" has been enabled on your Jenkins instance.
# JENKINS_URL=[root URL of Jenkins master]
# JENKINS_CRUMB is needed if your Jenkins master has CSRF protection enabled as it
should
JENKINS_CRUMB=`curl
"$JENKINS_URL/crumbsIssuer/api/xml?xpath=concat(//crumbRequestField,\":\",//crumb)"`
curl -X POST -H $JENKINS_CRUMB -F "jenkinsfile=<Jenkinsfile" $JENKINS_URL/pipeline-
model-converter/validate
```

Examples

Below are two examples of the Pipeline Linter in action. This first example shows the output of the linter when it is passed an invalid `Jenkinsfile`, one that is missing part of the `agent` declaration.

Jenkinsfile

```
pipeline {
  agent
  stages {
    stage ('Initialize') {
      steps {
        echo 'Placeholder.'
      }
    }
  }
}
```

Linter output for invalid Jenkinsfile

```
# pass a Jenkinsfile that does not contain an "agent" section
ssh -p 8675 localhost declarative-linter < ./Jenkinsfile
Errors encountered validating Jenkinsfile:
WorkflowScript: 2: Not a valid section definition: "agent". Some extra configuration
is required. @ line 2, column 3.
    agent
    ^

WorkflowScript: 1: Missing required section "agent" @ line 1, column 1.
pipeline &#125;
^
```

In this second example, the **Jenkinsfile** has been updated to include the missing **any** on **agent**. The linter now reports that the Pipeline is valid.

Jenkinsfile

```
pipeline {
  agent any
  stages {
    stage ('Initialize') {
      steps {
        echo 'Placeholder.'
      }
    }
  }
}
```

Linter output for valid Jenkinsfile

```
ssh -p 8675 localhost declarative-linter < ./Jenkinsfile
Jenkinsfile successfully validated.
```


"Replay" Pipeline Runs with Modifications

Typically a Pipeline will be defined inside of the classic Jenkins web UI, or by committing to a [Jenkinsfile](#) in source control. Unfortunately, neither approach is ideal for rapid iteration, or prototyping, of a Pipeline. The "Replay" feature allows for quick modifications and execution of an existing Pipeline without changing the Pipeline configuration or creating a new commit.

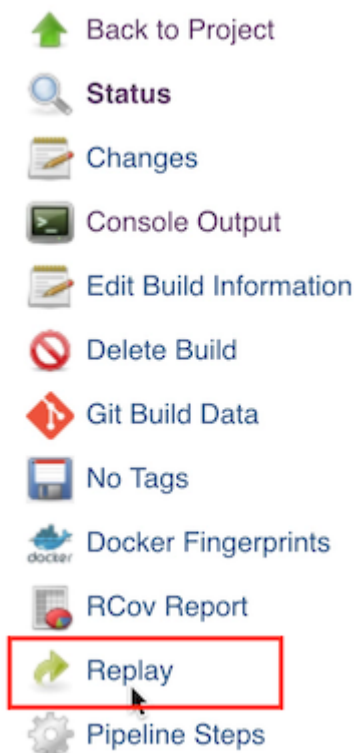
Usage

To use the "Replay" feature:

1. Select a previously completed run in the build history.



2. Click "Replay" in the left menu



3. Make modifications and click "Run". In this example, we changed "ruby-2.3" to "ruby-2.4".

/declarative/html > #4 > Replay

Replay #4

Allows you to replay a Pipeline build with a modified script. If any load steps were run, you can also modify the scripts they loaded.

```
Main Script
1  #!groovy
2
3  pipeline {
4  agent {
5    // Use docker container
6  docker {
7    image 'ruby:2.4'
8  }
9  }
10 options {
11   // Only keep the 10 most recent builds
12   buildDiscarder(logRotator(numToKeepStr:'10'))
13 }
14 stages {
15   stage ('Install') {
16
```

[Pipeline Syntax](#)

Run

4. Check the results of changes

Once you are satisfied with the changes, you can use Replay to view them again, copy them back to your Pipeline job or [Jenkinsfile](#), and then commit them using your usual engineering processes.

Features

- **Can be called multiple times on the same run** - allows for easy parallel testing of different changes.
- **Can also be called on Pipeline runs that are still in-progress** - As long as a Pipeline contained syntactically correct Groovy and was able to start, it can be Replayed.
- **Referenced Shared Library code is also modifiable** - If a Pipeline run references a [Shared Library](#), the code from the shared library will also be shown and modifiable as part of the Replay page.

Limitations

- **Pipeline runs with syntax errors cannot be replayed** - meaning their code cannot be viewed and any changes made in them cannot be retrieved. When using Replay for more significant modifications, save your changes to a file or editor outside of Jenkins before running them. See [JENKINS-37589](#)
- **Replayed Pipeline behavior may differ from runs started by other methods** - For Pipelines that are not part of a Multi-branch Pipeline, the commit information may differ for the original run and the Replayed run. See [JENKINS-36453](#)

Pipeline Unit Testing Framework

NOTE

The Pipeline Unit Testing Framework is a third-party tool that is not supported by the Jenkins Project.

The [Pipeline Unit Testing Framework](#) allows you to [unit test](#) Pipelines and [Shared Libraries](#) before running them in full. It provides a mock execution environment where real Pipeline steps are replaced with mock objects that you can use to check for expected behavior. New and rough around the edges, but promising. The [README](#) for that project contains examples and usage instructions.

Pipeline Syntax

This section builds on the information introduced in [Getting Started](#), and should be treated solely as a reference. For more information on how to use Pipeline syntax in practical examples, refer to [The Jenkinsfile](#) section of this chapter. As of version 2.5 of the Pipeline plugin, Pipeline supports two discrete syntaxes which are detailed below. For the pros and cons of each, see the [Syntax Comparison](#).

As discussed in [Getting Started](#), the most fundamental part of a Pipeline is the "step." Basically, steps tell Jenkins *what* to do, and serve as the basic building block for both Declarative and Scripted Pipeline syntax.

For an overview of available steps, please refer to the [Pipeline Steps reference](#) which contains a comprehensive list of steps built into Pipeline as well as steps provided by plugins.

Declarative Pipeline

Declarative Pipeline is a relatively recent addition to Jenkins Pipeline [19: Version 2.5 of the "Pipeline plugin" introduces support for Declarative Pipeline syntax] which presents a more simplified and opinionated syntax on top of the Pipeline sub-systems.

All valid Declarative Pipelines must be enclosed within a `pipeline` block, for example:

```
pipeline {
    /* insert Declarative Pipeline here */
}
```

The basic statements and expressions which are valid in Declarative Pipeline follow the same rules as [Groovy's syntax](#) with the following exceptions:

- The top-level of the Pipeline must be a *block*, specifically: `pipeline { }`
- No semicolons as statement separators. Each statement has to be on its own line
- Blocks must only consist of [Sections](#), [Directives](#), [Steps](#), or assignment statements.
- A property reference statement is treated as no-argument method invocation. So for example, `input` is treated as `input()`

Sections

Sections in Declarative Pipeline typically contain one or more [Directives](#) or [Steps](#).

agent

The `agent` section specifies where the entire Pipeline, or a specific stage, will execute in the Jenkins environment depending on where the `agent` section is placed. The section must be defined at the top-level inside the `pipeline` block, but stage-level usage is optional.

Required	Yes
Parameters	Described below
Allowed	In the top-level <code>pipeline</code> block and each <code>stage</code> block.

Parameters

In order to support the wide variety of use-cases Pipeline authors may have, the `agent` section supports a few different types of parameters. These parameters can be applied at the top-level of the `pipeline` block, or within each `stage` directive.

any

Execute the Pipeline, or stage, on any available agent. For example: `agent any`

none

When applied at the top-level of the `pipeline` block no global agent will be allocated for the entire Pipeline run and each `stage` section will need to contain its own `agent` section. For example: `agent none`

label

Execute the Pipeline, or stage, on an agent available in the Jenkins environment with the provided label. For example: `agent { label 'my-defined-label' }`

node

`agent { node { label 'labelName' } }` behaves the same as `agent { label 'labelName' }`, but `node` allows for additional options (such as `customWorkspace`).

docker

Execute the Pipeline, or stage, with the given container which will be dynamically provisioned on a `node` pre-configured to accept Docker-based Pipelines, or on a node matching the optionally defined `label` parameter. `docker` also optionally accepts an `args` parameter which may contain arguments to pass directly to a `docker run` invocation, and an `alwaysPull` option, which will force a `docker pull` even if the image name is already present. For example: `agent { docker 'maven:3-alpine' }` or

```
agent {
  docker {
    image 'maven:3-alpine'
    label 'my-defined-label'
    args '-v /tmp:/tmp'
  }
}
```

dockerfile

Execute the Pipeline, or stage, with a container built from a `Dockerfile` contained in the source repository. In order to use this option, the `Jenkinsfile` must be loaded from either a Multibranch Pipeline, or a "Pipeline from SCM." Conventionally this is the `Dockerfile` in the root of the source repository: `agent { dockerfile true }`. If building a `Dockerfile` in another directory, use the `dir` option: `agent { dockerfile { dir 'someSubDir' } }`. You can pass additional arguments to the `docker build ...` command with the `additionalBuildArgs` option, like `agent { dockerfile { additionalBuildArgs '--build-arg foo=bar' } }`.

Common Options

These are a few options that can be applied two or more `agent` implementations. They are not required unless explicitly stated.

label

A string. The label on which to run the Pipeline or individual `stage`.

This option is valid for `node`, `docker` and `dockerfile`, and is required for `node`.

customWorkspace

A string. Run the Pipeline or individual **stage** this **agent** is applied to within this custom workspace, rather than the default. It can be either a relative path, in which case the custom workspace will be under the workspace root on the node, or an absolute path. For example:

```
agent {
  node {
    label 'my-defined-label'
    customWorkspace '/some/other/path'
  }
}
```

This option is valid for **node**, **docker** and **dockerfile**.

reuseNode

A boolean, false by default. If true, run the container on the node specified at the top-level of the Pipeline, in the same workspace, rather than on a new node entirely.

This option is valid for **docker** and **dockerfile**, and only has an effect when used on an **agent** for an individual **stage**.

Example

```
// Declarative //
pipeline {
  agent { docker 'maven:3-alpine' } ①
  stages {
    stage('Example Build') {
      steps {
        sh 'mvn -B clean verify'
      }
    }
  }
}
// Script //
```

- ① Execute all the steps defined in this Pipeline within a newly created container of the given name and tag (**maven:3-alpine**).

Stage-level agent section

```

// Declarative //
pipeline {
  agent none ①
  stages {
    stage('Example Build') {
      agent { docker 'maven:3-alpine' } ②
      steps {
        echo 'Hello, Maven'
        sh 'mvn --version'
      }
    }
    stage('Example Test') {
      agent { docker 'openjdk:8-jre' } ③
      steps {
        echo 'Hello, JDK'
        sh 'java -version'
      }
    }
  }
}
// Script //

```

- ① Defining **agent none** at the top-level of the Pipeline ensures that **an Executor** will not be assigned unnecessarily. Using **agent none** also forces each **stage** section contain its own **agent** section.
- ② Execute the steps in this stage in a newly created container using this image.
- ③ Execute the steps in this stage in a newly created container using a different image from the previous stage.

post

The **post** section defines one or more additional **steps** that are run upon the completion of a Pipeline's or stage's run (depending on the location of the **post** section within the Pipeline). **post** can support one of the following **post-condition** blocks: **always**, **changed**, **failure**, **success**, **unstable**, and **aborted**. These condition blocks allow the execution of steps within the **post** section depending on the completion status of the Pipeline or stage.

Required	No
Parameters	<i>None</i>
Allowed	In the top-level pipeline block and each stage block.

Conditions

always

Run the steps in the **post** section regardless of the completion status of the Pipeline's or stage's run.

changed

Only run the steps in **post** if the current Pipeline's or stage's run has a different completion status from its previous run.

failure

Only run the steps in **post** if the current Pipeline's or stage's run has a "failed" status, typically denoted by red in the web UI.

success

Only run the steps in **post** if the current Pipeline's or stage's run has a "success" status, typically denoted by blue or green in the web UI.

unstable

Only run the steps in **post** if the current Pipeline's or stage's run has an "unstable" status, usually caused by test failures, code violations, etc. This is typically denoted by yellow in the web UI.

aborted

Only run the steps in **post** if the current Pipeline's or stage's run has an "aborted" status, usually due to the Pipeline being manually aborted. This is typically denoted by gray in the web UI.

Example

```
// Declarative //
pipeline {
  agent any
  stages {
    stage('Example') {
      steps {
        echo 'Hello World'
      }
    }
  }
  post { ①
    always { ②
      echo 'I will always say Hello again!'
    }
  }
}
// Script //
```

① Conventionally, the **post** section should be placed at the end of the Pipeline.

② **Post-condition** blocks contain **steps** the same as the **[steps]** section.

stages

Containing a sequence of one or more **[stage]** directives, the **stages** section is where the bulk of the "work" described by a Pipeline will be located. At a minimum it is recommended that **stages** contain at least one **[stage]** directive for each discrete part of the continuous delivery process, such as Build, Test, and Deploy.

Required	Yes
Parameters	<i>None</i>
Allowed	Only once, inside the pipeline block.

Example

```
// Declarative //
pipeline {
  agent any
  stages { ①
    stage('Example') {
      steps {
        echo 'Hello World'
      }
    }
  }
}
// Script //
```

① The **stages** section will typically follow the directives such as **agent**, **options**, etc.

steps

The **steps** section defines a series of one or more **steps** to be executed in a given **stage** directive.

Required	Yes
Parameters	<i>None</i>
Allowed	Inside each stage block.

Example

```
// Declarative //
pipeline {
  agent any
  stages {
    stage('Example') {
      steps { ①
        echo 'Hello World'
      }
    }
  }
}
// Script //
```

① The `steps` section must contain one or more steps.

Directives

environment

The `environment` directive specifies a sequence of key-value pairs which will be defined as environment variables for the all steps, or stage-specific steps, depending on where the `environment` directive is located within the Pipeline.

This directive supports a special helper method `credentials()` which can be used to access pre-defined Credentials by their identifier in the Jenkins environment. For Credentials which are of type "Secret Text", the `credentials()` method will ensure that the environment variable specified contains the Secret Text contents. For Credentials which are of type "Standard username and password", the environment variable specified will be set to `username:password` and two additional environment variables will be automatically be defined: `MYVARNAME_USR` and `MYVARNAME_PSW` respective.

Required	No
Parameters	<i>None</i>
Allowed	Inside the <code>pipeline</code> block, or within <code>stage</code> directives.

Example

```

// Declarative //
pipeline {
  agent any
  environment { ①
    CC = 'clang'
  }
  stages {
    stage('Example') {
      environment { ②
        AN_ACCESS_KEY = credentials('my-prefined-secret-text') ③
      }
      steps {
        sh 'printenv'
      }
    }
  }
}
// Script //

```

- ① An `environment` directive used in the top-level `pipeline` block will apply to all steps within the Pipeline.
- ② An `environment` directive defined within a `stage` will only apply the given environment variables to steps within the `stage`.
- ③ The `environment` block has a helper method `credentials()` defined which can be used to access pre-defined Credentials by their identifier in the Jenkins environment.

options

The `options` directive allows configuring Pipeline-specific options from within the Pipeline itself. Pipeline provides a number of these options, such as `buildDiscarder`, but they may also be provided by plugins, such as `timestamps`.

Required	No
Parameters	<i>None</i>
Allowed	Only once, inside the <code>pipeline</code> block.

Available Options

`buildDiscarder`

Persist artifacts and console output for the specific number of recent Pipeline runs. For example:
`options { buildDiscarder(logRotator(numToKeepStr: '1')) }`

`disableConcurrentBuilds`

Disallow concurrent executions of the Pipeline. Can be useful for preventing simultaneous

accesses to shared resources, etc. For example: `options { disableConcurrentBuilds() }`

overrideIndexTriggers

Allows overriding default treatment of branch indexing triggers. If branch indexing triggers are disabled at the multibranch or organization label, `options { overrideIndexTriggers(true) }` will enable them for this job only. Otherwise, `options { overrideIndexTriggers(false) }` will disable branch indexing triggers for this job only.

skipDefaultCheckout

Skip checking out code from source control by default in the `agent` directive. For example: `options { skipDefaultCheckout() }`

skipStagesAfterUnstable

Skip stages once the build status has gone to UNSTABLE. For example: `options { skipStagesAfterUnstable() }`

timeout

Set a timeout period for the Pipeline run, after which Jenkins should abort the Pipeline. For example: `options { timeout(time: 1, unit: 'HOURS') }`

retry

On failure, retry the entire Pipeline the specified number of times. For example: `options { retry(3) }`

timestamps

Prepend all console output generated by the Pipeline run with the time at which the line was emitted. For example: `options { timestamps() }`

Example

```
// Declarative //
pipeline {
  agent any
  options {
    timeout(time: 1, unit: 'HOURS') ①
  }
  stages {
    stage('Example') {
      steps {
        echo 'Hello World'
      }
    }
  }
}
// Script //
```

① Specifying a global execution timeout of one hour, after which Jenkins will abort the Pipeline run.

NOTE | A comprehensive list of available options is pending the completion of [INFRA-1503](#).

parameters

The `parameters` directive provides a list of parameters which a user should provide when triggering the Pipeline. The values for these user-specified parameters are made available to Pipeline steps via the `params` object, see the [Example](#) for its specific usage.

Required	No
Parameters	<i>None</i>
Allowed	Only once, inside the <code>pipeline</code> block.

Available Parameters

string

A parameter of a string type, for example: `parameters { string(name: 'DEPLOY_ENV', defaultValue: 'staging', description: '') }`

booleanParam

A boolean parameter, for example: `parameters { booleanParam(name: 'DEBUG_BUILD', defaultValue: true, description: '') }`

Example

```
// Declarative //
pipeline {
  agent any
  parameters {
    string(name: 'PERSON', defaultValue: 'Mr Jenkins', description: 'Who should I
say hello to?')
  }
  stages {
    stage('Example') {
      steps {
        echo "Hello ${params.PERSON}"
      }
    }
  }
}
// Script //
```

NOTE | A comprehensive list of available parameters is pending the completion of [INFRA-1503](#).

triggers

The `triggers` directive defines the automated ways in which the Pipeline should be re-triggered. For Pipelines which are integrated with a source such as GitHub or BitBucket, `triggers` may not be necessary as webhooks-based integration will likely already be present. The triggers currently available are `cron`, `pollSCM` and `upstream`.

Required	No
Parameters	<i>None</i>
Allowed	Only once, inside the <code>pipeline</code> block.

cron

Accepts a cron-style string to define a regular interval at which the Pipeline should be re-triggered, for example: `triggers { cron('H */4 * * 1-5') }`

pollSCM

Accepts a cron-style string to define a regular interval at which Jenkins should check for new source changes. If new changes exist, the Pipeline will be re-triggered. For example: `triggers { pollSCM('H */4 * * 1-5') }`

upstream

Accepts a comma separated string of jobs and a threshold. When any job in the string finishes with the minimum threshold, the Pipeline will be re-triggered. For example: `triggers { upstream(upstreamProjects: 'job1,job2', threshold: hudson.model.Result.SUCCESS) }`

NOTE The `pollSCM` trigger is only available in Jenkins 2.22 or later.

Example

```
// Declarative //
pipeline {
  agent any
  triggers {
    cron('H */4 * * 1-5')
  }
  stages {
    stage('Example') {
      steps {
        echo 'Hello World'
      }
    }
  }
}
// Script //
```

stage

The `stage` directive goes in the `stages` section and should contain a `[steps]` section, an optional `agent` section, or other stage-specific directives. Practically speaking, all of the real work done by a Pipeline will be wrapped in one or more `stage` directives.

Required	At least one
Parameters	One mandatory parameter, a string for the name of the stage.
Allowed	Inside the <code>stages</code> section.

Example

```
// Declarative //
pipeline {
  agent any
  stages {
    stage('Example') {
      steps {
        echo 'Hello World'
      }
    }
  }
}
// Script //
```

tools

A section defining tools to auto-install and put on the `PATH`. This is ignored if `agent none` is specified.

Required	No
Parameters	<i>None</i>
Allowed	Inside the <code>pipeline</code> block or a <code>stage</code> block.

Supported Tools

maven

jdk

gradle

Example

```
// Declarative //
pipeline {
  agent any
  tools {
    maven 'apache-maven-3.0.1' ①
  }
  stages {
    stage('Example') {
      steps {
        sh 'mvn --version'
      }
    }
  }
}
// Script //
```

① The tool name must be pre-configured in Jenkins under **Manage Jenkins** → **Global Tool Configuration**.

when

The **when** directive allows the Pipeline to determine whether the stage should be executed depending on the given condition. The **when** directive must contain at least one condition. If the **when** directive contains more than one condition, all the child conditions must return true for the stage to execute. This is the same as if the child conditions were nested in an **allOf** condition (see the [examples](#) below).

More complex conditional structures can be built using the nesting conditions: **not**, **allOf**, or **anyOf**. Nesting conditions may be nested to any arbitrary depth.

Required	No
Parameters	<i>None</i>
Allowed	Inside a stage directive

Built-in Conditions

branch

Execute the stage when the branch being built matches the branch pattern given, for example: **when { branch 'master' }**. Note that this only works on a multibranch Pipeline.

environment

Execute the stage when the specified environment variable is set to the given value, for example: **when { environment name: 'DEPLOY_TO', value: 'production' }**

expression

Execute the stage when the specified Groovy expression evaluates to true, for example: `when { expression { return params.DEBUG_BUILD } }`

not

Execute the stage when the nested condition is false. Must contain one condition. For example: `when { not { branch 'master' } }`

allOf

Execute the stage when all of the nested conditions are true. Must contain at least one condition. For example: `when { allOf { branch 'master'; environment name: 'DEPLOY_TO', value: 'production' } }`

anyOf

Execute the stage when at least one of the nested conditions is true. Must contain at least one condition. For example: `when { anyOf { branch 'master'; branch 'staging' } }`

Examples

Single condition

```
// Declarative //
pipeline {
  agent any
  stages {
    stage('Example Build') {
      steps {
        echo 'Hello World'
      }
    }
    stage('Example Deploy') {
      when {
        branch 'production'
      }
      steps {
        echo 'Deploying'
      }
    }
  }
}
// Script //
```

Multiple condition

```
// Declarative //
pipeline {
  agent any
  stages {
    stage('Example Build') {
      steps {
        echo 'Hello World'
      }
    }
    stage('Example Deploy') {
      when {
        branch 'production'
        environment name: 'DEPLOY_TO', value: 'production'
      }
      steps {
        echo 'Deploying'
      }
    }
  }
}
// Script //
```

Nested condition (same behavior as previous example)

```
// Declarative //
pipeline {
  agent any
  stages {
    stage('Example Build') {
      steps {
        echo 'Hello World'
      }
    }
    stage('Example Deploy') {
      when {
        allOf {
          branch 'production'
          environment name: 'DEPLOY_TO', value: 'production'
        }
      }
      steps {
        echo 'Deploying'
      }
    }
  }
}
// Script //
```

Multiple condition and nested condition

```
// Declarative //
pipeline {
  agent any
  stages {
    stage('Example Build') {
      steps {
        echo 'Hello World'
      }
    }
    stage('Example Deploy') {
      when {
        branch 'production'
        anyOf {
          environment name: 'DEPLOY_TO', value: 'production'
          environment name: 'DEPLOY_TO', value: 'staging'
        }
      }
      steps {
        echo 'Deploying'
      }
    }
  }
}
// Script //
```



```
// Declarative //
pipeline {
  agent any
  stages {
    stage('Example Build') {
      steps {
        echo 'Hello World'
      }
    }
    stage('Example Deploy') {
      when {
        expression { BRANCH_NAME ==~ /(production|staging)/ }
        anyOf {
          environment name: 'DEPLOY_TO', value: 'production'
          environment name: 'DEPLOY_TO', value: 'staging'
        }
      }
      steps {
        echo 'Deploying'
      }
    }
  }
}
// Script //
```

Parallel

Stages in Declarative Pipeline may declare a number of nested stages within them, which will be executed in parallel. Note that a stage must have one and only one of either `steps` or `parallel`. The nested stages cannot contain further `parallel` stages themselves, but otherwise behave the same as any other `stage`. Any stage containing `parallel` cannot contain `agent` or `tools`, since those are not relevant without `steps`.

In addition, you can force your `parallel` stage's to all be aborted when one of them fails, by adding `'failFast true'` to the `stage` containing the `parallel`.

Example

```

// Declarative //
pipeline {
  agent any
  stages {
    stage('Non-Parallel Stage') {
      steps {
        echo 'This stage will be executed first.'
      }
    }
    stage('Parallel Stage') {
      when {
        branch 'master'
      }
      failFast true
      parallel {
        stage('Branch A') {
          agent {
            label "for-branch-a"
          }
          steps {
            echo "On Branch A"
          }
        }
        stage('Branch B') {
          agent {
            label "for-branch-b"
          }
          steps {
            echo "On Branch B"
          }
        }
      }
    }
  }
}

// Script //

```

Steps

Declarative Pipelines may use all the available steps documented in the [Pipeline Steps reference](#), which contains a comprehensive list of steps, with the addition of the steps listed below which are **only supported** in Declarative Pipeline.

script

The `script` step takes a block of `[scripted-pipeline]` and executes that in the Declarative Pipeline. For most use-cases, the `script` step should be unnecessary in Declarative Pipelines, but it can provide a

useful "escape hatch." `script` blocks of non-trivial size and/or complexity should be moved into [Shared Libraries](#) instead.

Example

```
// Declarative //
pipeline {
  agent any
  stages {
    stage('Example') {
      steps {
        echo 'Hello World'

        script {
          def browsers = ['chrome', 'firefox']
          for (int i = 0; i < browsers.size(); ++i) {
            echo "Testing the ${browsers[i]} browser"
          }
        }
      }
    }
  }
}
// Script //
```

Scripted Pipeline

Scripted Pipeline, like [\[declarative-pipeline\]](#), is built on top of the underlying Pipeline sub-system. Unlike Declarative, Scripted Pipeline is effectively a general purpose DSL [20: Domain-specific Language] built with [Groovy](#). Most functionality provided by the Groovy language is made available to users of Scripted Pipeline, which means it can be a very expressive and flexible tool with which one can author continuous delivery pipelines.

Flow Control

Scripted Pipeline is serially executed from the top of a [Jenkinsfile](#) downwards, like most traditional scripts in Groovy or other languages. Providing flow control therefore rests on Groovy expressions, such as the [if/else](#) conditionals, for example:

```
// Scripted //
node {
    stage('Example') {
        if (env.BRANCH_NAME == 'master') {
            echo 'I only execute on the master branch'
        } else {
            echo 'I execute elsewhere'
        }
    }
}
// Declarative //
```

Another way Scripted Pipeline flow control can be managed is with Groovy's exception handling support. When [Steps](#) fail for whatever reason they throw an exception. Handling behaviors on-error must make use of the [try/catch/finally](#) blocks in Groovy, for example:

```
// Scripted //
node {
    stage('Example') {
        try {
            sh 'exit 1'
        }
        catch (exc) {
            echo 'Something failed, I should sound the klaxons!'
            throw
        }
    }
}
// Declarative //
```

Steps

As discussed in [Getting Started](#), the most fundamental part of a Pipeline is the "step." Fundamentally, steps tell Jenkins *what* to do, and serve as the basic building block for both Declarative and Scripted Pipeline syntax.

Scripted Pipeline does **not** introduce any steps which are specific to its syntax; [Pipeline Steps reference](#) which contains a comprehensive list of steps provided by Pipeline and plugins.

Differences from plain Groovy

In order to provide *durability*, which means that running Pipelines can survive a restart of the Jenkins [master](#), Scripted Pipeline must serialize data back to the master. Due to this design requirement, some Groovy idioms such as `collection.each { item → /* perform operation */ }` are not fully supported. See [JENKINS-27421](#) and [JENKINS-26481](#) for more information.

Syntax Comparison

When Jenkins Pipeline was first created, Groovy was selected as the foundation. Jenkins has long shipped with an embedded Groovy engine to provide advanced scripting capabilities for admins and users alike. Additionally, the implementors of Jenkins Pipeline found Groovy to be a solid foundation upon which to build what is now referred to as the "Scripted Pipeline" DSL. [1: [Domain-Specific Language](#)].

As it is a fully featured programming environment, Scripted Pipeline offers a tremendous amount of flexibility and extensibility to Jenkins users. The Groovy learning-curve isn't typically desirable for all members of a given team, so Declarative Pipeline was created to offer a simpler and more opinionated syntax for authoring Jenkins Pipeline.

The two are both fundamentally the same Pipeline sub-system underneath. They are both durable implementations of "Pipeline as code." They are both able to use steps built into Pipeline or provided by plugins. Both are able utilize [Shared Libraries](#)

Where they differ however is in syntax and flexibility. Declarative limits what is available to the user with a more strict and pre-defined structure, making it an ideal choice for simpler continuous delivery pipelines. Scripted provides very few limits, insofar that the only limits on structure and syntax tend to be defined by Groovy itself, rather than any Pipeline-specific systems, making it an ideal choice for power-users and those with more complex requirements. As the name implies, Declarative Pipeline is encourages a declarative programming model. [21: [Declarative Programming](#)] Whereas Scripted Pipelines follow a more imperative programming model.. [22: [Imperative Programming](#)]

Blue Ocean

This chapter will cover all aspects of Blue Ocean, from the Dashboard, to viewing branches and results from individual Pipeline runs, to using the Visual Editor to modify Pipelines as code.

This chapter is intended for Jenkins users of all skill levels, but beginners may need to refer to some sections of "[Using Jenkins](#)" to understand some topics covered in this chapter.

If you are not yet familiar with Jenkins terminology and features, start with [Getting Started with Jenkins](#).

What is Blue Ocean?

Blue Ocean rethinks the user experience of Jenkins. Designed from the ground up for [Jenkins Pipeline](#), but still compatible with Freestyle jobs, Blue Ocean reduces clutter and increases clarity for every member of the team.

- **Sophisticated visualizations** of continuous delivery (CD) Pipelines, allowing for fast and intuitive comprehension of pipeline's status.
- **Pipeline editor** makes creation of Pipelines approachable by guiding the user through an intuitive and visual process to create a Pipeline.
- **Personalization** to suit the role-based needs of each member of the team.
- **Pinpoint precision** when intervention is needed and/or issues arise. Blue Ocean shows where in the pipeline attention is needed, facilitating exception handling and increasing productivity.
- **Native integration for branch and pull requests** enables maximum developer productivity when collaborating on code with others in GitHub and Bitbucket.

To start using Blue Ocean, see [Getting Started with Blue Ocean](#).

Join the community

There are a few ways you can join the community:

1. Chat with the community and development team on Gitter [chat on gitter](#)
2. Request features or report bugs against the [blueocean-plugin component in JIRA](#).
3. Subscribe and ask questions on the [Jenkins Users mailing list](#).
4. Developer? We've [labeled a few issues](#) that are great for anyone wanting to get started developing Blue Ocean. Don't forget to drop by the Gitter chat and introduce yourself!

Frequently Asked Questions

Why does Blue Ocean exist?

The world has moved on from developer tools that are purely functional to developer tools being part of a "developer experience". That is to say, it's no longer about a single tool but the many tools developers use throughout the day and how they work together to achieve a workflow that's beneficial for the developer - this is Developer Experience.

Developer tools companies like Heroku, Atlassian and Github have raised the bar for what is considered good developer experience, and developers are increasingly expecting exceptional design. In recent years developers are becoming more rapidly attracted to tools that are not only functional but are designed to fit into their workflow seamlessly and are a joy to use. This shift represents a higher standard of design and user experience that Jenkins needs to rise to meet.

Creating and visualising continuous delivery pipelines is something valuable for many Jenkins users and this is demonstrated in the 5+ plugins that the community has created to meet their needs. To us this indicates a need to revisit how Jenkins currently expresses these concepts and consider delivery pipelines as a central theme to the Jenkins user experience.

It's not just continuous delivery concepts but the tools that developers use every day – Github, Bitbucket, Slack, HipChat, Puppet or Docker. It's about more than Jenkins – it's the developer workflow that surrounds Jenkins that spans multiple tools.

New teams have little time for learning to assemble their own Jenkins experience – they want to improve their time to market by shipping better software faster. Assembling that ideal Jenkins experience is something we can work together as a community of Jenkins users and contributors to define. As time progresses, developers' expectations of good user experience will change and the mission of Blue Ocean will enable the Jenkins project to respond.

The Jenkins community has poured its sweat and tears into building the most technically capable and extensible software automation tool in existence. Not doing anything to revolutionize the Jenkins developer experience today is just inviting someone else – in closed source – to do it.

Where is the name from?

The name Blue Ocean comes from the book [Blue Ocean Strategy](#) where instead of looking at strategic problems within a contested space you look at problems in the larger uncontested space. To put this more simply, consider this quote from ice hockey legend Wayne Gretzky: "skate to where the puck is going to be, not where it has been".

Does Blue Ocean support Freestyle jobs?

Blue Ocean aims to deliver a great experience around Pipeline and be compatible with any Freestyle jobs that you have configured in your system. However, they won't be able to benefit from any of the features built for Pipelines – for example, Pipeline visualization.

As Blue Ocean is designed to be extensible it will be possible for the community to extend it for

other job types in the future.

What does this mean for the classic Jenkins UI?

The intention is that as Blue Ocean matures there will be less and less reasons for users to go back to the existing UI.

For example, in the first version we will mainly be targeting Pipeline jobs. You might be able to see your existing non-pipeline jobs in Blue Ocean but it might not be possible to configure them from the new UI for some time. This means users will have to jump back to the classic UI for configuration of non-pipeline jobs.

There are likely going to be more examples of this and that's why the classic UI will still be important in the long term.

What does this mean for my plugins?

Extensibility is a pretty core concept to Jenkins, so being able to extend the Blue Ocean UI is important. Based on some research, we worked out a way to allow `<ExtensionPoint name=..>` to be used in the markup of Blue Ocean, leaving places for plugins to contribute to the UI (plugins can have their own Blue Ocean extension points, just like they do today in Jenkins). Blue Ocean itself (as it is so far) is implemented using these extension points. Extensions are delivered by plugins, as normal, only if they wish to contribute to the Blue Ocean experience they will have some additional javascript that provides extensions.

What technologies are currently in use?

Blue Ocean is built as a collection of Jenkins plugins itself. There is one key difference, however. It provides both its own endpoint for http requests and delivers up html/javascript via a different path, without the existing Jenkins UI markup/scripts. React.js and ES6 are used to deliver the javascript components of Blue Ocean. Inspired by this excellent open source project ([react-plugins](#)) an `<ExtensionPoint>` pattern was established, that allows extensions to come from any Jenkins plugin (only with Javascript) and should they fail to load, have failures isolated.

Where can I find the source code?

The source code can be found on Github:

- [Blue Ocean](#)
 - [Jenkins Design Language](#)
-

Getting Started with Blue Ocean

This section will show how to start using Blue Ocean. It will include instructions for installing and configuring the Blue Ocean plugin, and how to get switch into and out of the Blue Ocean UI.

Installing

Blue Ocean can be installed in an existing Jenkins environment or be run [with Docker](#).

In an Existing Jenkins Environment

To start using the plugin:blueocean[Blue Ocean plugin] in an existing Jenkins environment, it must be running Jenkins 2.7.x or later.:

1. Login to your Jenkins server
2. Click **Manage Jenkins** in the sidebar then **Manage Plugins**
3. Choose the **Available** tab and use the search bar to find **Blue Ocean**
4. Click the checkbox in the Install column
5. Click either **Install without restart** or **Download now and install after restart**

Filter:

Updates **Available** Installed Advanced

Install	Name	Version
<input type="checkbox"/>	Blue Ocean Blue Ocean is a new project that rethinks the user experience of Jenkins. Designed from the ground up for Jenkins Pipeline and compatible with Freestyle jobs, Blue Ocean reduces clutter and increases clarity for every member of your team.	1.0.0
<input type="checkbox"/>	Common API for Blue Ocean	1.0.0
<input type="checkbox"/>	Config API for Blue Ocean	1.0.0
<input type="checkbox"/>	Dashboard for Blue Ocean	1.0.0
<input type="checkbox"/>	Events API for Blue Ocean	1.0.0
<input type="checkbox"/>	Git Pipeline for Blue Ocean	1.0.0
<input type="checkbox"/>	GitHub Pipeline for Blue Ocean	1.0.0
<input type="checkbox"/>	i18n for Blue Ocean	1.0.0
<input type="checkbox"/>	JWT for Blue Ocean	1.0.0
<input type="checkbox"/>	Personalization for Blue Ocean	1.0.0
<input type="checkbox"/>	Pipeline REST API for Blue Ocean	1.0.0
<input type="checkbox"/>	REST API for Blue Ocean	1.0.0
<input type="checkbox"/>	REST Implementation for Blue Ocean	1.0.0
<input type="checkbox"/>	Web for Blue Ocean	1.0.0
<input type="checkbox"/>	Blue Ocean Pipeline Editor The Blue Ocean Pipeline Editor is the simplest way for anyone wanting to get started with creating Pipelines in Jenkins	0.2.0

Install without restart

Download now and install after restart

Check now

For in-depth description on how to install and manage plugins, refer to the [Managing Plugins](#) section.





The majority of Blue Ocean requires no additional configuration after installation. Existing Pipelines and Jobs will continue to work as usual. However, the first time a [Pipeline is created or added](#), Blue Ocean will ask for permissions to access your repositories (either Git or GitHub) in order to create pipelines based on those repositories.

With Docker

Read more about this in the [Docker](#) section of the [Installing Jenkins](#) page.

Starting Blue Ocean

Once a Jenkins environment has Blue Ocean installed, users can start using Blue Ocean by clicking the **Open Blue Ocean** in the navigation bar of the Jenkins web UI. Alternatively, users can navigate directly to Blue Ocean at the `/blue` URL for their Jenkins environment, for example `JENKINS_URL/blue`.

-  Manage Jenkins
-  My Views
-  Open Blue Ocean
-  Credentials

Build Queue
No builds in the queue

If Pipelines are already present on the current Jenkins instance, this will bring up the [Blue Ocean Dashboard](#).

If this is a new Jenkins instance, Blue Ocean will display a box offering to "[Create a new pipeline](#)."



Welcome to Jenkins

It's time to create your first Pipeline.

Create a new Pipeline

Navigation bar


Blue Ocean uses a common Navigation bar at the top of most Blue Ocean views. It includes five buttons:

- **Jenkins** - Navigate to the Dashboard (reload if already viewing it)
- **Pipelines** - Navigate to the Dashboard (do nothing if already viewing it)
- **Administration** - [Manage](#) this Jenkins instance (using the Classic UI)
- **Switch to "Classic" UI** - [Switch to the "Classic" Jenkins UI](#)
- **Logout** - Logout the current user, return to the Jenkins login page

Views that use the standard navigation bar will add another bar below it with options specific to that view. Some views replace the common navigation bar with one specifically suited to that view.

Switching to the "Classic" UI

Blue Ocean may not support some legacy or administrative functions which are necessary to some users. For those wishing to exit the Blue Ocean user experience, an "exit" icon is located at the top of most pages in Blue Ocean. Clicking the exit icon will navigate to the most relevant page in "classic" which parallels the current page in Blue Ocean.




Some links in Blue Ocean, like **Administration**, will also navigate to the classic web UI when there is no Blue Ocean equivalent. In these cases, Blue Ocean will automatically take the user into the classic web UI as necessary.

Creating Pipelines

Blue Ocean makes it easy to create Pipelines in Jenkins. Pipelines can be created from existing `Jenkinsfile`s or from new `Jenkinsfile`s created with the [Blue Ocean Pipeline Editor](#). The Pipeline Creation workflow guides users through this process in clear, easy to understand steps.

Starting Pipeline Creation

At the top of the [Blue Ocean Dashboard](#), is a "New Pipeline" button that launch the Pipeline Creation workflow.



On new a Jenkins instance, where there are no jobs or Pipelines and the Dashboard is empty, Blue Ocean will also display a message box offering to "Create a New Pipeline".



Welcome to Jenkins

It's time to create your first Pipeline.

Create a new Pipeline

Creating a Pipeline for a Git Repository

To create a Pipeline from a Git repository, start by selecting **Git** as the Source Control system.




Then enter the URL for the Git Repository and optionally select which credentials to use. If the dropdown does not show the desired credentials, they can be added using the "Add" button.

When done, click "Create Pipeline". Blue ocean will look at all branches for the selected repository, and will start a Pipeline run for each branch containing a **Jenkinsfile**.



Creating Pipelines for GitHub Repositories

To create a Pipeline from a GitHub, start by selecting **GitHub** as the Source Control system.




Provide a GitHub Access Token

If this is the first time Pipeline Creation has been run by the currently logged in user, Blue Ocean will ask for a [GitHub Access Token](#) to allow Blue Ocean to access your organizations and repositories.



If you have not already created a access token, click on the link provided and Blue Ocean will navigate to [the right page on GitHub](#), automatically selecting the appropriate permissions it will need.



New personal access token

Personal access tokens function like ordinary OAuth access tokens. They can be used instead of a for Git over HTTPS, or can be used to [authenticate to the API over Basic Authentication](#).

Token description

Blue Ocean Token

What's this token for?

Select scopes

Scopes define the access for personal tokens. [Read more about OAuth scopes](#).

- | | |
|--|--------------------------------------|
| <input checked="" type="checkbox"/> repo | Full control of private repositories |
| <input checked="" type="checkbox"/> repo:status | Access commit status |
| <input checked="" type="checkbox"/> repo_deployment | Access deployment status |
| <input checked="" type="checkbox"/> public_repo | Access public repositories |

Select a GitHub Account or Organization

All repositories on Github are grouped by owner, either an account or organization. When creating Pipelines, Blue Ocean mirrors that structure, asking users to select an account or organization which owns the repositories from which it will add Pipelines.



From here, Blue Ocean offers two styles of Pipeline creation, either "[single Pipeline](#)" or "[discover all](#)

Pipelines".

New Pipeline from a Single Repository

Selecting "New Pipeline" allows the user select and create a Pipeline for a single Repository.



After selecting a repository Blue Ocean will scan all the branches in that Repository and will create a Pipeline for each branch containing a "Jenkinsfile" in the root folder. Blue Ocean will then run the Pipeline created for each branch in this process.

If no branches in the selected repository have a "Jenkinsfile", Blue Ocean will offer to "Create a New Pipeline" for that repository, taking the user to the [Blue Ocean Pipeline Editor](#) to create a new **Jenkinsfile** and add a new Pipeline based on that.

Auto-discover Pipelines

Selecting "Auto-discover Pipelines" scans all repositories belonging to the selected owner, and will create a Pipeline for each branch containing a "Jenkinsfile" in the root folder.



This option is useful for adding Pipelines for all the repositories in an organization, when those repositories already have **Jenkinsfile** entries in them. Repositories that do not contain **Jenkinsfile** entries are ignored. To create a new **Jenkinsfile** in a single repository that does not have one, use the "New Pipeline" option instead.

Dashboard

The Blue Ocean Dashboard shows an overview of all Pipelines on a Jenkins instance. The Dashboard is the default view shown when you open Blue Ocean. It consists of a Navigation bar at the top, a Favorites list, and a Pipelines list.

Jenkins Pipelines Administration [Logout](#)

Pipelines [New Pipeline](#)

Favorites

- ✕ bitwise-jenkins / git-plugin master a few seconds ago ↺ ⏪ ⭐
- bitwise-jenkins / junit-plugin blog/blue-ocean-editor #d3d9a39 12 days ago ↺ ⏪ ⭐
- ✓ bitwise-jenkins / junit-plugin PR-7 #b518058 12 days ago ↺ ⏪ ⭐

Name	Health	Branches	PR
bitwise-jenkins / git-plugin		-	
bitwise-jenkins / hermann		-	
bitwise-jenkins / JS-Nightwatch.js		-	
bitwise-jenkins / junit-plugin		1 passing	2 passing
bitwise-jenkins / logstash-plugin		-	

Navigation Bar

The Dashboard includes the [standard navigation bar](#) at the top, with a local navigation bar below that. The local navigation bar includes only one button:

- **New Pipeline** - Open the [Pipeline Creation workflow](#)

Favorites

The "Favorites" list is displayed above the regular Pipelines list. The "Favorites" list shows the [status](#) and other key details about the latest Run for branches or pull requests that include items that are important to the current user. Blue Ocean automatically adds branches or pull requests to this list when they contain a Run that has changes authored by the current user. Users can also manually remove items from their favorites by clicking on the solid star " " in this list.

Clicking on an item in the Favorites list will bring up the [Pipeline Run Details](#) for latest Run in that branch or pull request.

Pipelines

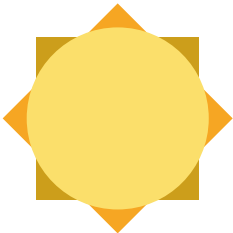
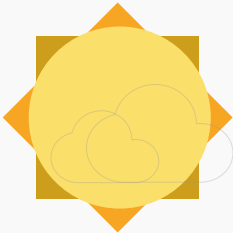


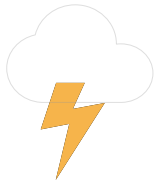
The "Pipelines" list shows the overall state of each Pipeline in this Jenkins instance, including the [health](#) of the pipeline, counts of branches and pull requests that are passing or failing, and a star (either solid " " or outline " ") indicating whether the default branch for this Pipeline has been manually added to this user's Favorites. Clicking on the star will toggle whether the default branch for that Pipeline is shown in the ["Favorites" list](#) for this user.

Clicking on an item in the Pipelines list will bring up the [Pipeline Activity View](#) for that Pipeline.

Health Icons

Blue Ocean represents the overall health of a Pipeline or Pipeline branch using weather icons, which change depending on the number of recent builds that have Passed. Health icons on the [Dashboard](#) represent overall Pipeline health. Health icons in the [Branches tab of the Activity View](#) represent health for each branch.

Table 1. Pipeline Health Icons (best to worst)

Icon	Health
	Sunny , more than 80% of Runs passing
	Partially Sunny , 61% to 80% of Runs passing
	Cloudy , 41% to 60% of Runs passing
	Raining , 21% to 40% of Runs passing
	Storm , less than 21% of Runs passing

Pipeline Run Status

Blue Ocean represents Run Status using a consistent set of icons throughout.

Table 2. Pipeline Run Status Icons

Icon	Status
	In Progress
	Passed
	Unstable
	Failed
	Aborted

Activity View

The Blue Ocean Activity View shows the all activity related to one Pipeline.

Jenkins Pipelines Administration Logout

bitwise-jenkins / junit-plugin ☆ Activity Branches Pull Requests

Status	Run	Commit	Branch	Message	Duration	Completed	
✓	3	b518058	PR-7	-	4m 51s	a minute ago	↺
✓	1	63a7f47	master	-	5s	8 minutes ago	↺
✗	1	86b2229	PR-7	-	48s	7 minutes ago	↺
—	6	d3d9a39	blog/blue-ocean-editor	-	1m 52s	12 minutes ago	↺
✓	5	d3d9a39	blog/blue-ocean-editor	-	11m 8s	4 hours ago	↺

Navigation Bar

The Activity View includes the [standard navigation bar](#) at the top, with a local navigation bar below that. The local navigation bar includes:

- **Pipeline Name** - Clicking on this displays the [default activity tab](#)
- **Favorites Toggle** - Clicking the "Favorite" symbol (a star outline " ") adds a branch to the favorites list shown on the [Dashboard's "Favorites" list](#) for this user.
- **Tabs** ([Activity](#), [Branches](#), [Pull Requests](#)) - Clicking one of these will display that tab of the Activity View.

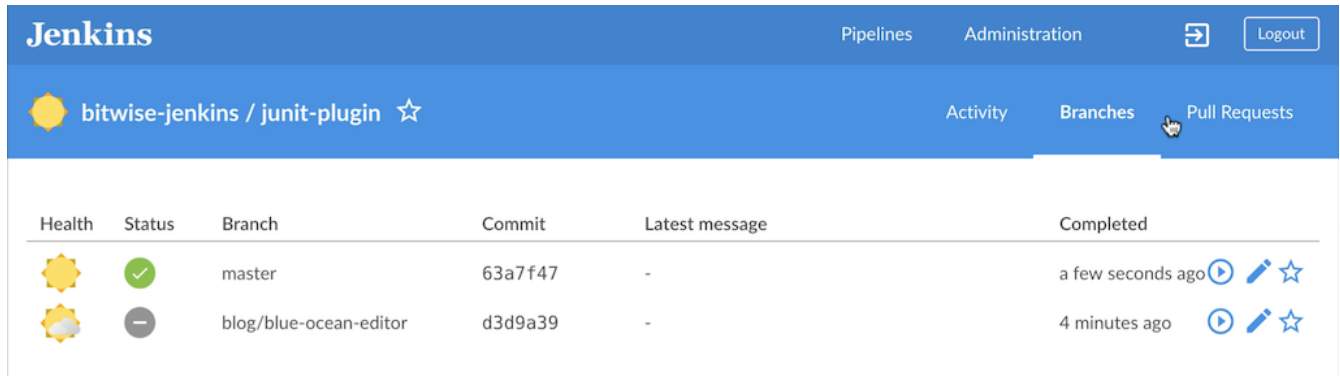
Activity

The default tab of the Activity View, the "Activity" tab, shows a list of the latest completed or in-progress Runs. Each line in the list shows the [status](#) of the Run, id number, commit information, duration, and when the run completed. Clicking on a Run will bring up the [Pipeline Run Details](#) for that Run. "In Progress" Runs can be aborted from this list by clicking on the "Stop" symbol (a square " " inside a circle). Runs that have completed can be re-run by clicking the "Re-run" symbol (a counter-clockwise arrow " "). The list can be filtered by branch or pull request by clicking on the "branch" drop-down in the list header.

This list does not allow runs to be edited or marked as favorites. Those actions can be done from the ["branches"](#) tab.

Branches

The "Branches" tab shows a list of all branches that have a completed or in-progress Run in the current Pipeline. Each line in the list corresponds to a branch in source control, [23: en.wikipedia.org/wiki/Source_control_management] showing overall health of the branch based on recent runs, status of the most recent run, id number, commit information, duration, and when the run completed.

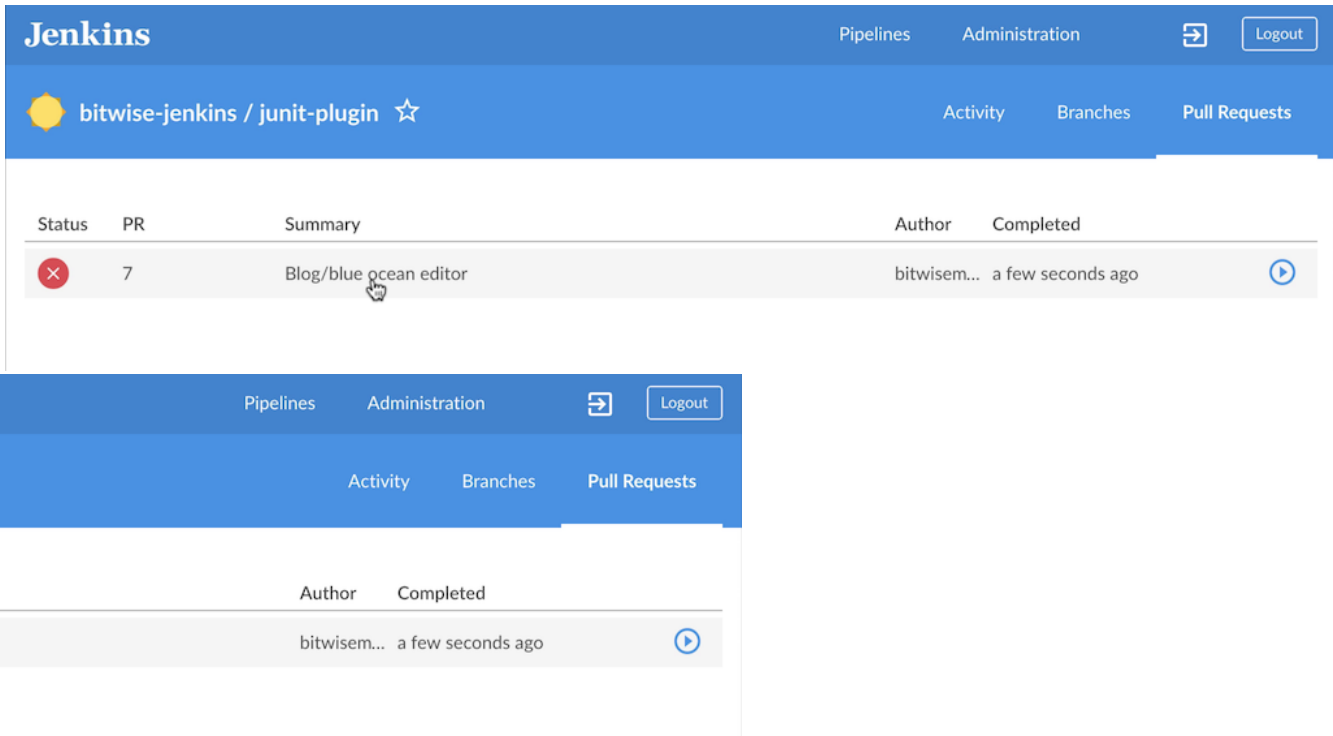


Health	Status	Branch	Commit	Latest message	Completed
		master	63a7f47	-	a few seconds ago
		blog/blue-ocean-editor	d3d9a39	-	4 minutes ago

Clicking on a branch in this list will bring up the [Pipeline Run Details](#) for the latest completed or in-progress Run of that branch. "In Progress" runs can be aborted from this list by clicking on the "Stop" symbol (a square " " inside a circle). Pull requests whose latest run has completed can be run again by clicking the "Play" symbol (a triangle "▶" inside a circle). Clicking the "Edit" symbol (similar to a pencil " ") opens the [pipeline editor](#) on the Pipeline for that branch. Clicking the "Favorite" symbol (a star outline " ") adds a branch to the favorites list shown on the [Dashboard's "Favorites" list](#) for this user. A favorite branch will show a solid star " " and clicking it removes this branch from the favorites.

Pull Requests

The "Pull Requests" tab shows a list of all Pull Requests for the current Pipeline that have a completed or in-progress Run. (Some source control systems call these "Merge Requests", others do not support them at all.) Each line in the list corresponds to a pull request in source control, showing the status of the most recent run, id number, commit information, duration, and when the run completed.



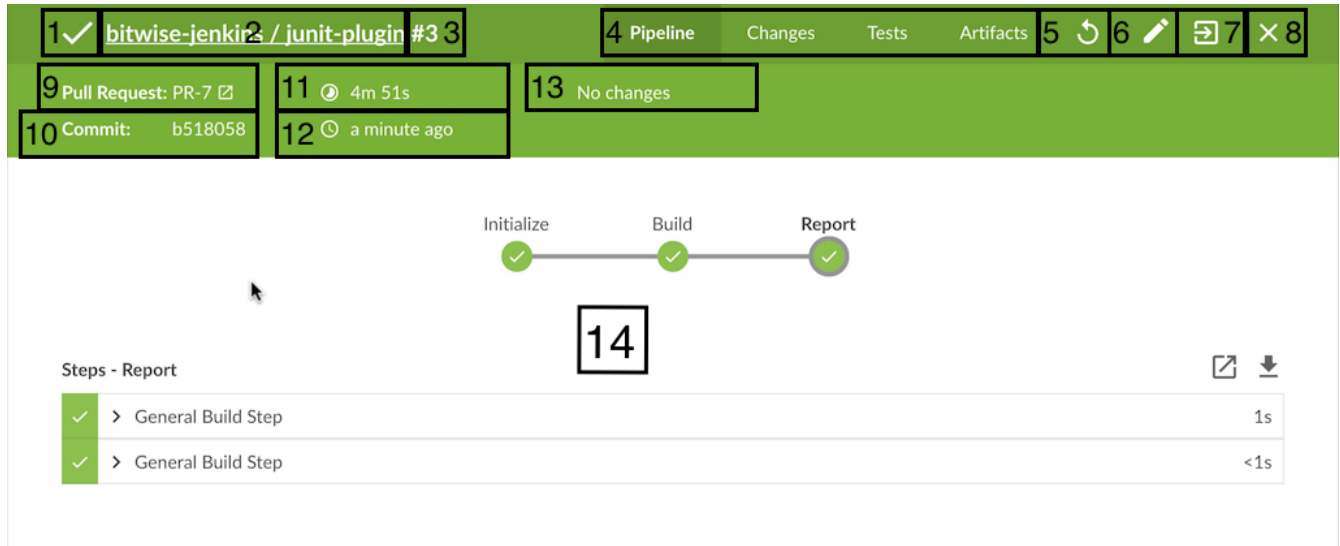
Blue Ocean displays pull requests separately from branches, but otherwise the Pull Requests list behaves similar to the Branches list. Clicking on a pull request in this list will bring up the [Pipeline Run Details](#) for the latest completed or in-progress Run of that pull request. "In Progress" runs can be aborted from this list by clicking on the "Stop" symbol (a square " " inside a circle). Pull requests whose latest run has completed can be run again by clicking the "Play" symbol (a triangle "▶" inside a circle). Pull request do not display "Heath Icons" and cannot be edited or marked as favorites.

NOTE

By default, when a Pull Request is closed, Jenkins will remove the Pipeline from Jenkins (to be cleaned up at a later date), and runs for that Pull Request will not longer be accessible from Jenkins. That can be changed by changing the configuration of the underlying Multi-branch Pipeline job.

Pipeline Run Details View

The Blue Ocean Pipeline Run Details view shows the information related to a single Pipeline Run and allows users to edit or replay that run. Below is a detailed overview of the parts of the Run Details view.



1. **Run Status** - This icon, along with the background color of the top menu bar, indicates the status of this Pipeline run.
2. **Pipeline Name** - The name of this run's Pipeline.
3. **Run Number** - The id number for this Pipeline run. Id numbers unique for each Branch (and Pull Request) of a Pipeline.
4. **Tab Selector** - View one of the detail [tabs](#) for this run. The default is "[Pipeline](#)".
5. **Re-run Pipeline** - Execute this run's Pipeline again.
6. **Edit Pipeline** - Open this run's Pipeline in the [Pipeline Editor](#).
7. **Go to Classic** - Switch to the "Classic" UI view of the details for this run.
8. **Close Details** - This closes the Details view and returns the user to the <<activity, Activity view> for this Pipeline.
9. **Branch or Pull Request** - the branch or pull request for this run.
10. **Commit Id** - Commit id for this run.
11. **Duration** - The duration of this run.
12. **Completed Time** - How long ago the this run completed.
13. **Change Author** - Names of the authors with changes in this run.
14. **Tab View** - Shows the information for the selected tab.

Pipeline Run Status

Blue Ocean makes it easy to see the status of the current Pipeline Run by changing the color of the top menu bar to match the status: blue for "In progress", green for "Passed", yellow for "Unstable", red for "Failed", and gray for "Aborted".

Special cases

Blue Ocean is optimized for working with Pipelines in Source Control, but it can display details for other kinds of projects as well. Blue Ocean offers the same [tabs](#) for all supported projects types, but those tabs may display different information.

Pipelines outside of Souce Control

For Pipelines that are not based on Source Control, Blue Ocean still shows the "Commit Id", "Branch", and "Changes", but those fields are left blank. In this case, the top menu bar does not include the "Edit" option.

Freestyle Projects

For Freestyle projects, Blue Ocean still offers the same [tabs](#), but the Pipeline tab only shows the console log output. The "Rerun" or "Edit" options are also not shown in the top menu bar.

Matrix projects

Matrix projects are not supported in Blue Ocean. Viewing a Matrix project will redirect to the "Classic UI" view for that project.

Tabs

Each of the tabs of the Run Detail view provides information on a specific aspect of a run.

Pipeline

This is the default tab and gives an overall view of the flow of this Pipeline Run. It shows each stage and parallel branch, the steps in those stages, and the console output from those steps. The overview image above shows a successful Pipeline run. If a particular step during the run fails, this tab will automatically default to showing the console log from the failed step. The image below shows a failed Run.

bitwise-jenkins / junit-plugin #1

Pipeline Changes Tests Artifacts

Pull Request: PR-7 48s No changes

Commit: 86b2229 a minute ago

Initialize Build Report

Steps - Initialize

- Fetches the environment variables for a given tool in a list of 'FOO=bar' strings suitable for the withEnv step. <1s
- Use a tool from a predefined Tool Installation <1s
- Fetches the environment variables for a given tool in a list of 'FOO=bar' strings suitable for the withEnv step. <1s
- Use a tool from a predefined Tool Installation <1s
- General SCM 21s
- Shell Script <1s

```
1 [e-jenkins_junit-plugin_PR-7-4ER2Y6SMXFXVYX35EAW7IBEK7PUBWEKHLIPYA6JQPSF2NMLDUTDA] Running shell script
2 + echo PATH =
  /Library/Java/JavaVirtualMachines/jdk1.7.0_80.jdk/Contents/Home/bin:/var/jenkins_home/tools/hudson.tasks.Maven_MavenI
  PATH
```

Changes

bitwise-jenkins / junit-plugin #3

Pipeline Changes Tests Artifacts

Pull Request: PR-8 5m 40s Changes by Liam Newman

Commit: a307242 a few seconds ago

Commit	Author	Message	Date
73fa016	Liam Newman	Fixed failing test	8 minutes ago

Tests

The "Tests" tab shows information about test results for this run. This tab will only contain information if a test result publishing step, such as the "Publish JUnit test results" ([junit](#)) step. If no results are recorded this table will say that, If all tests pass, this tab will report the total number of

passing tests. In the case of failures, the tab will display logs details from the failures as shown below.

The screenshot shows the Jenkins 'Tests' tab for a pipeline named 'bitwise-jenkins / junit-plugin #2'. The top bar is orange and contains navigation links for Pipeline, Changes, Tests, and Artifacts. Below the bar, the status is 'No changes' and the duration is '6m 59s'. A commit hash 'ef7a0ec' is shown with a timestamp of 'a few seconds ago'. The main area displays four statistics: 0 Fixed (green), 1 New Failures (red), 1 Failures (red), and 414 Passing (green). A section titled 'New failing - 1' shows a list item for 'testIssue20090516 - hudson.tasks.junit.CaseResultTest' with a timestamp of 'a few seconds'. The item is expanded to show an error message: 'Error' with the text '1 expected:<...rFirmKeyForVendorRep[Wrong]> but was:<...rFirmKeyForVendorRep[]>'. Below the error is a stacktrace: '1 org.junit.ComparisonFailure: expected:<...rFirmKeyForVendorRep[Wrong]> but was:<...rFirmKeyForVendorRep[]> at hudson.tasks.junit.CaseResultTest.testIssue20090516(CaseResultTest.java:78)'. At the bottom, there is a 'Standard Output' section.

When the previous Run had failures and the current run fixes those failures, this tab will note the fixed texts and display their logs as well.

The screenshot shows the Jenkins 'Tests' tab for a pipeline named 'bitwise-jenkins / junit-plugin #3'. The top bar is green and contains navigation links for Pipeline, Changes, Tests, and Artifacts. Below the bar, the status is 'Changes by Liam Newman' and the duration is '5m 40s'. A commit hash 'a307242' is shown with a timestamp of 'a few seconds ago'. The main area features a large blue banner with a white checkmark icon and the text 'All tests are passing' and 'Nice one! All 415 tests for this pipeline are passing.'. Below the banner, there is a 'Fixed' section with a list item for 'testIssue20090516 - hudson.tasks.junit.CaseResultTest' with a timestamp of 'a few seconds'. The item is expanded to show a green checkmark icon and a mouse cursor pointing to it.

Artifacts

The "Artifacts" tabs show a list of any artifacts saved using the "Archive Artifacts" ([archive](#)) step. Clicking on a item in the list will download it. The full output log from the Run can be downloaded from this list.



Pull Request: PR-8 [↗](#)

🕒 5m 40s

Changes by Liam Newman

Commit: a307242

🕒 a few seconds ago

Name	Size	
pipeline.log	-	
target/junit.hpi	328.1 KB	
target/junit.jar	401.6 KB	

Download All

Pipeline Editor

The Blue Ocean Pipeline Editor is the simplest way for anyone to get started with creating Pipelines in Jenkins. It's also a great way for existing Jenkins users to start adopting Pipeline.

The editor allows users to create and edit Declarative Pipelines, add stages and parallelized tasks that can run at the same time, depending on their needs. When finished, the editor saves the Pipeline to a source code repository as a **Jenkinsfile**. If the Pipeline needs to be changed again, Blue Ocean makes it easy to jump back in into the visual editor to modify the Pipeline at any time.

The screenshot displays the Jenkins Pipeline Editor interface. At the top, the Jenkins logo is on the left, and navigation links for 'Pipelines' and 'Administration' are on the right, along with a 'Logout' button. Below the navigation bar, the breadcrumb path 'bitwise-jenkins / logstash-plugin / master' is shown, with 'Cancel' and 'Save' buttons to its right. The main area is divided into two sections. On the left is a pipeline diagram starting with a 'Start' node, followed by a 'Build' stage, a 'Test' stage, and a 'Deploy' stage. The 'Test' stage is parallelized, containing three sub-steps: 'Platform 1-A', 'Platform 2-A', and 'Platform 3-A'. Each stage has a plus sign icon below it, indicating the ability to add more tasks. On the right is the 'Pipeline Settings' panel, which includes an 'Agent' dropdown menu set to 'any', an 'Environment' section with a table for 'Name' and 'Value', and a plus sign icon to add new environment variables.

Starting the editor

To use the editor a user must first have [created a pipeline in Blue Ocean](#) or have one or more existing Pipelines already created in Jenkins. If editing an existing pipeline, the credentials for that pipeline must allow pushing of changes to the target repository.

The editor can be launched via:

- Dashboard "New Pipeline" button
- Activity View for Single Run
- Pipeline Run Details

Limitations

- SCM-based Declarative Pipelines only
- Credentials must have write permission
- Does not have full parity with Declarative Pipeline
- Pipeline re-ordered and comments removed

Navigation bar

The Pipeline Editor includes the [standard navigation bar](#) at the top, with a local navigation bar below that. The local navigation bar includes:

- **Pipeline Name** - This will include the branch depending on how
- **Cancel** - Discard changes made to the pipeline.
- **Save** - Open the [Save Pipeline Dialog](#).

Pipeline settings

By default, the right side of editor shows the "Pipeline Settings". This sheet can be accessed by clicking anywhere in the [Stage editor](#) that is not a Stage or one of the "Add Stage" buttons.

Agent


The "Agent" section controls what agent the Pipeline will use. This is the same as the ["agent" directive](#).

Environment

The "Environment" sections lets us set environment variables for the Pipeline. This is the same as the ["environment" directive](#).


Stage editor

The left side editor screen contains the Stage editor, used for creating the stages of a Pipeline.



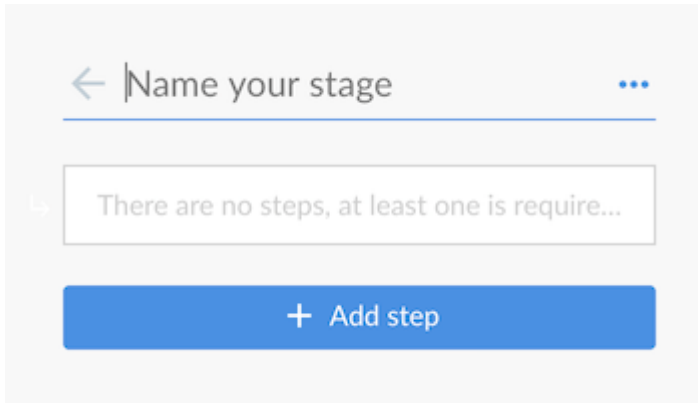
Stages can be added to the Pipeline by clicking the "+" button to the right of an existing stage. Parallel stages can be added by clicking the "+" button below an existing Stage. Stages can be deleted using the [context menu in the stage configuration sheet](#).

The Stage editor will display the name of each Stage once it has been set. Stages that contain incomplete or invalid information will display a warning symbol. Pipelines can have validation errors while they are being edited, but cannot be saved until the errors are fixed.



Stage configuration

Selecting a stage in the Stage editor will open the "Stage Configuration" sheet on the right side. Here we can change the name of the Stage, delete the Stage, and add steps to the Stage.

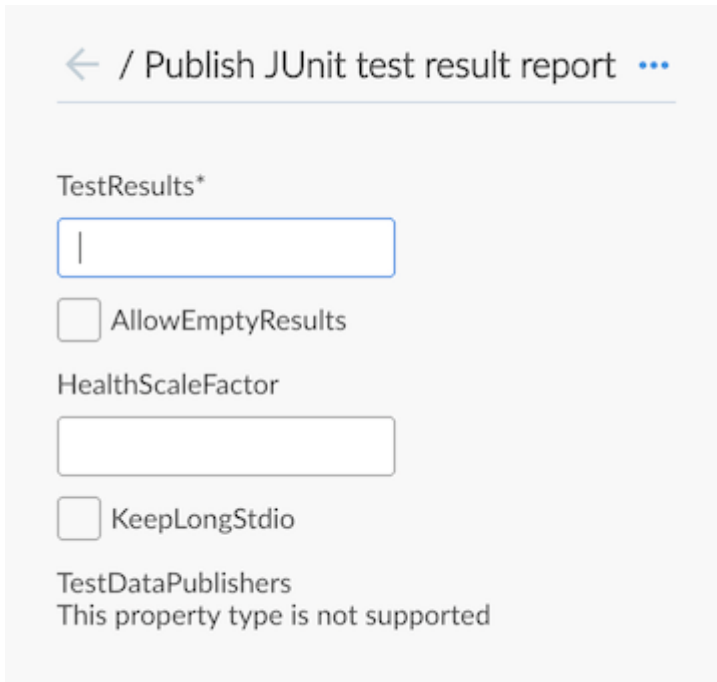


The name of the Stage can be set at the top of the Stage Configuration sheet. The context menu (three dots on the upper right), can be used to delete the current stage. Clicking "Add step" will display the list of available Steps types with a search bar at the top. Steps can be deleted using the [context menu in the step configuration sheet](#). Adding a step or selecting an existing step will open the [step configuration sheet](#).

``

Step configuration

Selecting a step from the Stage configuration sheet will open the Step Configuration sheet.



← / Publish JUnit test result report ...

TestResults*

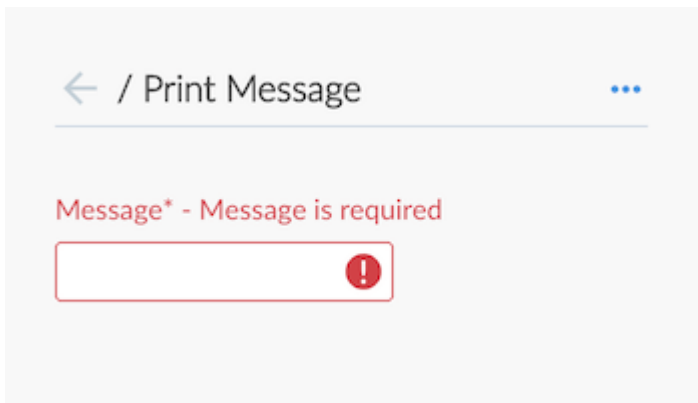
AllowEmptyResults

HealthScaleFactor

KeepLongStdio

TestDataPublishers
This property type is not supported

This sheet will differ depending on the step type, containing whatever fields or controls are needed. The name of the Step cannot be changed. The context menu (three dots on the upper right), can be used to delete the current step. Fields that contain incomplete or invalid information will display a warning symbol. Pipelines can have validation errors while they are being edited, but cannot be saved until the errors are fixed.



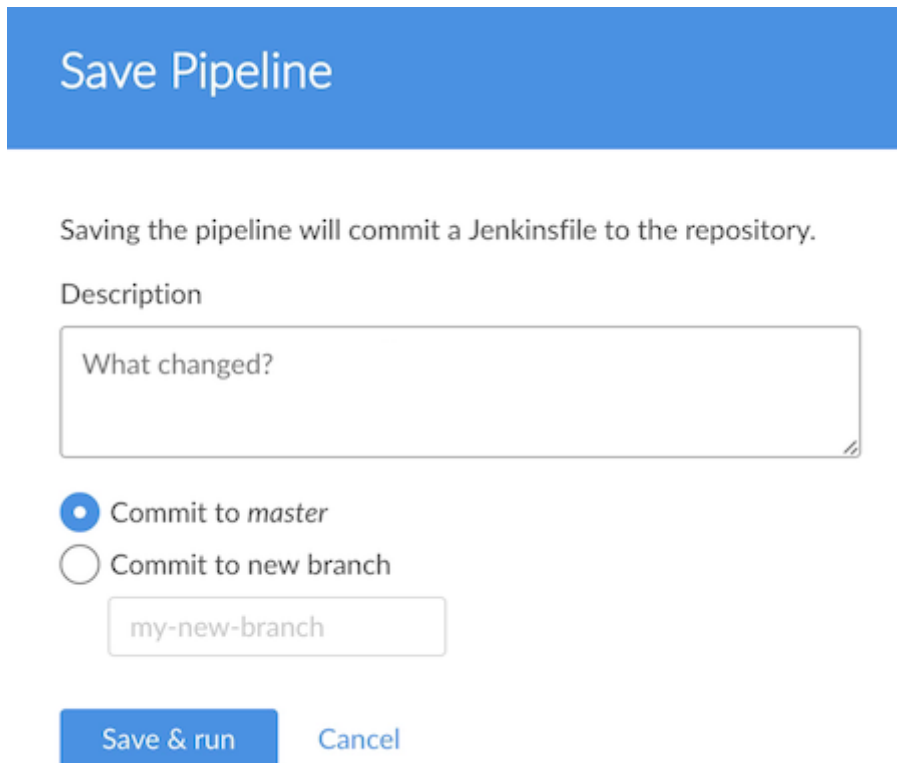
← / Print Message ...

Message* - Message is required

 !

Save Pipeline dialog

In order to be run, changes to a Pipeline must be saved in source control. The "Save Pipeline" dialog controls saving of changes to source control.



Save Pipeline

Saving the pipeline will commit a Jenkinsfile to the repository.

Description

What changed?

Commit to *master*

Commit to new branch

my-new-branch

Save & run Cancel

A helpful description of the changes can be added or left blank. The dialog also supports saving changes the same branch or entering a new branch to save to. Clicking on "Save & run" will save any changes to the Pipeline as a new commit, will start a new Pipeline Run based on those changes, and will navigate to the [Activity View](#) for this pipeline.

Managing Jenkins

This chapter cover how to manage and configure Jenkins masters and nodes.

This chapter is intended for Jenkins administrators. More experienced users may find this information useful, but only to the extent that they will understand what is and is not possible for administrators to do. Individual sections may assume knowledge of information from previous sections, but such assumptions will be explicitly called out and cross-referenced.

If you are not yet familiar with basic Jenkins terminology and features, start with [Getting Started with Jenkins](#).

If you are already familiar with Jenkins basics and would like to delve deeper into how to use specific features, see [Using Jenkins](#).

If you are a system administrator and want learn how to back-up, restore, maintain as Jenkins servers and nodes, see [Jenkins System Administration](#).

Configuring the System

NOTE | This is still very much a work in progress

Managing Security

Jenkins is used everywhere from workstations on corporate intranets, to high-powered servers connected to the public internet. To safely support this wide spread of security and threat profiles, Jenkins offers many configuration options for enabling, editing, or disabling various security features.

As of Jenkins 2.0, many of the security options were enabled by default to ensure that Jenkins environments remained secure unless an administrator explicitly disabled certain protections.

This section will introduce the various security options available to a Jenkins administrator, explaining the protections offered, and trade-offs to disabling some of them.

Enabling Security

When the **Enable Security** checkbox is checked, which has been the default since Jenkins 2.0, users can log in with a username and password in order to perform operations not available to anonymous users. Which operations require users to log in depends on the chosen authorization strategy and its configuration; by default anonymous users have no permissions, and logged in users have full control. This checkbox should **always** be enabled for any non-local (test) Jenkins environment.

The Enable Security section of the web UI allows a Jenkins administrator to enable, configure, or disable key security features which apply to the entire Jenkins environment.



Configure Global Security

Enable security ?

TCP port for JNLP agents Fixed : Random Disable ?

Disable remember me ?

Access Control

Security Realm

Delegate to servlet container ?

Jenkins' own user database ?

Allow users to sign up ?

LDAP ?

Unix user/group database ?

Authorization

Anyone can do anything ?

Legacy mode ?

Logged-in users can do anything ?

Allow anonymous read access ?

Matrix-based security ?

Project-based Matrix Authorization Strategy ?

Markup Formatter ?

Treats all input as plain text. HTML unsafe characters like < and & are escaped to their respective character entities.

JNLP TCP Port

Jenkins uses a TCP port to communicate with agents launched via the JNLP protocol, such as Windows-based agents. As of Jenkins 2.0, by default this port is disabled.

For administrators wishing to use JNLP-based agents, the two port options are:

1. **Random:** The JNLP port is chosen random to avoid collisions on the Jenkins [master](#). The downside to randomized JNLP ports is that they're chosen during the boot of the Jenkins master, making it difficult to manage firewall rules allowing JNLP traffic.
2. **Fixed:** The JNLP port is chosen by the Jenkins administrator and is consistent across reboots of

the Jenkins master. This makes it easier to manage firewall rules allowing JNLP-based agents to connect to the master.

Access Control

Access Control is the primary mechanism for securing a Jenkins environment against unauthorized usage. Two facets of configuration are necessary for configuring Access Control in Jenkins:

1. A **Security Realm** which informs the Jenkins environment how and where to pull user (or identity) information from. Also commonly known as "authentication."
2. **Authorization** configuration which informs the Jenkins environment as to which users and/or groups can access which aspects of Jenkins, and to what extent.

Using both the Security Realm and Authorization configurations it is possible to configure very relaxed or very rigid authentication and authorization schemes in Jenkins.

Additionally, some plugins such as the `plugin:role-strategy`[Role-based Authorization Strategy] plugin can extend the Access Control capabilities of Jenkins to support even more nuanced authentication and authorization schemes.

Security Realm

By default Jenkins includes support for a few different Security Realms:

Delegate to servlet container

For delegating authentication a servlet container running the Jenkins master, such as [Jetty](#). If using this option, please consult the servlet container's authentication documentation.

Jenkins' own user database

Use Jenkins's own built-in user data store for authentication instead of delegating to an external system. This is enabled by default with new Jenkins 2.0 or later installations and is suitable for smaller environments.

LDAP

Delegate all authentication to a configured LDAP server, including both users and groups. This option is more common for larger installations in organizations which already have configured an external identity provider such as LDAP. This also supports Active Directory installations.

NOTE

This feature is provided by the `plugin:ldap`[LDAP plugin] that may not be installed on your instance.

Unix user/group database

Delegates the authentication to the underlying Unix OS-level user database on the Jenkins master. This mode will also allow re-use of Unix groups for authorization. For example, Jenkins can be configured such that "Everyone in the `developers` group has administrator access." To support this feature, Jenkins relies on [PAM](#) which may need to be configured external to the Jenkins environment.

CAUTION

Unix allows an user and a group to have the same name. In order to disambiguate, use the @ prefix to force the name to be interpreted as a group. For example, @dev would mean the dev group and not the dev user.

Plugins can provide additional security realms which may be useful for incorporating Jenkins into existing identity systems, such as:

- plugin:active-directory[Active Directory]
- plugin:github-oauth[GitHub Authentication]
- plugin:crowd2[Atlassian Crowd 2]

Authorization

The Security Realm, or authentication, indicates *who* can access the Jenkins environment. The other piece of the puzzle is **Authorization**, which indicates *what* they can access in the Jenkins environment. By default Jenkins supports a few different Authorization options:

Anyone can do anything

Everyone gets full control of Jenkins, including anonymous users who haven't logged in. **Do not use this setting** for anything other than local test Jenkins masters.

Legacy mode

Behaves exactly the same as Jenkins <1.164. Namely, if a user has the "admin" role, they will be granted full control over the system, and otherwise (including anonymous users) will only have the read access. **Do not use this setting** for anything other than local test Jenkins masters.

Logged in users can do anything

In this mode, every logged-in user gets full control of Jenkins. Depending on an advanced option, anonymous users get read access to Jenkins, or no access at all. This mode is useful to force users to log in before taking actions, so that there is an audit trail of users' actions.

Matrix-based security

This authorization scheme allows for granular control over which users and groups are able to perform which actions in the Jenkins environment (see the screenshot below).

Project-based Matrix Authorization Strategy

This authorization scheme is an extension to Matrix-based security which allows additional access control lists (ACLs) to be defined for **each project** separately in the Project configuration screen. This allows granting specific users or groups access only to specified projects, instead of all projects in the Jenkins environment. The ACLs defined with Project-based Matrix Authorization are additive such that access grants defined in the Configure Global Security screen will be combined with project-specific ACLs.

NOTE Matrix-based security and Project-based Matrix Authorization Strategy are provided by the plugin:matrix-auth[Matrix Authorization Strategy Plugin] and may not be installed on your Jenkins.

For most Jenkins environments, Matrix-based security provides the most security and flexibility so it is recommended as a starting point for "production" environments.

Authorization

- Anyone can do anything
- Legacy mode
- Logged-in users can do anything
- Matrix-based security

User/group	Overall					
	Administer	Configure	UpdateCenter	Read	RunScripts	UploadPlugins
Anonymous	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
admin	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

User/group to add:

Figure 2. Matrix-based security

The table shown above can get quite wide as each column represents a permission provided by Jenkins core or a plugin. Hovering the mouse over a permission will display more information about the permission.

Each row in the table represents a user or group (also known as a "role"). This includes special entries named "anonymous" and "authenticated." The "anonymous" entry represents permissions granted to all unauthenticated users accessing the Jenkins environment. Whereas "authenticated" can be used to grant permissions to all authenticated users accessing the environment.

The permissions granted in the matrix are additive. For example, if a user "kohsuke" is in the groups "developers" and "administrators", then the permissions granted to "kohsuke" will be a union of all those permissions granted to "kohsuke", "developers", "administrators", "authenticated", and "anonymous."

Markup Formatter

Jenkins allows user-input in a number of different configuration fields and text areas which can lead to users inadvertently, or maliciously, inserting unsafe HTML and/or JavaScript.

By default the **Markup Formatter** configuration is set to **Plain Text** which will escape unsafe characters such as `<` and `&` to their respective character entities.

Using the **Safe HTML** Markup Formatter allows for users and administrators to inject useful and information HTML snippets into Project Descriptions and elsewhere.

Cross Site Request Forgery

A cross site request forgery (or CSRF/XSRF) [24: www.owasp.org/index.php/Cross-Site_Request_Forgery] is an exploit that enables an unauthorized third party to perform requests against a web application by impersonating another, authenticated, user. In the context of a Jenkins environment, a CSRF attack could allow an malicious actor to delete projects, alter builds, or modify Jenkins' system configuration. To guard against this class of vulnerabilities, CSRF protection has been enabled by default with all Jenkins versions since 2.0.



The screenshot shows the Jenkins configuration page for 'Crumb Algorithm'. The option 'Prevent Cross Site Request Forgery exploits' is checked and highlighted with a red box. Below it, the 'Crumb Algorithm' section is visible, with 'Default Crumb Issuer' selected and 'Enable proxy compatibility' unchecked. There are help icons (question marks) next to the main option and the 'Enable proxy compatibility' option.

When the option is enabled, Jenkins will check for a CSRF token, or "crumb", on any request that may change data in the Jenkins environment. This includes any form submission and calls to the remote API, including those using "Basic" authentication.

It is **strongly recommended** that this option be left **enabled**, including on instances operating on private, fully trusted networks.

Caveats

CSRF protection *may* result in challenges for more advanced usages of Jenkins, such as:

- Some Jenkins features, like the remote API, are more difficult to use when this option is enabled. Consult the [Remote API](#) documentation for more information.
- Accessing Jenkins through a poorly-configured reverse proxy may result in the CSRF HTTP header being stripped from requests, resulting in protected actions failing.
- Out-dated plugins, not tested with CSRF protection enabled, may not properly function.

More information about CSRF exploits can be found [on the OWASP website](#).

Agent/Master Access Control

Conceptually, the Jenkins master and agents can be thought of as a cohesive system which happens to execute across multiple discrete processes and machines. This allows an agent to ask the master process for information available to it, for example, the contents of files, etc.

For larger or mature Jenkins environments where a Jenkins administrator might enable agents provided by other teams or organizations, a flat agent/master trust model is insufficient.

The Agent/Master Access Control system was introduced [25: Starting with 1.587, and 1.580.1, releases] to allow Jenkins administrators to add more granular access control definitions between the Jenkins master and the connected agents.



As of Jenkins 2.0, this subsystem has been turned on by default.

Customizing Access

For advanced users who may wish to allow certain access patterns from the agents to the Jenkins master, Jenkins allows administrators to create specific exemptions from the built-in access control rules.



By following the link highlighted above, an administrator may edit **Commands** and **File Access** Agent/Master access control rules.

Commands

"Commands" in Jenkins and its plugins are identified by their fully-qualified class names. The majority of these commands are intended to be executed on agents by a request of a master, but some of them are intended to be executed on a master by a request of an agent.

Plugins not yet updated for this subsystem may not classify which category each command falls into, such that when an agent requests that the master execute a command which is not explicitly allowed, Jenkins will err on the side of caution and refuse to execute the command.

In such cases, Jenkins administrators may "whitelist" [26: en.wikipedia.org/wiki/Whitelist] certain commands as acceptable for execution on the master.

Agent → Master Access Control

Jenkins master is now more strict about what commands its agents can send to the master. Unfortunately, this prevents some plugins from functioning correctly, as those plugins do not specify which commands are open for agents to execute and which ones are not. While plugin developers work on updating this, as an administrator, you can mark commands as OK for agents to execute (aka "whitelisting".)

⚠ Please see [the discussion of this feature](#) to understand the security implication of this.

Currently Whitelisted Commands

The following commands are currently whitelisted for agents to execute them on the master. Type in any fully-qualified class names to white list them:

Advanced

Administrators may also whitelist classes by creating files with the `.conf` extension in the directory `JENKINS_HOME/secrets/whitelisted-callables.d/`. The contents of these `.conf` files should list command names on separate lines.

The contents of all the `.conf` files in the directory will be read by Jenkins and combined to create a `default.conf` file in the directory which lists all known safe command. The `default.conf` file will be re-written each time Jenkins boots.

Jenkins also manages a file named `gui.conf`, in the `whitelisted-callables.d` directory, where commands added via the web UI are written. In order to disable the ability of administrators to change whitelisted commands from the web UI, place an empty `gui.conf` file in the directory and change its permissions such that it is not writeable by the operating system user Jenkins runs as.

File Access Rules

The File Access Rules are used to validate file access requests made from agents to the master. Each File Access Rule is a triplet which must contain each of the following elements:

1. **allow / deny**: if the following two parameters match the current request being considered, an **allow** entry would allow the request to be carried out and a **deny** entry would deny the request to be rejected, regardless of what later rules might say.
2. **operation**: Type of the operation requested. The following 6 values exist. The operations can also be combined by comma-separating the values. The value of **all** indicates all the listed operations are allowed or denied.
 - **read**: read file content or list directory entries
 - **write**: write file content
 - **mkdirs**: create a new directory
 - **create**: create a file in an existing directory
 - **delete**: delete a file or directory

- `stat`: read metadata of a file/directory, such as timestamp, length, file access modes.
3. *file path*: regular expression that specifies file paths that matches this rule. In addition to the base regexp syntax, it supports the following tokens:
- `<JENKINS_HOME>` can be used as a prefix to match the master's `JENKINS_HOME` directory.
 - `<BUILDDIR>` can be used as a prefix to match the build record directory, such as `/var/lib/jenkins/job/foo/builds/2014-10-17_12-34-56`.
 - `<BUILDDID>` matches the timestamp-formatted build IDs, like `2014-10-17_12-34-56`.

The rules are ordered, and applied in that order. The earliest match wins. For example, the following rules allow access to all files in `JENKINS_HOME` except the `secrets` folders:

```
# To avoid hassle of escaping every '\' on Windows, you can use / even on Windows.
deny all <JENKINS_HOME>/secrets/*
allow all <JENKINS_HOME>/*
```

Ordering is very important! The following rules are incorrectly written because the 2nd rule will never match, and allow all agents to access all files and folders under `JENKINS_HOME`:

```
allow all <JENKINS_HOME>/*
deny all <JENKINS_HOME>/secrets/*
```

Advanced

Administrators may also add File Access Rules by creating files with the `.conf` extension in the directory `JENKINS_HOME/secrets/filepath-filters.d/`. Jenkins itself generates the `30-default.conf` file on boot in this directory which contains defaults considered the best balance between compatibility and security by the Jenkins project. In order to disable these built-in defaults, replace `30-default.conf` with an empty file which is not writable by the operating system user Jenkins run as.

On each boot, Jenkins will read all `.conf` files in the `filepath-filters.d` directory in alphabetical order, therefore it is good practice to name files in a manner which indicates their load order.

Jenkins also manages `50-gui.conf`, in the `filepath-filters/` directory, where File Access Rules added via the web UI are written. In order to disable the ability of administrators to change the File Access Rules from the web UI, place an empty `50-gui.conf` file in the directory and change its permissions such that is not writeable by the operating system user Jenkins run as.

Disabling

While it is not recommended, if all agents in a Jenkins environment can be considered "trusted" to the same degree that the master is trusted, the Agent/Master Access Control feature may be disabled.

Additionally, all the users in the Jenkins environment should have the same level of access to all configured projects.

An administrator can disable Agent/Master Access Control in the web UI by un-checking the box on the **Configure Global Security** page. Alternatively an administrator may create a file in `JENKINS_HOME/secrets` named `slave-to-master-security-kill-switch` with the contents of `true` and restart Jenkins.

CAUTION

Most Jenkins environments grow over time requiring their trust models to evolve as the environment grows. Please consider scheduling regular "check-ups" to review whether any disabled security settings should be re-enabled.

Managing Tools

NOTE | This is still very much a work in progress

Built-in tool providers

Ant

Ant build step

Git

JDK

Maven

Managing Plugins

Plugins are the primary means of enhancing the functionality of a Jenkins environment to suit organization- or user-specific needs. There are [over a thousand different plugins](#) which can be installed on a Jenkins master and to integrate various build tools, cloud providers, analysis tools, and much more.

Plugins can be automatically downloaded, with their dependencies, from the [Update Center](#). The Update Center is a service operated by the Jenkins project which provides an inventory of open source plugins which have been developed and maintained by various members of the Jenkins community.

This section will cover everything from the basics of managing plugins within the Jenkins web UI, to making changes on the [master's](#) file system.

Installing a plugin

Jenkins provides a couple of different methods for installing plugins on the master:

1. Using the "Plugin Manager" in the web UI.
2. Using the [Jenkins CLI](#) `install-plugin` command.

Each approach will result in the plugin being loaded by Jenkins but may require different levels of access and trade-offs in order to use.

The two approaches require that the Jenkins master be able to download meta-data from an Update Center, whether the primary Update Center operated by the Jenkins project [27: updates.jenkins.io], or a custom Update Center.

The plugins are packaged as self-contained `.hpi` files, which have all the necessary code, images, and other resources which the plugin needs to operate successfully.

From the web UI

The simplest and most common way of installing plugins is through the **Manage Jenkins > Manage Plugins** view, available to administrators of a Jenkins environment.

Under the **Available** tab, plugins available for download from the configured Update Center can be searched and considered:

Filter:

Updates **Available** Installed Advanced

Install	Name	Version
<input type="checkbox"/>	Violations plugin This plug-in generates reports static code violation detectors such as checkstyle, pmd, cpd, findbugs, codenarc, fxcop, stylecop and simian.	0.7.11
<input type="checkbox"/>	Static Analysis Collector Plug-in This plug-in is an add-on for the plug-ins Checkstyle , Dry , FindBugs , PMD , Task Scanner , and Warnings : the plug-in collects the different analysis results and shows the results in a combined trend graph. Additionally, the plug-in provides health reporting and build stability based on these combined results.	1.49
<input type="checkbox"/>	FindBugs Plug-in This plugin generates the trend report for FindBugs , an open source program which uses static analysis to look for bugs in Java code.	4.69

Update information obtained: 1 hr 11 min ago

Most plugins can be installed and used immediately by checking the box adjacent to the plugin and clicking **Install without restart**.

CAUTION

If the list of available plugins is empty, the master might be incorrectly configured or has not yet downloaded plugin meta-data from the Update Center. Clicking the **Check now** button will force Jenkins to attempt to contact its configured Update Center.

Using the Jenkins CLI

Administrators may also use the [Jenkins CLI](#) which provides a command to install plugins. Scripts to manage Jenkins environments, or configuration management code, may need to install plugins without direct user interaction in the web UI. The Jenkins CLI allows a command line user or automation tool to download a plugin and its dependencies.

```
java -jar jenkins-cli.jar -s http://localhost:8080/ install-plugin SOURCE ... [-  
deploy] [-name VAL] [-restart]
```

Installs a plugin either from a file, an URL, or from update center.

SOURCE : If this points to a local file, that file will be installed. If this is an URL, Jenkins downloads the URL and installs that as a plugin. Otherwise the name is assumed to be the short name of the plugin in the existing update center (like "findbugs"), and the plugin will be installed from the update center.

-deploy : Deploy plugins right away without postponing them until the reboot.

-name VAL : If specified, the plugin will be installed as this short name (whereas normally the name is inferred from the source name automatically).

-restart : Restart Jenkins upon successful installation.

Advanced installation

The Update Center only allows the installation of the most recently released version of a plugin. In cases where an older release of the plugin is desired, a Jenkins administrator can download an older `.hpi` archive and manually install that on the Jenkins master.

From the web UI

Assuming a `.hpi` file has been downloaded, a logged-in Jenkins administrator may upload the file from within the web UI:

1. Navigate to the **Manage Jenkins > Manage Plugins** page in the web UI.
2. Click on the **Advanced** tab.
3. Choose the `.hpi` file under the **Upload Plugin** section.
4. **Upload** the plugin file.

HTTP Proxy Configuration

Server

Port

User name

Password

No Proxy Host

Advanced...

Submit

Upload Plugin

You can upload a .hpi file to install a plugin from outside the central plugin repository.

File: No file chosen

Upload

Once a plugin file has been uploaded, the Jenkins master must be manually restarted in order for the changes to take effect.

On the master

Assuming a .hpi file has been explicitly downloaded by a systems administrator, the administrator can manually place the .hpi file in a specific location on the file system.

Copy the downloaded .hpi` file into the `JENKINS_HOME/plugins` directory on the Jenkins master (for example, on Debian systems `JENKINS_HOME` is generally `/var/lib/jenkins`).

The master will need to be restarted before the plugin is loaded and made available in the Jenkins environment.

NOTE

The names of the plugin directories in the Update Site [27: updates.jenkins.io] are not always the same as the plugin’s display name. Searching plugins.jenkins.io for the desired plugin will provide the appropriate link to the .hpi files.

Updating a plugin

Updates are listed in the **Updates** tab of the **Manage Plugins** page and can be installed by checking the checkboxes of the desired plugin updates and clicking the **Download now and install after restart** button.

Filter:

Updates Available Installed Advanced

Install	Name ↓	Version	Installed
<input checked="" type="checkbox"/>	Blue Ocean beta A new user experience for Jenkins	1.0.0-b14	1.0.0-b13

Download now and install after restart Update information obtained: 1 day 19 hr ago **Check now**

By default, the Jenkins master will check for updates from the Update Center once every 24 hours. To manually trigger a check for updates, simply click on the **Check now** button in the **Updates** tab.

Removing a plugin

When a plugin is no longer used in a Jenkins environment, it is prudent to remove the plugin from the Jenkins master. This provides a number of benefits such as reducing memory overhead at boot or runtime, reducing configuration options in the web UI, and removing the potential for future conflicts with new plugin updates.

Uninstalling a plugin

The simplest way to uninstall a plugin is to navigate to the **Installed** tab on the **Manage Plugins** page. From there, Jenkins will automatically determine which plugins are safe to uninstall, those which are not dependencies of other plugins, and present a button for doing so.



Enabled	Name ↓	Version	Previously installed version	Pinned	Uninstall
<input checked="" type="checkbox"/>	Ant Plugin This plugin adds Apache Ant support to Jenkins.	1.4			Uninstall
<input checked="" type="checkbox"/>	bouncycastle API Plugin Provides an stable API to Bouncy Castle related tasks. Plugins using Bouncy Castle should depend on this plugin and not directly on Bouncy Castle	2.16.0			Uninstall
<input checked="" type="checkbox"/>	Structs Plugin Library plugin for DSL plugins that need names for Jenkins objects.	1.5			Uninstall

A plugin may also be uninstalled by removing the corresponding `.hpi` file from the `JENKINS_HOME/plugins` directory on the master. The plugin will continue to function until the master has been restarted.

CAUTION

If a plugin `.hpi` file is removed but required by other plugins, the Jenkins master may fail to boot correctly.

Uninstalling a plugin does **not** remove the configuration that the plugin may have created. If there are existing jobs/nodes/views/builds/etc configurations that reference data created by the plugin, during boot Jenkins will warn that some configurations could not be fully loaded and ignore the unrecognized data.

Since the configuration(s) will be preserved until they are overwritten, re-installing the plugin will result in those configuration values reappearing.

Removing old data

Jenkins provides a facility for purging configuration left behind by uninstalled plugins. Navigate to **Manage Jenkins** and then click on **Manage Old Data** to review and remove old data.

Disabling a plugin

Disabling a plugin is a softer way to retire a plugin. Jenkins will continue to recognize that the plugin is installed, but it will not start the plugin, and no extensions contributed from this plugin will be visible.

A Jenkins administrator may disable a plugin by unchecking the box on the **Installed** tab of the **Manage Plugins** page (see below).



Enabled	Name ↓	Version	Previously installed version	Pinned	Uninstall
<input checked="" type="checkbox"/>	Ant Plugin This plugin adds Apache Ant support to Jenkins.	1.4			Uninstall
<input checked="" type="checkbox"/>	bouncycastle API Plugin Provides an stable API to Bouncy Castle related tasks. Plugins using Bouncy Castle should depend on this plugin and not directly on Bouncy Castle	2.16.0			Uninstall
<input checked="" type="checkbox"/>	Structs Plugin Library plugin for DSL plugins that need names for Jenkins objects.	1.5			Uninstall

A systems administrator may also disable a plugin by creating a file on the Jenkins master, such as: `JENKINS_HOME/plugins/PLUGIN_NAME.hpi.disabled`.

The configuration(s) created by the disabled plugin behave as if the plugin were uninstalled, insofar that they result in warnings on boot but are otherwise ignored.

Pinned plugins

CAUTION

Pinned plugins feature was removed in Jenkins 2.0. Versions later than Jenkins 2.0 do not bundle plugins, instead providing a wizard to install the most useful plugins.

The notion of **pinned plugins** applies to plugins that are bundled with Jenkins 1.x, such as the plugin:matrix-auth[**Matrix Authorization plugin**].

By default, whenever Jenkins is upgraded, its bundled plugins overwrite the versions of the plugins that are currently installed in `JENKINS_HOME`.

However, when a bundled plugin has been manually updated, Jenkins will mark that plugin as pinned to the particular version. On the file system, Jenkins creates an empty file called `JENKINS_HOME/plugins/PLUGIN_NAME.hpi.pinned` to indicate the pinning.

Pinned plugins will never be overwritten by bundled plugins during Jenkins startup. (Newer versions of Jenkins do warn you if a pinned plugin is *older* than what is currently bundled.)

It is safe to update a bundled plugin to a version offered by the Update Center. This is often necessary to pick up the newest features and fixes. The bundled version is occasionally updated, but not consistently.

The Plugin Manager allows plugins to be explicitly unpinned. The `JENKINS_HOME/plugins/PLUGIN_NAME.hpi.pinned` file can also be manually created/deleted to control the pinning behavior. If the `pinned` file is present, Jenkins will use whatever plugin version the user has specified. If the file is absent, Jenkins will restore the plugin to the default version on startup.

Jenkins CLI

Jenkins has a built-in command line interface that allows users and administrators to access Jenkins from a script or shell environment. This can be convenient for scripting of routine tasks, bulk updates, troubleshooting, and more.

The command line interface can be accessed over SSH or with the Jenkins CLI client, a `.jar` file distributed with Jenkins.

WARNING

Use of the CLI client distributed with Jenkins 2.53 and older and Jenkins LTS 2.46.1 and older is **not recommended** for security reasons: while there are no currently known vulnerabilities, several have been reported and patched in the past, and the Jenkins Remoting protocol it uses is inherently vulnerable to remote code execution bugs, even “preauthentication” exploits (by anonymous users able to physically access the Jenkins network).

The client distributed with Jenkins 2.54 and newer and Jenkins LTS 2.46.2 and newer is considered secure in its default (`-http`) or `-ssh` modes, as is using the standard `ssh` command.

Using the CLI over SSH

In a new Jenkins installation, the SSH service is disabled by default. Administrators may choose to set a specific port or ask Jenkins to pick a random port in the [Configure Global Security](#) page. In order to determine the randomly assigned SSH port, inspect the headers returned on a Jenkins URL, for example:

```
% curl -Lv https://JENKINS_URL/login 2>&1 | grep 'X-SSH-Endpoint'  
< X-SSH-Endpoint: localhost:53801  
%
```

With the random SSH port (53801 in this example), and [Authentication](#) configured, any modern SSH client may securely execute CLI commands.

Authentication

Whichever user used for authentication with the Jenkins master must have the [Overall/Read](#) permission in order to *access* the CLI. The user may require additional permissions depending on the commands executed.

Authentication relies on SSH-based public/private key authentication. In order to add an SSH public key for the appropriate user, navigate to [JENKINS_URL/user/USERNAME/configure](#) and paste an SSH public key into the appropriate text area.

Full Name 

Description

API Token

[Show API Token...](#)

My Views

Default View

The view selected by default when navigating to the users private views

Password

Password:

Confirm Password:

SSH Public Keys

SSH Public Keys

Setting for search

Case-sensitivity Insensitive search tool

[Save](#)

[Apply](#)

Common Commands

Jenkins has a number of built-in CLI commands which can be found in every Jenkins environment, such as `build` or `list-jobs`. Plugins may also provide CLI commands; in order to determine the full list of commands available in a given Jenkins environment, execute the CLI `help` command:

```
% ssh -l kohsuke -p 53801 localhost help
```

The following list of commands is not comprehensive, but it is a useful starting point for Jenkins CLI usage.

build

One of the most common and useful CLI commands is **build**, which allows the user to trigger any job or Pipeline for which they have permission.

The most basic invocation will simply trigger the job or Pipeline and exit, but with the additional options a user may also pass parameters, poll SCM, or even follow the console output of the triggered build or Pipeline run.

```
% ssh -l kohsuke -p 53801 localhost help build

java -jar jenkins-cli.jar build JOB [-c] [-f] [-p] [-r N] [-s] [-v] [-w]
Starts a build, and optionally waits for a completion. Aside from general
scripting use, this command can be used to invoke another job from within a
build of one job. With the -s option, this command changes the exit code based
on the outcome of the build (exit code 0 indicates a success) and interrupting
the command will interrupt the job. With the -f option, this command changes
the exit code based on the outcome of the build (exit code 0 indicates a
success) however, unlike -s, interrupting the command will not interrupt the
job (exit code 125 indicates the command was interrupted). With the -c option,
a build will only run if there has been an SCM change.
JOB : Name of the job to build
-c  : Check for SCM changes before starting the build, and if there's no
      change, exit without doing a build
-f  : Follow the build progress. Like -s only interrupts are not passed
      through to the build.
-p  : Specify the build parameters in the key=value format.
-s  : Wait until the completion/abortion of the command. Interrupts are passed
      through to the build.
-v  : Prints out the console output of the build. Use with -s
-w  : Wait until the start of the command
% ssh -l kohsuke -p 53801 localhost build build-all-software -f -v
Started build-all-software #1
Started from command line by admin
Building in workspace /tmp/jenkins/workspace/build-all-software
[build-all-software] $ /bin/sh -xe /tmp/hudson1100603797526301795.sh
+ echo hello world
hello world
Finished: SUCCESS
Completed build-all-software #1 : SUCCESS
%
```

console

Similarly useful is the **console** command, which retrieves the console output for the specified build or Pipeline run. When no build number is provided, the **console** command will output the last completed build's console output.

```
% ssh -l kohsuke -p 53801 localhost help console

java -jar jenkins-cli.jar console JOB [BUILD] [-f] [-n N]
Produces the console output of a specific build to stdout, as if you are doing 'cat
build.log'
JOB   : Name of the job
BUILD : Build number or permalink to point to the build. Defaults to the last
       build
-f    : If the build is in progress, stay around and append console output as
       it comes, like 'tail -f'
-n N  : Display the last N lines
% ssh -l kohsuke -p 53801 localhost console build-all-software
Started from command line by kohsuke
Building in workspace /tmp/jenkins/workspace/build-all-software
[build-all-software] $ /bin/sh -xe /tmp/hudson1100603797526301795.sh
+ echo hello world
yes
Finished: SUCCESS
%
```

who-am-i

The `who-am-i` command is helpful for listing the current user's credentials and permissions available to the user. This can be useful when debugging the absence of CLI commands due to the lack of certain permissions.

```
% ssh -l kohsuke -p 53801 localhost help who-am-i

java -jar jenkins-cli.jar who-am-i
Reports your credential and permissions.
% ssh -l kohsuke -p 53801 localhost who-am-i
Authenticated as: kohsuke
Authorities:
  authenticated
%
```

Using the CLI client

While the SSH-based CLI is fast and covers most needs, there may be situations where the CLI client distributed with Jenkins is a better fit. For example, the default transport for the CLI client is HTTP which means no additional ports need to be opened in a firewall for its use.

Downloading the client

The CLI client can be downloaded directly from a Jenkins master at the URL `/jnlPjars/jenkins-cli.jar`, in effect `JENKINS_URL/jnlPjars/jenkins-cli.jar`

While a CLI `.jar` can be used against different versions of Jenkins, should any compatibility issues arise during use, please re-download the latest `.jar` file from the Jenkins master.

Using the client

The general syntax for invoking the client is as follows:

```
java -jar jenkins-cli.jar [-s JENKINS_URL] [global options...] command [command options...] [arguments...]
```

The `JENKINS_URL` can be specified via the environment variable `$JENKINS_URL`. Summaries of other general options can be displayed by running the client with no arguments at all.

Client connection modes

There are three basic modes in which the 2.54+ / 2.46.2+ client may be used, selectable by global option: `-http`; `-ssh`; and `-remoting`.

HTTP connection mode

This is the default mode as of 2.54 and 2.46.2, though you may pass the `-http` option explicitly for clarity.

Authentication is preferably with an `-auth` option, which takes a `username:apitoken` argument. Get your API token from `/me/configure`:

```
java -jar jenkins-cli.jar [-s JENKINS_URL] -auth kohsuke:abc1234ffe4a command ...
```

(Actual passwords are also accepted, but this is discouraged.)

You can also precede the argument with `@` to load the same content from a file:

```
java -jar jenkins-cli.jar [-s JENKINS_URL] -auth @/home/kohsuke/.jenkins-cli command
...
```

Generally no special system configuration need be done to enable HTTP-based CLI connections. If you are running Jenkins behind an HTTP(S) reverse proxy, ensure it does not buffer request or response bodies.

WARNING

The HTTP(S) connection mode of the CLI in Jenkins 2.54 and newer does not work correctly behind an Apache HTTP reverse proxy server using `mod_proxy`. Workarounds include using a different reverse proxy such as Nginx or HAProxy, or using the SSH connection mode where possible. See [JENKINS-47279](#).

SSH connection mode

Authentication is via SSH keypair. You must select the Jenkins user ID as well:

```
java -jar jenkins-cli.jar [-s JENKINS_URL] -ssh -user kohsuke command ...
```

In this mode, the client acts essentially like a native `ssh` command.

By default the client will try to connect to an SSH port on the same host as is used in the `JENKINS_URL`. If Jenkins is behind an HTTP reverse proxy, this will not generally work, so run Jenkins with the system property `-Dorg.jenkinsci.main.modules.sshd.SSHD.hostName=ACTUALHOST` to define a hostname or IP address for the SSH endpoint.

Remoting connection mode

This was the only mode supported by clients downloaded from a pre-2.54 / pre-2.46.2 Jenkins server (prior to the introduction of the `-remoting` option). Its use is deprecated for security and performance reasons. That said, certain commands or command modes can *only* run in Remoting mode, typically because the command functionality involves running server-supplied code on the client machine.

This mode is disabled on the server side for new installations of 2.54+ and 2.46.2. If you must use it, and accept the risks, it may be enabled in [Configure Global Security](#).

Authentication is preferably via SSH keypair. A `login` command and `--username` / `--password` command (note: **not global**) options are also available; these are discouraged since they cannot work with a non-password-based security realm, certain command arguments will not be properly parsed if anonymous users lack overall or job read access, and saving human-chosen passwords for use in scripts is considered insecure.

Note that there are two transports available for this mode: over HTTP, or over a dedicated TCP socket. If the [TCP port is enabled](#) and seems to work, the client will use this transport. If the TCP port is disabled, or such a port is advertised but does not accept connections (for example because you are using an HTTP reverse proxy with a firewall), the client will automatically fall back to the

less efficient HTTP transport.

Common Problems with the Remoting-based client

There are a number of common problems that may be experienced when running the CLI client.

Operation timed out

Check that the HTTP or TCP port is opened if you are using a firewall on your server. You can configure its value in Jenkins configuration. By default it is set to use a random port.

```
% java -jar jenkins-cli.jar -s JENKINS_URL help
Exception in thread "main" java.net.ConnectException: Operation timed out
    at java.net.PlainSocketImpl.socketConnect(Native Method)
    at java.net.PlainSocketImpl.doConnect(PlainSocketImpl.java:351)
    at java.net.PlainSocketImpl.connectToAddress(PlainSocketImpl.java:213)
    at java.net.PlainSocketImpl.connect(PlainSocketImpl.java:200)
    at java.net.SocksSocketImpl.connect(SocksSocketImpl.java:432)
    at java.net.Socket.connect(Socket.java:529)
    at java.net.Socket.connect(Socket.java:478)
    at java.net.Socket.<init>(Socket.java:375)
    at java.net.Socket.<init>(Socket.java:189)
    at hudson.cli.CLI.<init>(CLI.java:97)
    at hudson.cli.CLI.<init>(CLI.java:82)
    at hudson.cli.CLI._main(CLI.java:250)
    at hudson.cli.CLI.main(CLI.java:199)
```

No X-Jenkins-CLI2-Port

Go to **Manage Jenkins > Configure Global Security** and choose "Fixed" or "Random" under **TCP port for JNLP agents**.

```

java.io.IOException: No X-Jenkins-CLI2-Port among [X-Jenkins, null, Server, X-Content-
Type-Options, Connection,
    X-You-Are-In-Group, X-Hudson, X-Permission-Implied-By, Date, X-Jenkins-
Session, X-You-Are-Authenticated-As,
    X-Required-Permission, Set-Cookie, Expires, Content-Length, Content-Type]
    at hudson.cli.CLI.getCliTcpPort(CLI.java:284)
    at hudson.cli.CLI.<init>(CLI.java:128)
    at hudson.cli.CLIConnectionFactory.connect(CLIConnectionFactory.java:72)
    at hudson.cli.CLI._main(CLI.java:473)
    at hudson.cli.CLI.main(CLI.java:384)
    Suppressed: java.io.IOException: Server returned HTTP response code: 403 for URL:
http://citest.gce.px/cli
        at
sun.net.www.protocol.http.HttpURLConnection.getInputStream0(HttpURLConnection.java:184
0)
        at
sun.net.www.protocol.http.HttpURLConnection.getInputStream(HttpURLConnection.java:1441
)
        at hudson.cli.FullDuplexHttpStream.<init>(FullDuplexHttpStream.java:78)
        at hudson.cli.CLI.connectViaHttp(CLI.java:152)
        at hudson.cli.CLI.<init>(CLI.java:132)
        ... 3 more

```

Server key did not validate

You may get the error below and find a log entry just below that concerning **mismatched keys**:

```

org.apache.sshd.common.SshException: Server key did not validate
    at
org.apache.sshd.client.session.AbstractClientSession.checkKeys(AbstractClientSession.j
ava:523)
    at
org.apache.sshd.common.session.helpers.AbstractSession.handleKexMessage(AbstractSessio
n.java:616)
    ...

```

This means your SSH configuration does not recognize the public key presented by the server. It's often the case when you run Jenkins in dev mode and multiple instances of the application are run under the same SSH port over time.

In a development context, access your `~/.ssh/known_hosts` (or in `C:/Users/<your_name>/.ssh/known_hosts` for Windows) and remove the line corresponding to your current SSH port (e.g. `[localhost]:3485`). In a production context, check with the Jenkins administrator if the public key of the server changed recently. If so, ask the administrator to do the the steps described above.

UsernameNotFoundException

If your client displays a stacktrace that looks like:

```
org.acegisecurity.userdetails.UsernameNotFoundException: <name_you_used>
    ...
```

This means your SSH keys were recognized and validated against the stored users but the username is not valid for the security realm your application is using at the moment. This could occur when you were using the Jenkins database initially, configured your users, and then switched to another security realm (like LDAP, etc.) where the defined users do not exist yet.

To solve the problem, ensure your users exist in your configured security realm.

Troubleshooting logs

To get more information about the authentication process:

1. Go into **Manage Jenkins** > **System Log** > **Add new log recorder**.
2. Enter any name you want and click on **Ok**.
3. Click on **Add**
4. Type `org.jenkinsci.main.modules.sshd.PublicKeyAuthenticatorImpl` (or type `PublicKeyAuth` and then select the full name)
5. Set the level to **ALL**.
6. Repeat the previous three steps for `hudson.model.User`
7. Click on **Save**

When you try to authenticate, you can then refresh the page and see what happen internally.

Script Console

NOTE | This is still very much a work in progress

Managing Nodes

NOTE | This is still very much a work in progress

In-process Script Approval

Jenkins, and a number of plugins, allow users to execute Groovy scripts *in* Jenkins. These scripting capabilities are provided by:

- [Script Console](#).
- [Jenkins Pipeline](#).
- The plugin:email-ext[Extended Email plugin].
- The plugin:groovy[Groovy plugin] - when using the "Execute system Groovy script" step.
- The plugin:job-dsl[JobDSL plugin] as of version 1.60 and later.

To protect Jenkins from execution of malicious scripts, these plugins execute user-provided scripts in a [Groovy Sandbox](#) that limits what internal APIs are accessible. Administrators can then use the "In-process Script Approval" page, provided by the plugin:script-security[Script Security plugin], to manage which unsafe methods, if any, should be allowed in the Jenkins environment.



[About Jenkins](#)

See the version and license information.



[Manage Old Data](#)

Scrub configuration files to remove remnants from old plugins and earlier versions.



[Manage Users](#)

Create/delete/modify users that can log in to this Jenkins



[In-process Script Approval](#)

Allows a Jenkins administrator to review proposed scripts (written e.g. in Groovy) which run inside the Jenkins process and so could bypass security restrictions.



[Prepare for Shutdown](#)

Stops executing new builds, so that the system can be eventually shut down safely.

Getting Started

The `plugin:script-security`[Script Security plugin] is installed automatically by the [Post-install Setup Wizard](#), although initially no additional scripts or operations are approved for use.

IMPORTANT

Older versions of this plugin may not be safe to use. Please review the security warnings listed on `plugin:script-security`[the Script Security plugin page] in order to ensure that the `plugin:script-security`[Script Security plugin] is up to date.

Security for in-process scripting is provided by two different mechanisms: the [Groovy Sandbox](#) and [Script Approval](#). The first, the Groovy Sandbox, is enabled by default for [Jenkins Pipeline](#) allowing user-supplied Scripted and Declarative Pipeline to execute without prior Administrator intervention. The second, Script Approval, allows Administrators to approve or deny unsandboxed scripts, or allow sandboxed scripts to execute additional methods.

For most instances, the combination of the Groovy Sandbox and the [Script Security's built-in list](#) of approved method signatures, will be sufficient. It is strongly recommended that Administrators only deviate from these defaults if absolutely necessary.

Groovy Sandbox

To reduce manual interventions by Administrators, most scripts will run in a Groovy Sandbox by default, including all [Jenkins Pipelines](#). The sandbox only allows a subset of Groovy's methods deemed sufficiently safe for "untrusted" access to be executed without prior approval. Scripts using the Groovy Sandbox are **all** subject to the same restrictions, therefore a Pipeline authored by an Administrator is subject to the restrictions as one authorized by a non-administrative user.

When a script attempts to use features or methods unauthorized by the sandbox, a script is halted immediately, as shown below with Jenkins Pipeline

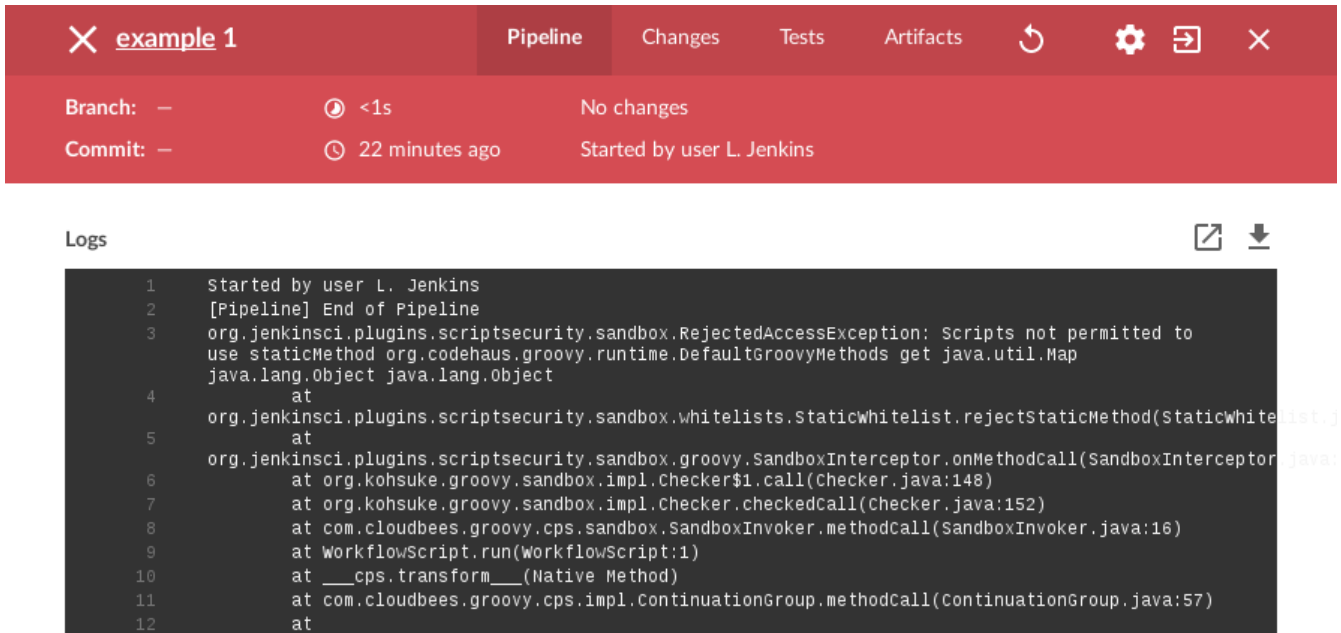


Figure 3. Unauthorized method signature rejected at runtime via Blue Ocean

The Pipeline above will not execute until an Administrator [approves the method signature](#) via the **In-process Script Approval** page.

In addition to adding approved method signatures, users may also disable the Groovy Sandbox entirely as shown below. Disabling the Groovy Sandbox requires that the **entire** script must be reviewed and [manually approved](#) by an administrator.

Pipeline

Definition Pipeline script

Script

```
1 node {  
2   echo 'Hello World'  
3 }
```

⚠ A Jenkins administrator will need to approve this script before it can be used.

Use Groovy Sandbox

[Pipeline Syntax](#)

Figure 4. Disabling the Groovy Sandbox for a Pipeline

Script Approval

Manual approval of entire scripts, or method signatures, by an administrator provides Administrators with additional flexibility to support more advanced usages of in-process scripting. When the [Groovy Sandbox](#) is disabled, or a method outside of the built-in list is invoked, the Script Security plugin will check the Administrator-managed list of approved scripts and methods.

For scripts which wish to execute outside of the [Groovy Sandbox](#), the Administrator must approve the **entire** script in the **In-process Script Approval** page:

/ Groovy script from [notadmin](#) in [unboxed-pipeline](#):

```
node {
  echo 'Hello World'
}
```

You can also remove all previous script approvals:

Figure 5. Approving an unsandboxed Scripted Pipeline

For scripts which use the [Groovy Sandbox](#), but wish to execute an currently unapproved method signature will also be halted by Jenkins, and require an Administrator to approve the specific method signature before the script is allowed to execute:

No pending script approvals.

You can also remove all previous script approvals:

/ / signature : staticMethod
org.codehaus.groovy.runtime.DefaultGroovyMethods get java.util.Map
java.lang.Object java.lang.Object

Signatures already approved:

Figure 6. Approving a new method signature

Approve assuming permissions check

Script approval provides three options: Approve, Deny, and "Approve assuming permissions check." While the purpose of the first two are self-evident, the third requires some additional understanding of what internal data scripts are able to access and how permissions checks inside of Jenkins function.

Consider a script which accesses the method `hudson.model.AbstractItem.getParent()`, which by itself is harmless and will return an object containing either the folder or root item which contains the currently executing Pipeline or Job. Following that method invocation, executing `hudson.model.ItemGroup.getItems()`, which will list items in the folder or root item, requires the `Job/Read` permission.

This could mean that approving the `hudson.model.ItemGroup.getItems()` method signature would allow a script to bypass built-in permissions checks.

Instead, it is usually more desirable to click **Approve assuming permissions check** which will cause the Script Approval engine to allow the method signature assuming the user running the script has the permissions to execute the method, such as the `Job/Read` permission in this example.

Managing Users

NOTE | This is still very much a work in progress

System Administration

This chapter is for system administrators of Jenkins servers and nodes. It will cover system maintenance topics including security, monitoring, and backup/restore.

Users not involved with system-level tasks will find this chapter of limited use. Individual sections may assume knowledge of information from previous sections, but such assumptions will be explicitly called out and cross-referenced.

If you are not yet familiar with basic Jenkins terminology and features, start with [Getting Started with Jenkins](#).

If you are already familiar with Jenkins basics and would like to delve deeper into how to use specific features, see [Using Jenkins](#).

If you are a Jenkins administrator and want to know more about managing Jenkins nodes and instances, see [Managing Jenkins](#).

Backing-up/Restoring Jenkins

NOTE | This is still very much a work in progress

Monitoring Jenkins

NOTE | This is still very much a work in progress

Securing Jenkins

NOTE | This is still very much a work in progress

In the default configuration of Jenkins 1.x, Jenkins does not perform any security checks. This means the ability of Jenkins to launch processes and access local files are available to anyone who can access Jenkins web UI and some more.

Securing Jenkins has two aspects to it.

- Access control, which ensures users are authenticated when accessing Jenkins and their activities are authorized.
- Protecting Jenkins against external threats

Access Control

You should lock down the access to Jenkins UI so that users are authenticated and appropriate set of permissions are given to them. This setting is controlled mainly by two axes:

- **Security Realm**, which determines users and their passwords, as well as what groups the users belong to.
- **Authorization Strategy**, which determines who has access to what.

These two axes are orthogonal, and need to be individually configured. For example, you might choose to use external LDAP or Active Directory as the security realm, and you might choose "everyone full access once logged in" mode for authorization strategy. Or you might choose to let Jenkins run its own user database, and perform access control based on the permission/user matrix.

- [Quick and Simple Security](#) --- if you are running Jenkins like `java -jar jenkins.war` and only need a very simple setup
- [Standard Security Setup](#) --- discusses the most common setup of letting Jenkins run its own user database and do finer-grained access control
- [Apache frontend for security](#) --- run Jenkins behind Apache and perform access control in Apache instead of Jenkins
- [Authenticating scripted clients](#) --- if you need to programmatically access security-enabled Jenkins web UI, use BASIC auth
- [Matrix-based security](#) | [Matrix-based security](#) --- Granting and denying finer-grained permissions

Protect users of Jenkins from other threats

There are additional security subsystems in Jenkins that protect Jenkins and users of Jenkins from indirect attacks.

The following topics discuss features that are **off by default**. We recommend you read them first and act on them immediately.

- [CSRF Protection](#) --- prevent a remote attack against Jenkins running inside your firewall. This feature is **off by default** in Jenkins 1.x and when upgrading to 2.x.
- [Security implication of building on master](#) --- protect Jenkins master from malicious builds
- [Slave To Master Access Control](#) --- protect Jenkins master from malicious build agents

The following topics discuss other security features that are on by default. You'll only need to look at them when they are causing problems.

- [Configuring Content Security Policy](#) --- protect users of Jenkins from malicious builds
- [Markup formatting](#) --- protect users of Jenkins from malicious users of Jenkins

Disabling Security

One may accidentally set up a security realm / authorization in such a way that you may no longer be able to reconfigure Jenkins.

When this happens, you can fix this by the following steps:

1. Stop Jenkins (the easiest way to do this is to stop the servlet container.)
2. Go to `$JENKINS_HOME` in the file system and find `config.xml` file.
3. Open this file in the editor.
4. Look for the `<useSecurity>true</useSecurity>` element in this file.
5. Replace `true` with `false`
6. Remove the elements `authorizationStrategy` and `securityRealm`
7. Start Jenkins
8. When Jenkins comes back, it will be in an unsecured mode where everyone gets full access to the system.

If this is still not working, trying renaming or deleting `config.xml`.

Managing Jenkins with Chef

NOTE | This is still very much a work in progress

Managing Jenkins with Puppet

NOTE | This is still very much a work in progress

Scaling Jenkins

This chapter will cover topics related to using and managing large scale Jenkins configurations: large numbers of users, nodes, agents, folders, projects, concurrent jobs, job results and logs, and even large numbers of masters.

The audience for this chapter is expert Jenkins users, administrators, and those planning large-scale installations.

If you are not yet familiar with basic Jenkins terminology and features, start with [Getting Started with Jenkins](#).

If you are already familiar with Jenkins basics and would like to delve deeper into generally how to use various features, see [Using Jenkins](#).

If you are a Jenkins administrator and want to know more about managing Jenkins nodes and instances, see [Managing Jenkins](#).

Appendix

These sections are generally intended for Jenkins administrators and system administrators. Each section covers a different topic independent of the other sections. They are advanced topics, reference material, and topics that do not fit into other chapters.

WARNING

To Contributors: Please consider adding material elsewhere before adding it here. In fact, topics that do not fit elsewhere may even be out of scope for this handbook. See [Contributing to Jenkins](#) for details of how to contact project contributors and discuss what you want to add.

Glossary

General Terms

Agent

An agent is typically a machine, or container, which connects to a Jenkins master and executes tasks when directed by the master.

Artifact

An immutable file generated during a [Build](#) or [Pipeline](#) run which is **archived** onto the Jenkins [Master](#) for later retrieval by users.

Build

Result of a single execution of a [Project](#)

Cloud

A System Configuration which provides dynamic [Agent](#) provisioning and allocation, such as that provided by the plugin:azure-vm-agents[Azure VM Agents] or plugin:ec2[Amazon EC2] plugins.

Core

The primary Jenkins application (`jenkins.war`) which provides the basic web UI, configuration, and foundation upon which [Plugins](#) can be built.

Downstream

A configured [Pipeline](#) or [Project](#) which is triggered as part of the execution of a separate Pipeline or Project.

Executor

A slot for execution of work defined by a [Pipeline](#) or [Project](#) on a [Node](#). A Node may have zero or more Executors configured which corresponds to how many concurrent Projects or Pipelines are able to execute on that Node.

Fingerprint

A hash considered globally unique to track the usage of an [Artifact](#) or other entity across multiple [Pipelines](#) or [Projects](#).

Folder

An organizational container for [Pipelines](#) and/or [Projects](#), similar to folders on a file system.

Item

An entity in the web UI corresponding to either a: [Folder](#), [Pipeline](#), or [Project](#).

Job

A deprecated term, synonymous with [Project](#).

Label

User-defined text for grouping [Agents](#), typically by similar functionality or capability. For example `linux` for Linux-based agents or `docker` for Docker-capable agents.

Master

The central, coordinating process which stores configuration, loads plugins, and renders the various user interfaces for Jenkins.

Node

A machine which is part of the Jenkins environment and capable of executing [Pipelines](#) or [Projects](#). Both the [Master](#) and [Agents](#) are considered to be Nodes.

Project

A user-configured description of work which Jenkins should perform, such as building a piece of software, etc.

Pipeline

A user-defined model of a continuous delivery pipeline, for more read the [Pipeline chapter](#) in this handbook.

Plugin

An extension to Jenkins functionality provided separately from Jenkins [Core](#).

Publisher

Part of a [Build](#) after the completion of all configured [Steps](#) which publishes reports, sends notifications, etc.

Stage

[stage](#) is part of Pipeline, and used for defining a conceptually distinct subset of the entire Pipeline, for example: "Build", "Test", and "Deploy", which is used by many plugins to visualize or present Jenkins Pipeline status/progress.

Step

A single task; fundamentally steps tell Jenkins *what* to do inside of a [Pipeline](#) or [Project](#).

Trigger

A criteria for triggering a new [Pipeline](#) run or [Build](#).

Update Center

Hosted inventory of plugins and plugin metadata to enable plugin installation from within Jenkins.

Upstream

A configured [Pipeline](#) or [Project](#) which triggers a separate Pipeline or Project as part of its execution.

Workspace

A disposable directory on the file system of a [Node](#) where work can be done by a [Pipeline](#) or [Project](#). Workspaces are typically left in place after a [Build](#) or [Pipeline](#) run completes unless specific Workspace cleanup policies have been put in place on the Jenkins [Master](#).