



# Red Hat CodeReady Workspaces 1.2

## Administration Guide

Installing and administering Red Hat CodeReady Workspaces 1.2



# Red Hat CodeReady Workspaces 1.2 Administration Guide

---

Installing and administering Red Hat CodeReady Workspaces 1.2

Supriya Takkhi  
sbharadw@redhat.com

Robert Kratky  
rkratky@redhat.com

Michal Maléř  
mmaler@redhat.com

## Legal Notice

Copyright © 2020 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux<sup>®</sup> is the registered trademark of Linus Torvalds in the United States and other countries.

Java<sup>®</sup> is a registered trademark of Oracle and/or its affiliates.

XFS<sup>®</sup> is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL<sup>®</sup> is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js<sup>®</sup> is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack<sup>®</sup> Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

## Abstract

Information for administrators installing and operating Red Hat CodeReady Workspaces.

## Table of Contents

<b>CHAPTER 1. UNDERSTANDING RED HAT CODEREADY WORKSPACES</b>	<b>8</b>
<b>CHAPTER 2. INSTALLING CODEREADY WORKSPACES ON OPENSIFT V3</b>	<b>9</b>
2.1. DOWNLOADING THE RED HAT OPENSIFT ORIGIN CLIENT TOOLS	10
2.2. DOWNLOADING THE CODEREADY WORKSPACES DEPLOYMENT SCRIPT	11
2.3. RUNNING THE CODEREADY WORKSPACES DEPLOYMENT SCRIPT	11
2.3.1. Deploying CodeReady Workspaces with default settings	13
2.3.2. Deploying CodeReady Workspaces with a self-signed certificate and OpenShift OAuth	14
2.3.3. Deploying CodeReady Workspaces with a public certificate	14
2.3.4. Deploying CodeReady Workspaces with external Red Hat Single Sign-On	14
2.3.5. Deploying CodeReady Workspaces with external Red Hat SSO and PostgreSQL	15
2.4. VIEWING CODEREADY WORKSPACES INSTALLATION LOGS	16
2.4.1. Viewing CodeReady Workspaces installation logs in the terminal	16
2.4.2. Viewing CodeReady Workspaces installation logs in the OpenShift console	17
2.5. CONFIGURING CODEREADY WORKSPACES TO WORK BEHIND A PROXY SERVER	18
Additional resources	19
2.6. CODEREADY WORKSPACES DEPLOYMENT SCRIPT PARAMETERS	19
<b>CHAPTER 3. INSTALLING CODEREADY WORKSPACES ON OPENSIFT 4 FROM OPERATORHUB</b>	<b>22</b>
<b>CHAPTER 4. INSTALLING CODEREADY WORKSPACES IN RESTRICTED ENVIRONMENTS</b>	<b>27</b>
4.1. PREPARING CODEREADY WORKSPACES DEPLOYMENT FROM A LOCAL REGISTRY	27
4.2. RUNNING THE CODEREADY WORKSPACES DEPLOYMENT SCRIPT IN A RESTRICTED ENVIRONMENT	28
4.3. STARTING WORKSPACES IN RESTRICTED ENVIRONMENTS	29
4.4. MAKING CODEREADY WORKSPACES IMAGES AVAILABLE FROM A LOCAL REGISTRY	33
<b>CHAPTER 5. INSTALLING CODEREADY WORKSPACES ON OPENSIFT DEDICATED</b>	<b>37</b>
<b>CHAPTER 6. UPGRADING CODEREADY WORKSPACES</b>	<b>38</b>
6.1. UPGRADING CODEREADY WORKSPACES USING A MIGRATION SCRIPT	38
6.2. RUNNING THE CODEREADY WORKSPACES MIGRATION SCRIPT IN RESTRICTED ENVIRONMENTS	38
6.3. UPGRADING CODEREADY WORKSPACES MANUALLY	40
<b>CHAPTER 7. UNINSTALLING CODEREADY WORKSPACES</b>	<b>44</b>
<b>CHAPTER 8. VIEWING CODEREADY WORKSPACES OPERATION LOGS</b>	<b>45</b>
8.1. VIEWING CODEREADY WORKSPACES OPERATION LOGS IN THE TERMINAL	45
8.2. VIEWING CODEREADY WORKSPACES OPERATION LOGS IN THE OPENSIFT CONSOLE	46
<b>CHAPTER 9. USING THE CHE 7 IDE IN CODEREADY WORKSPACES</b>	<b>48</b>
<b>CHAPTER 10. USING THE ANALYTICS PLUG-IN IN CODEREADY WORKSPACES</b>	<b>50</b>
<b>CHAPTER 11. USING VERSION CONTROL</b>	<b>51</b>
11.1. GENERATING AND UPLOADING SSH KEYS	51
11.1.1. Using existing SSH keys	51
11.2. CONFIGURING GITHUB OAUTH	54
11.3. CONFIGURING GITLAB OAUTH	56
11.4. SUBMITTING PULL REQUESTS USING THE BUILT-IN PULL REQUEST PANEL	56
11.5. SAVING COMMITTER NAME AND EMAIL	56
11.6. INTERACTING WITH GIT FROM A WORKSPACE	56
11.7. GIT STATUS HIGHLIGHTING IN THE PROJECT TREE AND EDITOR	57
11.8. PERFORMING GIT OPERATIONS	58

11.8.1. Committing	58
11.8.2. Pushing and pulling	59
11.8.3. Managing branches	60
11.9. REVIEWING CHANGED FILES	61
<b>CHAPTER 12. CODEREADY WORKSPACES ADMINISTRATION GUIDE</b>	<b>65</b>
12.1. RAM PREREQUISITES	65
12.1.1. Setting default workspace RAM limits	65
12.2. REQUIREMENTS FOR RESOURCE ALLOCATION AND QUOTAS	65
12.3. SETTING UP THE PROJECT WORKSPACE	65
12.3.1. Setting up a single OpenShift project	65
12.3.2. Setting up a multi OpenShift project	66
12.4. HOW THE CODEREADY WORKSPACES SERVER USES PVCS AND PVS FOR STORAGE	66
12.4.1. Storage requirements for CodeReady Workspaces infrastructure	66
12.4.2. Storage strategies for workspaces	66
12.4.3. Unique PVC strategy	67
12.4.4. Common PVC Strategy	67
12.4.4.1. How the common PVC strategy works	67
12.4.4.2. Enabling the common strategy	67
12.4.4.3. Restrictions on using the common PVC strategy	68
12.4.5. Per workspace PVC strategy	68
12.4.5.1. How the per-workspace PVC strategy works	68
12.5. UPDATING YOUR CODEREADY WORKSPACES DEPLOYMENT	68
12.6. SCALABILITY	69
12.7. GDPR	69
12.8. DEBUG MODE	69
12.9. PRIVATE DOCKER REGISTRIES	69
12.10. CODEREADY WORKSPACES SERVER LOGS	69
12.11. WORKSPACE LOGS	69
12.12. CODEREADY WORKSPACES MASTER STATES	70
12.13. WORKSPACE TERMINATION GRACE PERIOD	70
12.14. AUTO-STOPPING A WORKSPACE WHEN ITS PODS ARE REMOVED	70
12.15. UPDATING CODEREADY WORKSPACES WITHOUT STOPPING ACTIVE WORKSPACES	71
12.15.1. Performing a recreate update	71
12.15.2. Performing a rolling update	71
12.15.2.1. Known issues	72
12.15.3. Updating with database migrations or API incompatibility	72
12.16. DELETING DEPLOYMENTS	72
12.17. MONITORING CODEREADY WORKSPACES MASTER SERVER	73
12.18. CREATING WORKSPACE OBJECTS IN PERSONAL NAMESPACES	73
12.19. OPENSIFT IDENTITY PROVIDER REGISTRATION	74
12.20. CONFIGURING CODEREADY WORKSPACES	75
12.21. PROVIDING THE OPENSIFT CERTIFICATE TO KEYCLOAK	75
<b>CHAPTER 13. MANAGING USERS</b>	<b>77</b>
13.1. AUTHORIZATION AND USER MANAGEMENT	77
13.2. CONFIGURING CODEREADY WORKSPACES TO WORK WITH KEYCLOAK	77
13.3. CONFIGURING KEYCLOAK TOKENS	77
13.4. SETTING UP USER FEDERATION	78
13.5. ENABLING AUTHENTICATION WITH SOCIAL ACCOUNTS AND BROKERING	78
13.6. USING PROTOCOL-BASED PROVIDERS	80
13.7. MANAGING USERS	80
13.8. CONFIGURING SMTP AND EMAIL NOTIFICATIONS	80

---

13.9. CODEREADY WORKSPACES AUTHENTICATION	80
13.9.1. Authentication on CodeReady Workspaces Master	80
13.9.1.1. OpenId	80
13.9.1.1.1. User Profile	81
13.9.1.1.2. Obtaining Token From Keycloak	82
13.9.1.2. Other authentication implementations	82
13.9.1.3. OAuth	82
13.9.2. Authentication on CodeReady Workspaces Agents	83
13.9.2.1. Machine JWT Token	83
13.9.2.2. Authentication schema	84
13.9.2.3. Obtaining Machine Token	85
13.9.3. Using Swagger or REST Clients	85
13.10. PERMISSIONS	86
13.10.1. Overview	86
13.10.2. Workspace permissions	86
13.10.3. Organization permissions	87
13.10.4. System permissions	87
13.10.5. manageSystem permission	87
13.10.6. monitorSystem permission	88
13.10.7. Super-privileged mode	89
13.10.8. Stack permissions	89
13.10.9. Permissions API	90
13.10.10. Listing permissions	90
13.10.11. Listing permissions for a user	90
13.10.12. Listing permissions for all users	91
13.10.13. Assigning permissions	91
13.10.14. Sharing permissions	92
13.11. ORGANIZATIONS	92
13.11.1. Organizations in CodeReady Workspaces	92
13.11.1.1. Roles in an organization	92
13.11.1.2. Root organizations and sub-organizations	93
13.11.1.3. Creating an organization	93
13.11.1.4. Displaying the list of organizations	93
13.11.1.5. Adding members to organizations	93
13.11.1.6. Workspaces in organizations	93
13.11.1.7. Setting email notifications	93
13.11.1.8. Creating sub-organizations	94
13.11.1.9. Adding members to sub-organizations	94
13.11.1.10. Organization and sub-organization administration	94
13.11.1.11. Renaming an organization or sub-organization	94
13.11.1.12. Leaving an organization or sub-organization	95
13.11.1.13. Deleting an organization or sub-organization	95
13.11.1.14. Allocating resources for organizations	95
13.11.1.15. Managing limits	95
13.11.1.16. Updating organization and sub-organization member roles	95
13.11.1.17. Removing members from an organization and sub-organization	96
13.12. RESOURCE MANAGEMENT	96
13.12.1. Overview	96
13.12.2. Resource API	97
13.12.3. Distributing resources	97
13.12.4. Configuring workspaces and resources	97
13.12.5. Unit formats	98
13.12.6. Resource-free API	99

---

13.12.7. Organization Resource API	99
<b>CHAPTER 14. ADMINISTERING WORKSPACES</b>	<b>101</b>
14.1. WORKSPACE	101
14.2. ENVIRONMENT	101
14.3. MACHINE	101
14.4. RECIPE	101
14.5. BOOTSTRAPPER	101
14.6. INSTALLER	101
14.7. VOLUME	102
14.8. ENVIRONMENT VARIABLES	102
14.9. WHAT IS NEXT?	102
14.10. MANAGING WORKSPACES	102
14.10.1. Creating workspaces	102
14.10.1.1. Creating a workspace from stacks in the dashboard	102
14.10.1.2. Duplicating an existing stack	103
14.10.1.3. Creating a custom stack from a custom recipe	104
14.10.2. Starting workspaces	105
14.10.3. Managing a workspace	105
14.11. COMMANDS AND IDE MACROS	106
14.11.1. Command goals	106
14.11.2. Command context	106
14.11.3. Managing commands	107
14.11.4. Macros list	108
14.11.5. Auto-completing macros	109
14.11.6. Using commands	109
14.11.6.1. Using the Command Palette	110
14.11.6.2. Using the command toolbar	110
14.11.6.2.1. Using the Run and Debug buttons	111
14.11.6.2.2. Using the command controller	111
14.11.6.2.3. Using the Preview button	112
14.11.7. Authoring command instructions	112
14.11.8. Macros	113
14.11.9. Environment variables	115
14.12. STACKS	115
14.12.1. Stack overview	115
14.12.2. Importing community supported stacks and applications	115
14.12.3. Sharing stacks and system stacks	116
14.12.4. Loading stacks	117
14.12.5. Creating stacks in CodeReady Workspaces	117
14.12.5.1. Choosing a base image	118
14.12.5.1.1. Using a certified CodeReady Workspaces image	118
14.12.5.1.2. Using a custom image	118
14.12.5.2. Building a custom stack	119
14.12.5.2.1. Building a Docker image	119
14.12.5.2.2. Uploading an image to the registry	119
14.12.5.2.3. Creating a custom stack	119
14.12.5.3. Sharing stacks	120
14.13. RECIPES	120
14.13.1. Supported Recipe Formats	120
14.13.2. Docker-formatted container image requirements and limitations	120
14.13.3. Dockerfile definition and limitations	120
14.13.4. Running multi-container workspaces using Compose files	121



14.13.4.1. Accessing remote files	121
14.13.4.2. Using private repositories	121
14.13.4.3. Configuring privileged access	121
14.13.4.4. Special considerations when using Compose files	122
14.13.5. Kubernetes YAML limitations and restrictions	123
14.14. SERVERS	125
14.14.1. What are servers?	125
14.14.2. Preview URLs	126
14.14.3. Getting preview URLs	126
14.14.4. Exposing internal servers	127
14.14.5. Exposing secure servers	127
14.15. INSTALLERS	128
14.15.1. What are installers?	128
14.15.2. How installers work	128
14.15.3. What happens when enabling and disabling installers	128
14.15.4. Troubleshooting installer failures	129
14.15.4.1. Permission denied failure	129
14.15.4.2. Permission to files and directories failures	129
14.15.5. Installer registry and REST API	129
14.16. VOLUMES	130
14.16.1. Default volumes for workspace containers	130
14.16.2. Adding volumes	130
14.16.3. Configuring workspaces	131
14.17. ENVIRONMENT VARIABLES	131
14.18. PROJECTS	132
14.18.1. Creating projects in workspaces	132
14.18.2. Defining project types	133
14.18.3. Creating a sub-project	134
14.18.4. Navigating the project tree	134
14.19. TROUBLESHOOTING FAILURES IN STARTING THE WORKSPACE	134
14.19.1. Incorrect environment recipes	134
14.19.1.1. Viewing logs from a failed workspace start	135
14.19.2. Restrictive network settings	135
14.19.2.1. Troubleshooting network setting when workspace agent cannot be reached	135
14.19.3. Failure in bootstrapping	135
14.20. WORKSPACE DATA MODEL	136
14.20.1. Environment recipes	136
14.20.2. Projects	137
14.20.3. Commands	139
14.20.4. Runtime	139
14.21. GETTING STARTED WITH FACTORIES	140
14.21.1. Trying a factory	140
14.21.2. Using factories	140
14.21.3. Invoking factories using their unique hashcodes	141
14.21.4. Invoking a named factory	141
14.21.5. Invoking a factory for a specific git repository	141
14.21.6. Next steps	141
14.21.7. Creating Factories	141
14.21.7.1. Creating a factory in the dashboard	141
14.21.7.2. Creating a factory in the IDE	142
14.21.7.3. Creating a factory based on a repository	142
14.21.7.3.1. Customizing URL factories	142
14.21.7.4. Configuring factory policies	143

14.21.7.4.1. Setting factory limitations	143
14.21.7.4.2. Setting factory multiplicity	143
14.21.7.5. Customizing the IDE	143
14.21.7.6. Lifecycle Events	143
14.21.7.7. Factory actions	144
14.21.7.8. Finding and replacing variables	144
14.21.7.9. Pull request workflow	144
14.21.7.10. Repository badging	145
14.21.7.11. Next steps	145
14.21.8. Factories JSON Reference	145
14.21.8.1. Mixins	146
14.21.8.2. Pull Request mixin attributes	147
14.21.8.3. Policies	147
14.21.8.4. Limitations	148
14.21.8.5. Multiplicity	148
14.21.8.6. Customizing the IDE	148
14.21.8.7. Action: Open File	150
14.21.8.8. Action: Find and Replace	151
14.21.8.9. Creator	151



# CHAPTER 1. UNDERSTANDING RED HAT CODEREADY WORKSPACES

Red Hat CodeReady Workspaces is a developer workspace server and cloud IDE. Workspaces are defined as project code files and all of their dependencies necessary to edit, build, run, and debug them. Each workspace has its own private IDE hosted within it. The IDE is accessible through a browser. The browser downloads the IDE as a single-page web application.

Red Hat CodeReady Workspaces provides:

- Workspaces that include runtimes and IDEs
- RESTful workspace server
- A browser-based IDE
- Plugins for languages, framework, and tools
- An SDK for creating plugins and assemblies

## Additional resources

See Red Hat CodeReady Workspaces [CodeReady Workspaces 1.2 Release Notes and Known Issues](#) for more details about the current version.

## CHAPTER 2. INSTALLING CODEREADY WORKSPACES ON OPENSIFT V3

This section describes how to obtain installation files for Red Hat CodeReady Workspaces and how to use them to deploy the product on an instance of OpenShift (such as Red Hat OpenShift Container Platform) v3.

### Prerequisites

#### Minimum hardware requirements

- Minimum 5 GB RAM to run CodeReady Workspaces. The Red Hat Single Sign-On (Red Hat SSO) authorization server and the PostgreSQL database require extra RAM. CodeReady Workspaces uses RAM in the following distribution:
  - The CodeReady Workspaces server: Approximately 750 MB
  - Red Hat SSO: Approximately 1 GB
  - PostgreSQL: Approximately 515 MB
  - Workspaces: 2 GB of RAM per workspace. The total workspace RAM depends on the size of the workspace runtime(s) and the number of concurrent workspace pods.

#### Software requirements

- CodeReady Workspaces deployment script and configuration file
- Container images required for deployment:



#### IMPORTANT

- The container images are now published on the [registry.redhat.io](https://registry.redhat.io) registry and will only be available from [registry.redhat.io](https://registry.redhat.io). For details, see [Red Hat Container Registry Authentication](#).
- To use this registry you must be logged into it.
  - To authorize from the local Docker daemon, see [Using Authentication](#).
  - To authorize from an OpenShift cluster, see [Allowing Pods to Reference Images from Other Secured Registries](#).
- It is not necessary to download any of the referenced images manually.
  - All container images required for deployment are automatically downloaded by the CodeReady Workspaces deployment script.
  - Stack images are automatically downloaded by CodeReady Workspaces when new workspaces are created.
- [registry.redhat.io/codeready-workspaces/server-rhel8:1.2](https://registry.redhat.io/codeready-workspaces/server-rhel8:1.2)
- [registry.redhat.io/codeready-workspaces/server-operator-rhel8:1.2](https://registry.redhat.io/codeready-workspaces/server-operator-rhel8:1.2)
- [registry.redhat.io/rhsc/postgresql-96-rhel7:1-40](https://registry.redhat.io/rhsc/postgresql-96-rhel7:1-40)

- [registry.redhat.io/redhat-sso-7/sso73-openshift:1.0-11](https://registry.redhat.io/redhat-sso-7/sso73-openshift:1.0-11)
- [registry.redhat.io/ubi8-minimal:8.0-127](https://registry.redhat.io/ubi8-minimal:8.0-127)
- Container images with preconfigured stacks for creating workspaces:
  - [registry.redhat.io/codeready-workspaces/stacks-java-rhel8:1.2](https://registry.redhat.io/codeready-workspaces/stacks-java-rhel8:1.2)
  - [registry.redhat.io/codeready-workspaces/stacks-node-rhel8:1.2](https://registry.redhat.io/codeready-workspaces/stacks-node-rhel8:1.2)
  - [registry.redhat.io/codeready-workspaces/stacks-php-rhel8:1.2](https://registry.redhat.io/codeready-workspaces/stacks-php-rhel8:1.2)
  - [registry.redhat.io/codeready-workspaces/stacks-python-rhel8:1.2](https://registry.redhat.io/codeready-workspaces/stacks-python-rhel8:1.2)
  - [registry.redhat.io/codeready-workspaces/stacks-dotnet-rhel8:1.2](https://registry.redhat.io/codeready-workspaces/stacks-dotnet-rhel8:1.2)
  - [registry.redhat.io/codeready-workspaces/stacks-golang-rhel8:1.2](https://registry.redhat.io/codeready-workspaces/stacks-golang-rhel8:1.2)
  - [registry.redhat.io/codeready-workspaces/stacks-java-rhel8:1.2](https://registry.redhat.io/codeready-workspaces/stacks-java-rhel8:1.2)
  - [registry.redhat.io/codeready-workspaces/stacks-cpp-rhel8:1.2](https://registry.redhat.io/codeready-workspaces/stacks-cpp-rhel8:1.2)
  - [registry.redhat.io/codeready-workspaces/stacks-node:1.2](https://registry.redhat.io/codeready-workspaces/stacks-node:1.2)

## Other

To be able to download the CodeReady Workspaces deployment script, Red Hat asks that you register for the free [Red Hat Developer Program](#). This allows you to agree to license conditions of the product. For instructions on how to obtain the deployment script, see [Section 2.2, “Downloading the CodeReady Workspaces deployment script”](#).

## 2.1. DOWNLOADING THE RED HAT OPENSIFT ORIGIN CLIENT TOOLS

This procedure describes steps to obtain and unpack the archive with the Red Hat OpenShift Origin Client Tools.

The CodeReady Workspaces deployment and migration scripts require the use of OpenShift Origin Client Tools 3.11. Later versions may be supported in future but as deployment of CodeReady Workspaces to OpenShift Container Platform and OpenShift Dedicated 4 can be done via the embedded OperatorHub, using a deployment script is no longer necessary.

### Procedure

1. Change to a temporary directory. Create it if necessary. For example:

```
$ mkdir ~/tmp  
$ cd ~/tmp
```

2. Download the archive with the **oc** file from: [oc.tar.gz](https://oc.tar.gz).
3. Unpack the downloaded archive. The **oc** executable file is unpacked in your current directory:

```
$ tar xf oc.tar.gz && ./oc version
```

4. Add the **oc** file to your path.

## 2.2. DOWNLOADING THE CODEREADY WORKSPACES DEPLOYMENT SCRIPT

This procedure describes how to obtain and unpack the archive with the CodeReady Workspaces deployment shell script.

The CodeReady Workspaces deployment script uses the OpenShift Operator to deploy Red Hat Single Sign-On, the PostgreSQL database, and the CodeReady Workspaces server container images on an instance of Red Hat OpenShift Container Platform. The images are available in the Red Hat Container Catalog.

### Procedure

1. Change to a temporary directory. Create it if necessary. For example:

```
$ mkdir ~/tmp
$ cd ~/tmp
```

2. Download the archive with the deployment script and the **custom-resource.yaml** file using the browser with which you logged into the [Red Hat Developer Portal](#): [codeready-workspaces-1.2.2.GA-operator-installer.tar.gz](#).
3. Unpack the downloaded archive and change to the created directory:

```
$ tar xvf codeready-workspaces-1.2.2.GA-operator-installer.tar.gz \
&& cd codeready-workspaces-operator-installer/
```

### Next steps

Continue by configuring and running the deployment script. See [Section 2.3, “Running the CodeReady Workspaces deployment script”](#).

## 2.3. RUNNING THE CODEREADY WORKSPACES DEPLOYMENT SCRIPT

The CodeReady Workspaces deployment script uses command-line arguments and the **custom-resource.yaml** file to populate a set of configuration environment variables for the OpenShift Operator used for the actual deployment.

### Prerequisites

- Downloaded and unpacked deployment script and the configuration file. See [Section 2.2, “Downloading the CodeReady Workspaces deployment script”](#).
- A running instance of Red Hat OpenShift Container Platform 3.11 or OpenShift Dedicated 3.11. To install OpenShift Container Platform, see the [Getting Started with OpenShift Container Platform](#) guide.
- The OpenShift Origin Client Tools 3.11, **oc**, is in the path. See [Section 2.1, “Downloading the Red Hat OpenShift Origin Client Tools”](#).
- The user is logged in to the OpenShift instance (using, for example, **oc login**).

- CodeReady Workspaces is supported for use with Google Chrome 70.0.3538.110 (Official Build) (64bit).
- **cluster-admin** rights to successfully deploy CodeReady Workspaces using the deploy script. The following table lists the objects and the required permissions:

Type of object	Name of the object that the installer creates	Description	Permission required
CRD	-	Custom Resource Definition - CheCluster	<b>cluster-admin</b>
CR	<b>codeready</b>	Custom Resource of the CheCluster type of object	<b>cluster-admin</b> . Alternatively, you can create a <b>clusterrole</b> .
ServiceAccount	<b>codeready-operator</b>	Operator uses this service account to reconcile CodeReady Workspaces objects	The edit role in a target namespace.
Role	<b>codeready-operator</b>	Scope of permissions for the operator-service account	<b>cluster-admin</b>
RoleBinding	<b>codeready-operator</b>	Assignment of a role to the service account	The edit role in a target namespace.
Deployment	<b>codeready-operator</b>	Deployment with operator image in the template specification	The edit role in a target namespace.
ClusterRole	<b>codeready-operator</b>	<b>ClusterRole</b> allows you to create, update, delete oAuthClients	<b>cluster-admin</b>
ClusterRoleBinding	<b>_\${NAMESPACE}-codeready-operator</b>	<b>ClusterRoleBinding</b> allows you to create, update, delete oAuthClients	<b>cluster-admin</b>
Role	<b>secret-reader</b>	<b>Role</b> allows you to read secrets in the router namespace	<b>cluster-admin</b>
RoleBinding	<b>_\${NAMESPACE}-codeready-operator</b>	<b>RoleBinding</b> allows you to read secrets in router namespace	<b>cluster-admin</b>



By default, the operator-service account gets privileges to list, get, watch, create, update, and delete ingresses, routes, service accounts, roles, rolebindings, PVCs, deployments, configMaps, secrets. It also has privileges to run execs into pods, watch events, and read pod logs in a target namespace.

With self-signed certificates support enabled, the operator-service account gets privileges to read secrets in an OpenShift router namespace.

With OpenShift OAuth enabled, the operator-service account gets privileges to get, list, create, update, and delete oAuthclients at a cluster scope.

### 2.3.1. Deploying CodeReady Workspaces with default settings

1. Run the following command:

```
$ ./deploy.sh --deploy
```



#### NOTE

Run the `./deploy.sh --help` command to get a list of all available arguments. For a description of all the options, see [Section 2.6, "CodeReady Workspaces deployment script parameters"](#).

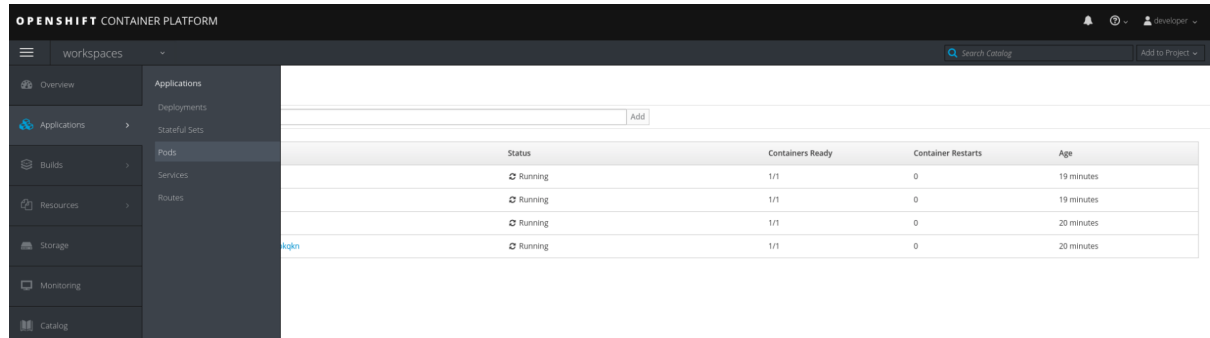
The following messages indicates that CodeReady Workspaces is getting installed:

```
[INFO]: Welcome to CodeReady Workspaces Installer
[INFO]: Found oc client in PATH
[INFO]: Checking if you are currently logged in...
[INFO]: Active session found. Your current context is: workspaces/192-168-42-231:8443/admin
[INFO]: Creating operator service account
[INFO]: Service account already exists
[INFO]: Create service account roles
[INFO]: Role Binding already exists
[INFO]: Self-signed certificate support enabled
[INFO]: Adding extra privileges for an operator service account
[INFO]: Creating secret-reader role and rolebinding in namespace default
[INFO]: Role secret-reader already exists
[INFO]: Creating role binding to let operator get secrets in namespace default
[INFO]: Role binding codeready-operator already exists in namespace default
[INFO]: Creating custom resource definition
[INFO]: Creating Operator Deployment
[INFO]: Existing operator deployment found. It will be deleted
[INFO]: Waiting for the Operator deployment to be scaled to 1. Timeout 5 minutes
[INFO]: Codeready Workspaces operator successfully deployed
[INFO]: Creating Custom resource. This will initiate CodeReady Workspaces deployment
[INFO]: CodeReady is going to be deployed with the following settings:
[INFO]: TLS support: false
[INFO]: OpenShift oAuth: false
[INFO]: Self-signed certs: true
[INFO]: Waiting for CodeReady Workspaces to boot. Timeout: 1200 seconds
[INFO]: CodeReady Workspaces successfully deployed and is available at http://codeready-workspaces.192.168.42.231.nip.io
```

The **CodeReady Workspaces successfully deployed and available at<URL>** message confirms that the deployment is successful.

2. Open the OpenShift web console.
3. In the **My Projects** pane, click **workspaces**.
4. Click **Applications > Pods**. The pods are shown running.

Figure 2.1. Pods for codeready shown running



### 2.3.2. Deploying CodeReady Workspaces with a self-signed certificate and OpenShift OAuth

To deploy CodeReady Workspaces with a self-signed certificate, run the following command:

```
$. /deploy.sh --deploy --oauth
```



#### NOTE

If you use the TLS mode with a self-signed certificate, ensure that your browser trusts the certificate. If it does not trust the certificate, the **Authorization token is missed** error is displayed on the login page and the running workspace may not work as intended.

### 2.3.3. Deploying CodeReady Workspaces with a public certificate

To deploy CodeReady Workspaces to a cluster configured with public certificates, run the following command:

```
$. /deploy.sh --deploy --public-certs
```

### 2.3.4. Deploying CodeReady Workspaces with external Red Hat Single Sign-On

To deploy with an external Red Hat Single Sign-On (Red Hat SSO) and enable a Red Hat SSO instance, take the following steps:

1. Update the following values in the **custom-resource.yaml** file:

```
auth:
  externalIdentityProvider: 'true'
  identityProviderURL: 'https://my-red-hat-sso.com'
  identityProviderRealm: 'myrealm'
  identityProviderClientId: 'myClient'
```

- 1 Instructs the operator on whether or not to deploy the Red Hat SSO instance. When set to **true**, provisions the connection details.
- 2 Retrieved from the respective route or ingress unless explicitly specified in CR (when the **externalIdentityProvider** variable is **true**).
- 3 Name of a Red Hat SSO realm. This realm is created when the **externalIdentityProvider** variable is **true**. Otherwise, it is passed to the CodeReady Workspaces server.
- 4 The ID of a Red Hat SSO client. This client is created when the **externalIdentityProvider** variable is **false**. Otherwise, it is passed to the CodeReady Workspaces server.

2. Run the deploy script:

```
$ ./deploy.sh --deploy
```

### 2.3.5. Deploying CodeReady Workspaces with external Red Hat SSO and PostgreSQL

The deploy script supports the following combinations of external Red Hat SSO and PostgreSQL:

- PostgreSQL and Red Hat SSO
- Red Hat SSO only

The deploy script does not support the external database and bundled Red Hat SSO combination currently. Provisioning of the database and the Red Hat SSO realm with the client happens only with bundled resources. If you are connecting your own database or Red Hat SSO, you should pre-create resources.

To deploy with the external PostgreSQL database and Red Hat SSO, take the following steps:

1. Update the following PostgreSQL database-related values in the **custom-resource.yaml** file:

```
database:
  externalDb:      'true'
  chePostgresHostname: 'http://postgres'
  chePostgresPort:  '5432'
  chePostgresUser:  'myuser'
  chePostgresPassword: 'mypass'
  chePostgresDb:    'mydb'
```

- 1 When set to **true** the operator skips deploying PostgreSQL and passes the connection details of the existing database to the CodeReady Workspaces server. Otherwise, a PostgreSQL deployment is created.
- 2 The PostgreSQL database hostname that the CodeReady Workspaces server uses to connect to. Defaults to **postgres**.
- 3 The PostgreSQL database port that the CodeReady Workspaces server uses to connect to. Defaults to **5432**.
- 4 The PostgreSQL user that the CodeReady Workspaces server when making a database connection. Defaults to **pgche**.

- 5 The password of a PostgreSQL user. Auto-generated when left blank.
- 6 The PostgreSQL database name that the CodeReady Workspaces server uses to connect to. Defaults to **dbche**.

2. Update the following Red Hat SSO-related values in the **custom-resource.yaml** file:

```
auth:
  externalIdentityProvider: 'true'
  identityProviderURL: 'https://my-red-hat-sso.com'
  identityProviderRealm: 'myrealm'
  identityProviderClientId: 'myClient'
```

- 1 Instructs the operator on whether or not to deploy Red Hat SSO instance. When set to **true**, provisions the connection details.
- 2 Retrieved from the respective route or ingress unless explicitly specified in CodeReady Workspaces (when **externalIdentityProvider** is **true**).
- 3 Name of a Red Hat SSO realm. This realm is created when **externalIdentityProvider** is **true**. Otherwise, passed to the CodeReady Workspaces server.
- 4 ID of a Red Hat SSO client. This client is created when **externalIdentityProvider** is **false**. Otherwise, passed to the CodeReady Workspaces server.

3. Run the deploy script:

```
$ ./deploy.sh --deploy
```

### Additional resources

- See [Section 2.6, "CodeReady Workspaces deployment script parameters"](#) for definitions of the deployment script parameters.

## 2.4. VIEWING CODEREADY WORKSPACES INSTALLATION LOGS

You can view the installation logs in the terminal or from the OpenShift console.

### 2.4.1. Viewing CodeReady Workspaces installation logs in the terminal

To view the installation logs on the terminal, take the following steps:

1. To obtain the names of the pods you must switch to project where CodeReady Workspaces is installed:

```
$ oc get pods -n=<OpenShift-project-name>
```

Following is an example output.

```
NAME                                READY STATUS RESTARTS AGE
codeready-76d985c5d8-4sqmm          1/1   Running 2    1d
codeready-operator-54b58f8ff7-fc88p 1/1   Running 3    1d
```

```
keycloak-7668cdb5f5-ss29s    1/1  Running 2    1d
postgres-7d94b544dc-nmhwp    1/1  Running 1    1d
```

- To view the logs for the pod, run:

```
$ oc logs <log-name>
```

The following is an example output:

```
Deployment of web application archive
[/home/jboss/codeready/tomcat/webapps/dashboard.war] has finished in [286] ms
2019-05-13 12:47:49,201[ost-startStop-1] [INFO ] [o.a.c.startup.HostConfig 957] -
Deploying web application archive [/home/jboss/codeready/tomcat/webapps/swagger.war]
2019-05-13 12:47:49,318[ost-startStop-1] [INFO ] [o.a.c.startup.HostConfig 1020] -
Deployment of web application archive
[/home/jboss/codeready/tomcat/webapps/swagger.war] has finished in [117] ms
2019-05-13 12:47:49,320[ost-startStop-1] [INFO ] [o.a.c.startup.HostConfig 957] -
Deploying web application archive [/home/jboss/codeready/tomcat/webapps/workspace-
loader.war]
2019-05-13 12:47:49,397[ost-startStop-1] [INFO ] [o.a.c.startup.HostConfig 1020] -
Deployment of web application archive [/home/jboss/codeready/tomcat/webapps/workspace-
loader.war] has finished in [77] ms
2019-05-13 12:47:49,403[main] [INFO ] [o.a.c.http11.Http11NioProtocol 588] -
Starting ProtocolHandler ["http-nio-8080"]
2019-05-13 12:47:49,419[main] [INFO ] [o.a.catalina.startup.Catalina 700] - Server
startup in 31229 ms
```

## 2.4.2. Viewing CodeReady Workspaces installation logs in the OpenShift console

To view installation logs in OpenShift console, take the following steps:

- Navigate to the OpenShift web console `.
- In the **My Projects** pane, click **workspaces**.
- Click **Applications** > **Pods**. Click the name of the pod for which you want to view the logs.
- Click **Logs** and click **Follow**.

```
time="2019-05-15T08:50:04Z" level=info msg="Deployment 'codeready' successfully scaled
to 1"
time="2019-05-15T08:50:04Z" level=info msg="Updating codeready CR with Keycloak URL
status: http://keycloak-workspaces.192.168.42.231.nip.io "
time="2019-05-15T08:50:05Z" level=info msg="Custom resource codeready updated"
time="2019-05-15T08:50:05Z" level=info msg="Updating codeready CR with status:
CodeReady Workspaces server: Available"
time="2019-05-15T08:50:05Z" level=info msg="Custom resource codeready updated"
time="2019-05-15T08:50:05Z" level=info msg="Updating codeready CR with CodeReady
Workspaces server URL: http://codeready-workspaces.192.168.42.231.nip.io "
time="2019-05-15T08:50:05Z" level=info msg="Custom resource codeready updated"
time="2019-05-15T08:50:05Z" level=info msg="CodeReady Workspaces is now available at:
http://codeready-workspaces.192.168.42.231.nip.io "
```

## 2.5. CONFIGURING CODEREADY WORKSPACES TO WORK BEHIND A PROXY SERVER

This procedure describes how to configure CodeReady Workspaces for use in a deployment behind a proxy server. To access external resources (for example, to download Maven artifacts to build Java projects), change the workspace configuration.

### Prerequisites

- OpenShift with a logged in **oc** client.
- Deployment script. See [Section 2.2, “Downloading the CodeReady Workspaces deployment script”](#).


### Procedure

1. Update the following values in the **custom-resource.yaml** file:

```
apiVersion: org.eclipse.che/v1
kind: CheCluster
metadata:
  name: codeready
spec:
  server:
    cheFlavor: codeready
    cheImage: ${SERVER_IMAGE_NAME}
    cheImageTag: ${SERVER_IMAGE_TAG}
    tlsSupport: ${TLS_SUPPORT}
    selfSignedCert: ${SELF_SIGNED_CERT}
    proxyURL: 'http://172.19.20.128'
    proxyPort: '3128'
    nonProxyHosts: 'localhost|172.30.0.1|*.172.19.20.240.nip.io'
    proxyUser: ""
    proxyPassword: ""
```

1. Substitute **http://172.19.20.128** for the protocol and hostname of your proxy server.
2. Substitute **3128** for the port of your proxy server.
3. Substitute **172.30.0.1** for the value of the **\$KUBERNETES\_SERVICE\_HOST** environment variable (run **echo \$KUBERNETES\_SERVICE\_HOST** in any container running in the cluster to obtain this value).

You may also have to add a custom **nonProxyHosts** value as required by your network. In this example, this value is **\*.172.19.20.240.nip.io** (the routing suffix of the OpenShift Container Platform installation).



## IMPORTANT

- Use correct indentation as shown above.
- Use the bar sign (|) as the delimiter for multiple **nonProxyHosts** values.
- You may need to list the same wildcard and abbreviated **nonProxyHosts** values more than once. For example:

```
nonProxyHosts: 'localhost | 127.0.0.1 | *.nip.io | .nip.io |
*.example.com | .example.com'
```

2. Run the following command:

```
$ ./deploy.sh --deploy
```

## Additional resources

- [Section 2.3, “Running the CodeReady Workspaces deployment script”](#)
- [Chapter 4, \*Installing CodeReady Workspaces in restricted environments\*](#)

## 2.6. CODEREADY WORKSPACES DEPLOYMENT SCRIPT PARAMETERS

The **custom-resource.yaml** file contains default values for the installation parameters. Those parameters that take environment variables as values can be overridden from a command line. Not all installation parameters are available as flags.

Before running the deployment script in a fast mode, review the **custom-resource.yaml** file. Run the **./deploy.sh --help** command to get a list of all available arguments.

The following is an annotated example of the **custom-resource.yaml** file with all available parameters:

### Server settings:

```
server:
  cheFlavor:      'codeready'      1
  cheImage:      '${SERVER_IMAGE_NAME}' 2
  cheImageTag:   '${SERVER_IMAGE_TAG}' 3
  tlsSupport:    '${TLS_SUPPORT}'    4
  selfSignedCert: '${SELF_SIGNED_CERT}' 5
  proxyURL:      ""                6
  proxyPort:     ""                7
  nonProxyHosts: ""                8
  proxyUser:     ""                9
  proxyPassword: ""               10
```

- 1 Defaults to **che**. When set to **codeready**, CodeReady Workspaces is deployed. The difference is in images, labels, and in exec commands.
- 2 The server image used in the Che deployment.
- 3 The tag of an image used in the Che deployment.

- 4 TLS mode for Che. Ensure that you either have public certificate or set the **selfSignedCert** environment variable to **true**. If you use the TLS mode with a self-signed certificate, ensure that
- 5 When set to **true**, the operator attempts to get a secret in the OpenShift router namespace to add it to the **ava** trust store of the CodeReady Workspaces server. Requires cluster-administrator privileges for the operator service account.
- 6 The protocol and hostname of a proxy server. Automatically added as **JAVA\_OPTS** variable and **https(s)\_proxy** to the CodeReady Workspaces server and workspaces containers.
- 7 The port of a proxy server.
- 8 A list of non-proxy hosts. Use | as a delimiter. Example: **localhost|my.host.com|123.42.12.32**.
- 9 The username for a proxy server.
- 10 The password for a proxy user.

### Storage settings:

```
storage:
  pvcStrategy: 'common' 1
  pvcClaimSize: '1Gi' 2
```

- 1 The persistent volume claim strategy for the CodeReady Workspaces server. Can be common (all workspaces PVCs in one volume), per-workspace (one PVC per workspace for all the declared volumes), or unique (one PVC per declared volume). Defaults to **common**.
- 2 The size of a persistent volume claim for workspaces. Defaults to **1Gi**.

### Database settings:

```
database:
  externalDb: 'false' 1
  chePostgresHostName: "" 2
  chePostgresPort: "" 3
  chePostgresUser: "" 4
  chePostgresPassword: "" 5
  chePostgresDb: "" 6
```

- 1 When set to **true**, the operator skips deploying PostgreSQL and passes the connection details of the existing database to the CodeReady Workspaces server. Otherwise, a PostgreSQL deployment is created.
- 2 The PostgreSQL database hostname that the CodeReady Workspaces server uses to connect to. Defaults to **postgres**.
- 3 The PostgreSQL database port that the CodeReady Workspaces server uses to connect to. Defaults to **5432**.
- 4 The Postgres user that the CodeReady Workspaces server when making a database connection. Defaults to **pgche**.



- 5 The password of a PostgreSQL user. Auto-generated when left blank.
- 6 The PostgreSQL database name that the CodeReady Workspaces server uses to connect to. Defaults to **dbche**.

#### auth settings:

```

auth:
  openShifttoAuth:    '${ENABLE_OPENSIFT_OAUTH}'      1
  externalIdentityProvider:  'false'                2
  identityProviderAdminUserName: 'admin'            3
  identityProviderPassword:  'admin'                4
  identityProviderURL:      ""                       5
  identityProviderRealm:    ""                       6
  identityProviderClientId: ""                       7

```

- 1 Instructs an operator to enable the OpenShift v3 identity provider in Red Hat SSO and create the respective oAuthClient and configure the Che configMap accordingly.
- 2 Instructs the operator on whether or not to deploy the RH SSO instance. When set to **true**, it provisions the connection details.
- 3 The desired administrator username of the Red Hat SSO administrator (applicable only when the **externalIdentityProvider** variable is **false**).
- 4 The desired password of the Red Hat SSO administrator (applicable only when the **externalIdentityProvider** variable is **false**).
- 5 Retrieved from the respective route or ingress unless explicitly specified in CR (when the **externalIdentityProvider** variable is **true**).
- 6 The name of a Red Hat SSO realm. This realm is created when the **externalIdentityProvider** variable is **true**. Otherwise, it is passed to the CodeReady Workspaces server.
- 7 The ID of a Red Hat SSO client. This client is created when the **externalIdentityProvider** variable is **false**. Otherwise, it is passed to the CodeReady Workspaces server.

## CHAPTER 3. INSTALLING CODEREADY WORKSPACES ON OPENSIFT 4 FROM OPERATORHUB

CodeReady Workspaces is now compatible with OpenShift 4.1 and has a dedicated operator compatible with the *Operator Lifecycle Manager* (OLM), which allows for easy installation and automated updates.



### NOTE

OLM is a management framework for extending Kubernetes with Operators. The OLM project is a component of the [Operator Framework](#), which is an open-source toolkit to manage Kubernetes-native applications, called operators, in an effective, automated, and scalable way.

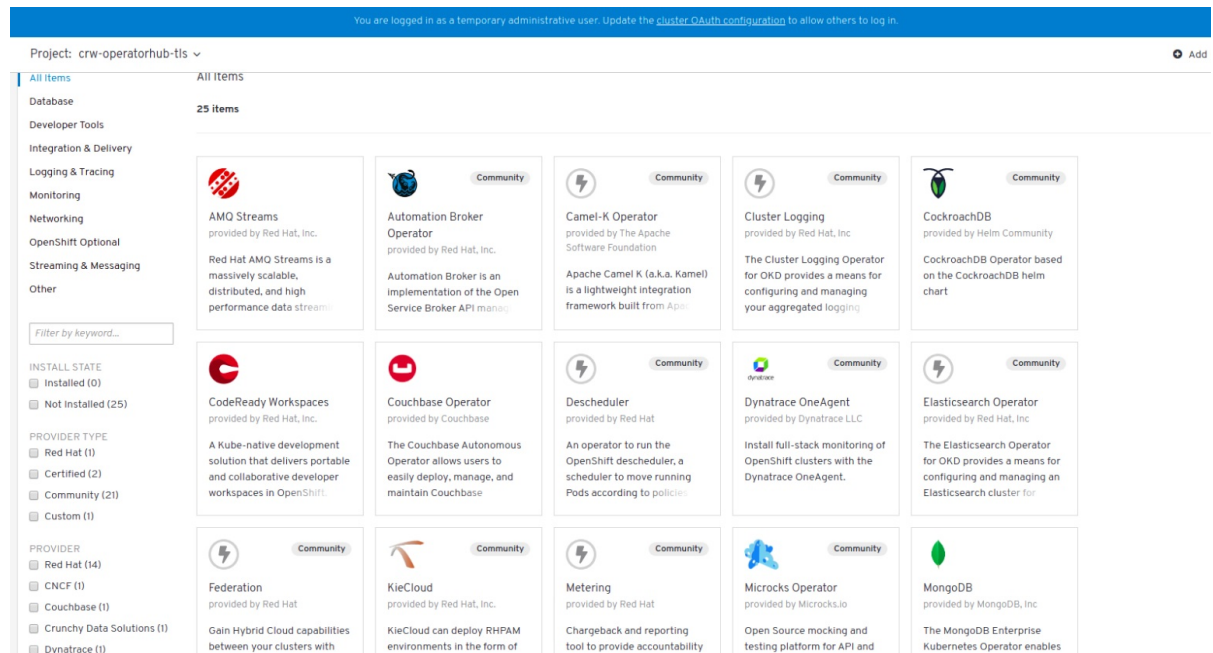
When using CodeReady Workspaces on OpenShift 4, updates are performed explicitly via OperatorHub. This is different from CodeReady Workspaces on OpenShift 3.11, where updates are performed using the `migrate*.sh` script or manual steps.

### Procedure

To install CodeReady Workspaces 1.2 from OperatorHub:

1. Launch the OpenShift Web Console.
2. In the console, click **CodeReady Workspaces** in the **OperatorHub** tab. The **CodeReady Workspaces 1.2** window is shown.

Figure 3.1. CodeReady Workspaces 1.2 listed on OperatorHub



3. Click **Install**.

Figure 3.2. Install button on the CodeReady Workspaces 1.2 window

**Install**

A collaborative Kubernetes-native development solution that delivers OpenShift workspaces and in-browser IDE for rapid cloud application development. This operator installs PostgreSQL, Red Hat SSO, and the CodeReady Workspaces server, as well as configures all three services.

**OPERATOR VERSION**  
1.1.0

**PROVIDER TYPE**  
Custom

**PROVIDER**  
Red Hat, Inc.

**REPOSITORY**  
N/A

**CONTAINER IMAGE**  
quay.io/crw/operator-container:1.1-25

**CREATED AT**  
2019-04-02 16:23:25

**SUPPORT**  
Red Hat, Inc.

## Pre-Reqs

In addition to a standard namespaced role, this operator may require a **ClusterRole** that allows the operator service account to:

- read secrets in openshift-ingress namespace (the operator will get TLS secret in openshift-ingress namespace to extract certificate and add it to Java trust store for CodeReady Workspaces server)
- list, get, create, update, patch, watch and delete oAuthclients at a cluster scope (the operator creates oAuthclient and OpenShift v3 Identity provider in Keycloak to enable Login With OpenShift in CodeReady Workspaces)

The operator service account will require extra privileges if you enable either `server.selfSignedCerts` or `auth.openShiftoAuth` which are false in CR template by default. After the operator is installed, grant the **codeready-operator** service account such privileges: When `auth.openShiftoAuth` is enabled:

```
oc create clusterrole codeready-operator --resource=oauthclients --verb=get,create,delete,update,list,watch
oc create clusterrolebinding codeready-operator --clusterrole=codeready-operator --serviceaccount=${NAMESPACE}:codeready-operator
```

When `server.selfSignedCerts` is enabled:

```
oc create role secret-reader --resource=secrets --verb=get -n=openshift-ingress
oc create rolebinding codeready-operator --role=secret-reader --serviceaccount=${NAMESPACE}:codeready-operator -n=openshift-ingress
```

`$(NAMESPACE)` is an OpenShift project where you installed the operator.

4. In the **Create Operator Subscription** window:
  - a. From the **A specific namespace in the cluster** drop-down list, select the **openshift-codeready-workspaces** namespace to install the Operator to.
  - b. Choose the appropriate approval strategy in the **Approval Strategy** field.
  - c. Click **Subscribe**.

Figure 3.3. Selections in the Create Operator Subscription window

You are logged in as a temporary administrative user. Update the [cluster OAuth configuration](#) to allow others to log in.

## Create Operator Subscription

Keep your service up to date by subscribing to a channel and update strategy from which to pull updates.

**Installation Mode \***

All namespaces on the cluster (default)  
This mode is not supported by this Operator

A specific namespace on the cluster  
Operator will be available in a single namespace only.

**Update Channel \***

final

**Approval Strategy \***

Automatic

Manual

A subscription is created in the *Operator Lifecycle Manager* (OLM), and the operator is installed in the chosen namespace. Successful installation implies that the following requirements in the *Cluster Service Version* (CSV) are created:

- Service account (SA)
- Role-based access control (RBAC)
- Deployment  
By default, a deployment creates only one instance of an application, scaled to 1. After the operator installation completes, the operator is marked as installed in the OperatorHub window.

5. Navigate to **Catalog > Installed Operators**. The CodeReady Workspaces operator with an **InstallSucceeded** message is displayed.




Figure 3.4. Installed Operators on OperatorHub

Project: crw-operatorhub-tls Add

### Installed Operators


Installed Operators are represented by Cluster Service Versions within this namespace. For more information, see the [Operator Lifecycle Manager documentation](#). Or create an Operator and Cluster Service Version using the [Operator SDK](#).

1 InstallSucceeded   1 of 1 Item

NAME	NAMESPACE	DEPLOYMENT	STATUS	PROVIDED APIS
 CodeReady Workspaces 1.1.0 provided by Red Hat, Inc.	 crw-operatorhub-tls	 codeready-operator	InstallSucceeded	CodeReady Workspaces Cluster

6. To deploy CodeReady Workspaces 1.2 on the operator:
  - a. In the selected namespace, create a custom resource Che Cluster that the operator will watch. To do so, click **Operator > CodeReady Workspaces Cluster**.
  - b. Click **Create Che Cluster**.

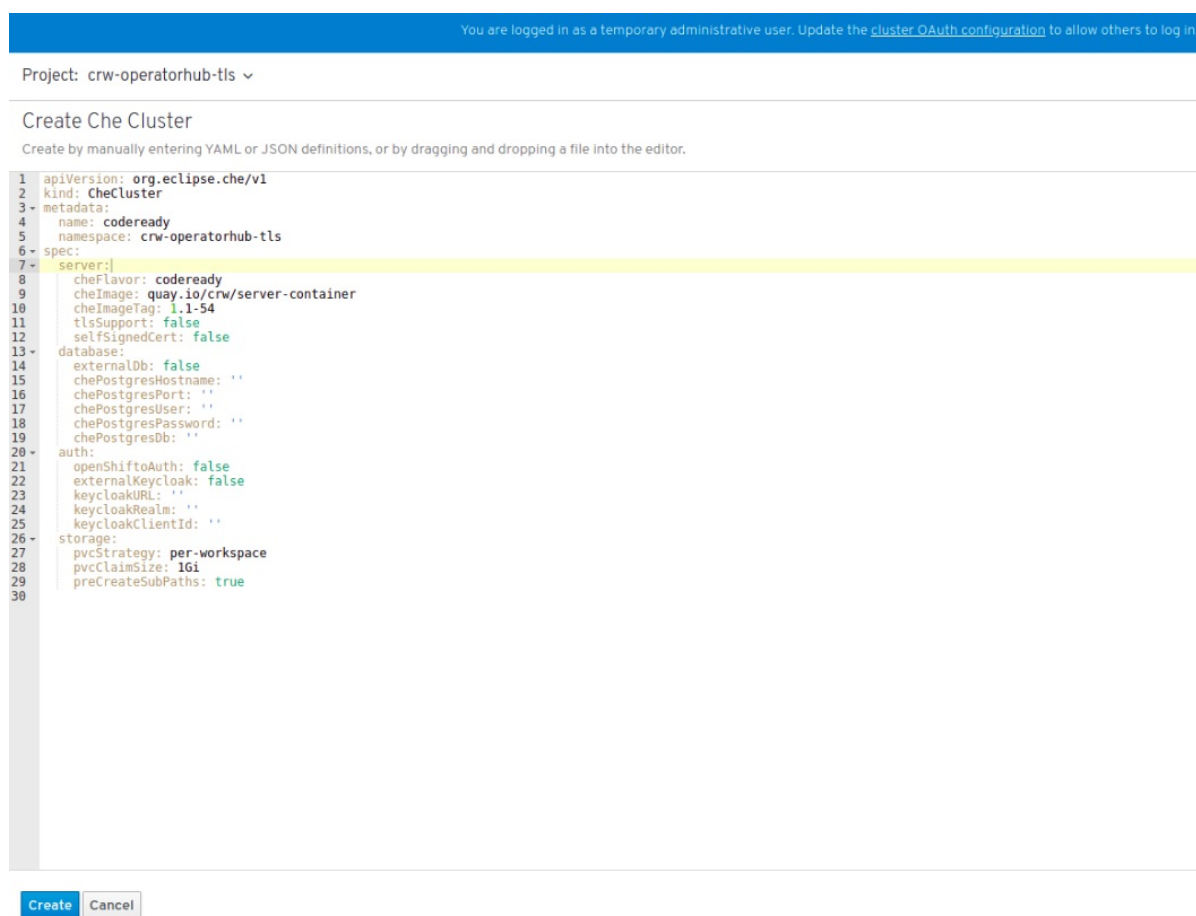
Figure 3.5. Clicking Create Che Cluster



A template for creating a Che-cluster type custom resource is available. When using OpenShift OAuth or a self-signed certificate, grant cluster administrator privileges to the codeready-operator service account. For instructions, see the operator description when you initially install the operator.

7. Click **Create**.

Figure 3.6. Clicking Create to create the Che cluster

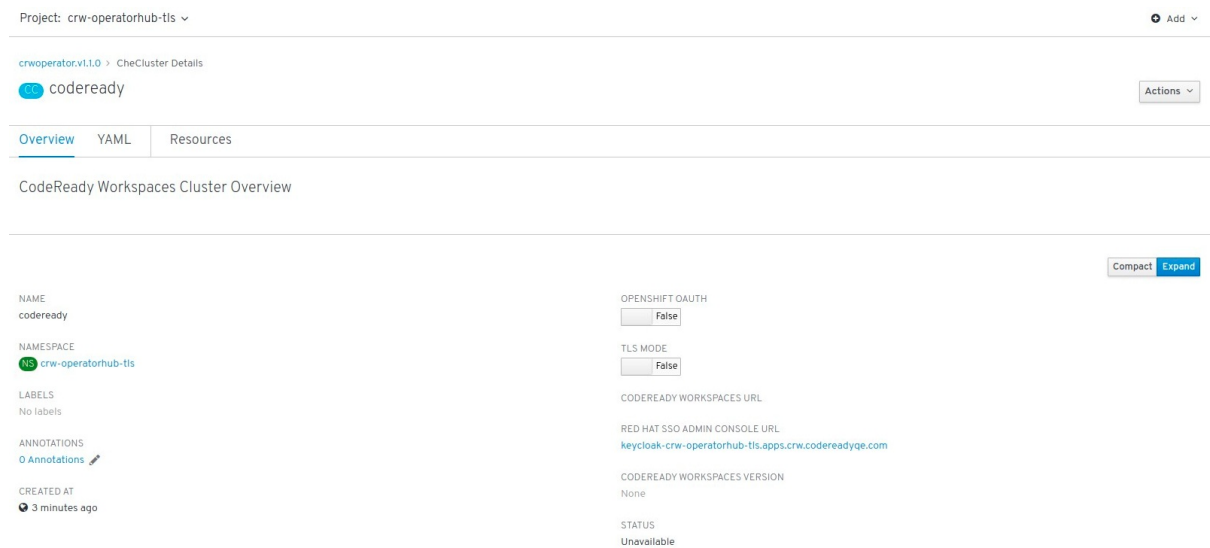


After the custom resource is created, the operator starts executing the following controller business logic:

- Creates the Kubernetes and OpenShift resources
- Provisions the database and Red Hat SSO resources
- Updates the resource status while the installation is in progress

8. To track the progress in the OLM UI, navigate to the resource details window.

Figure 3.7. Display of route for CodeReady Workspaces



9. Wait for the status to become **Available**. The CodeReady Workspaces route is displayed.
10. To track the installation progress, follow the operator logs. Notice that the codeready-operator pod is the same namespace. Follow the logs and wait until the operator updates the resource status to **Available** and sets the URLs.

Figure 3.8. CodeReady Operator pod shown running

```

103 lines
time="2019-04-04T15:57:07Z" level=info msg="waiting for deployment codeready, default timeout: 420 seconds"
time="2019-04-04T15:59:36Z" level=info msg="Deployment 'codeready' successfully scaled to 1"
time="2019-04-04T15:59:36Z" level=info msg="Updating codeready CR with Keycloak URL status: http://keycloak-crw-operhub.apps.crw.codereadyqe.com"
time="2019-04-04T15:59:36Z" level=info msg="Custom resource codeready updated"
time="2019-04-04T15:59:36Z" level=info msg="Updating codeready CR with status: CodeReady Workspaces server: Available"
time="2019-04-04T15:59:36Z" level=info msg="Custom resource codeready updated"
time="2019-04-04T15:59:36Z" level=info msg="Updating codeready CR with CodeReady Workspaces server URL: http://codeready-crw-operhub.apps.crw.codereadyqe.com"
time="2019-04-04T15:59:36Z" level=info msg="Custom resource codeready updated"
time="2019-04-04T15:59:36Z" level=info msg="CodeReady Workspaces is now available at: http://codeready-crw-operhub.apps.crw.codereadyqe.com"
time="2019-04-04T15:59:36Z" level=info msg="Updating codeready CR with version: 1.1-54"
time="2019-04-04T15:59:36Z" level=info msg="Custom resource codeready updated"
time="2019-04-04T16:01:47Z" level=info msg="Deployment codeready is in the rolling update state"
time="2019-04-04T16:01:47Z" level=info msg="Updating codeready CR with status: Available: Rolling update in progress"
time="2019-04-04T16:01:47Z" level=info msg="Custom resource codeready updated"
time="2019-04-04T16:01:47Z" level=info msg="Waiting for a successful rolling update of deployment codeready. Default timeout: 420 seconds"
time="2019-04-04T16:02:33Z" level=info msg="Rolling update of 'codeready' deployment finished"
time="2019-04-04T16:02:33Z" level=info msg="Updating codeready CR with Keycloak URL status: http://keycloak-crw-operhub.apps.crw.codereadyqe.com"
time="2019-04-04T16:02:33Z" level=info msg="Custom resource codeready updated"
time="2019-04-04T16:02:33Z" level=info msg="Updating codeready CR with status: CodeReady Workspaces server: Available"
time="2019-04-04T16:02:33Z" level=info msg="Custom resource codeready updated"

```

## Additional resources

- [CodeReady Workspaces scalability workaround on OCP 4.1/AWS](#)

## CHAPTER 4. INSTALLING CODEREADY WORKSPACES IN RESTRICTED ENVIRONMENTS

To install CodeReady Workspaces in a restricted environment that has no direct connection to the Internet available, all requisite container images must be downloaded prior to the deployment of the application and then used from a local container registry.

The installation of CodeReady Workspaces consists of the following three deployments:

- PostgreSQL (database)
- Red Hat Single Sign-On (Red Hat SSO)
- CodeReady Workspaces

Each of these deployments uses a container image in the container specification. Normally, the images are pulled from the Red Hat Container Catalog at [Red Hat CodeReady Workspaces](#). When the OpenShift Container Platform cluster used for the CodeReady Workspaces deployment does not have access to the Internet, the installation fails. To allow the installation to proceed, override the default image specification to point to your local registry.

### Prerequisites

- The CodeReady Workspaces deployment script. See [Section 2.2, “Downloading the CodeReady Workspaces deployment script”](#) for detailed instructions on how to obtain the script.
- A local (internal) container registry that can be accessed by the OpenShift instance where CodeReady Workspaces is to be deployed.
- Container images required for CodeReady Workspaces deployment imported to the local registry.
- The following images downloaded from [registry.redhat.io](#):

See [Section 4.1, “Preparing CodeReady Workspaces deployment from a local registry”](#) for detailed instructions.

## 4.1. PREPARING CODEREADY WORKSPACES DEPLOYMENT FROM A LOCAL REGISTRY

To install CodeReady Workspaces in a restricted environment without access to the Internet, the product deployment container images need to be imported from an external container registry into a local (internal) registry.

CodeReady Workspaces deployment requires the following images from the Red Hat Container Catalog at [registry.redhat.io](#):

- **codeready-workspaces/server-rhel8:1.2**: CodeReady Workspaces server
- **codeready-workspaces/server-operator-rhel8:1.2**: operator that installs and manages CodeReady Workspaces
- **rhscsl/postgresql-96-rhel7:1-40**: PostgreSQL database for persisting data
- **redhat-ss0-7/ss073-openshift:1.0-11**: Red Hat SSO for authentication

- **ubi8-minimal:8.0-127**: utility image used in preparing the PVCs (Persistent Volume Claims)

## Prerequisites

To import container images (create image streams) in your OpenShift Container Platform cluster, you need:

- **cluster-admin** rights

## Procedure


1. Import the required images from an external registry to a local registry that your OpenShift Container Platform cluster can reach. See [Section 4.4, "Making CodeReady Workspaces images available from a local registry"](#) for example instructions on how to do this.
2. Edit the **custom-resource.yaml** configuration file to specify that the CodeReady Workspaces deployment script should use images from your internal registry. Add the following specification fields to the respective blocks. Use the address of your internal registry and the name of the OpenShift project into which you imported the images. For example:

```
spec:

  storage:
    pvcJobsImage: '172.0.0.30:5000/openshift/ubi8-minimal:8.0-127'

  database:
    postgresImage: '172.0.0.30:5000/openshift/rhsc/postgresql-96-rhel7:1-40'

  auth:
    identityProviderImage: '172.0.0.30:5000/openshift/redhat-sso-7/sso73-openshift:1.0-11'
```



### IMPORTANT

- Make sure to use correct indentation as shown above.
- Substitute **172.0.0.30:5000** for the actual address of your local registry.
- Substitute **openshift** for the name of the OpenShift project into which you imported the images.

## Additional resources

- See [Section 4.4, "Making CodeReady Workspaces images available from a local registry"](#) for example instructions on how to transfer container images required for CodeReady Workspaces deployment to a restricted environment.
- See [Section 2.6, "CodeReady Workspaces deployment script parameters"](#) for an overview of all available configuration options in the **custom-resource.yaml** configuration file.

## 4.2. RUNNING THE CODEREADY WORKSPACES DEPLOYMENT SCRIPT IN A RESTRICTED ENVIRONMENT

To deploy CodeReady Workspaces in a restricted environment with access to the Internet, it is necessary to use container images from a local (internal) registry. The deployment script (**deploy.sh**) allows to specify a custom image to be used for installation.



Specification fields from the **custom-resource.yaml** configuration file and arguments supplied to the **deploy.sh** script are passed to the operator. The operator then constructs the deployment of CodeReady Workspaces with the images in the container specification.

### Prerequisites

- Imported container images required by CodeReady Workspaces to a local registry.
- Local addresses of imported images specified in the **custom-resource.yaml** file.

See [Section 4.1, “Preparing CodeReady Workspaces deployment from a local registry”](#) for detailed instructions.

### Procedure

To deploy CodeReady Workspaces from an internal registry, run the **./deploy.sh --deploy** command and specify custom (locally available) server and operator images.

1. Use the **--server-image** and **--version** parameters to specify the server image and the **--operator-image** parameter to specify the operator image. For example:

```
$ ./deploy.sh --deploy \
  --server-image=172.0.0.30:5000/openshift/codeready-workspaces/server-rhel8 \
  --version=1.2 \
  --operator-image=172.0.0.30:5000/openshift/codeready-workspaces/server-operator-rhel8:1.2
```



### IMPORTANT

- Substitute **172.0.0.30:5000** for the actual address of your local registry.
- Substitute **openshift** for the name of the OpenShift project into which you imported the images.

### Additional resources

- See [Section 2.3, “Running the CodeReady Workspaces deployment script”](#) for instructions on how to run the **deploy.sh** script in other situations.

## 4.3. STARTING WORKSPACES IN RESTRICTED ENVIRONMENTS

Starting a workspace in CodeReady Workspaces implies creating a new deployment. Different stacks use different images. All of these stacks are from the Red Hat Container Catalog at [Red Hat CodeReady Workspaces](#). For more information on stacks, see the [Stacks](#) chapter in this guide.

It is not possible to override stack images during the installation of CodeReady Workspaces. You need to manually edit preconfigured stacks. See [Creating stacks](#).

### Procedure

To start a workspace in a restricted environment:

1. Import all of the following images you need to an internal registry:
  - **registry.redhat.io/codeready-workspaces/stacks-java-rhel8:1.2**

- [registry.redhat.io/codeready-workspaces/stacks-node-rhel8:1.2](https://registry.redhat.io/codeready-workspaces/stacks-node-rhel8:1.2)
- [registry.redhat.io/codeready-workspaces/stacks-node:1.2](https://registry.redhat.io/codeready-workspaces/stacks-node:1.2)
- [registry.redhat.io/codeready-workspaces/stacks-php-rhel8:1.2](https://registry.redhat.io/codeready-workspaces/stacks-php-rhel8:1.2)
- [registry.redhat.io/codeready-workspaces/stacks-python-rhel8:1.2](https://registry.redhat.io/codeready-workspaces/stacks-python-rhel8:1.2)
- [registry.redhat.io/codeready-workspaces/stacks-dotnet-rhel8:1.2](https://registry.redhat.io/codeready-workspaces/stacks-dotnet-rhel8:1.2)
- [registry.redhat.io/codeready-workspaces/stacks-golang-rhel8:1.2](https://registry.redhat.io/codeready-workspaces/stacks-golang-rhel8:1.2)
- [registry.redhat.io/codeready-workspaces/stacks-java-rhel8:1.2](https://registry.redhat.io/codeready-workspaces/stacks-java-rhel8:1.2)
- [registry.redhat.io/codeready-workspaces/stacks-cpp-rhel8:1.2](https://registry.redhat.io/codeready-workspaces/stacks-cpp-rhel8:1.2)  
See [Section 4.4, “Making CodeReady Workspaces images available from a local registry”](#) for example instructions on how to transfer container images required for CodeReady Workspaces deployment to a restricted environment.

2. Modify the preconfigured stacks:


- Log in to CodeReady Workspaces as an administrator. (The default login credentials are username: **admin** and password: **admin**.)
- In the left pane, click **Stacks**.
- From the list of stacks, select a stack to edit.

Figure 4.1. Selecting a stack to edit

NAME	DESCRIPTION	COMPONENTS	ACTIONS
.NET	.NET 2.0 stack with .NET Core SDK and Runt...	.NET SDK, .NET Runtime, RHEL	🗑️ 🔗
Cpp	C and C++ Developer Tools stack	llvm-toolset, g++, gcc, cmake, make, ...	🗑️ 🔗
Go	Golang stack	Golang, RHEL	🗑️ 🔗
Java 1.8	RHEL 7 Java stack with OpenJDK 1.8 and Ma...	EAP, Maven, OpenJDK, RHEL	🗑️ 🔗
Java EAP	Java stack with EAP 7.1, OpenJDK 1.8 and M...	EAP, Maven, OpenJDK, RHEL	🗑️ 🔗
Node	NodeJS 8x stack	NPM, NodeJS, RHEL	🗑️ 🔗
PHP	PHP 7.1 stack, including most popular exten...	Composer, PHP, RHEL	🗑️ 🔗
Python	Python 3.6 stack	PIP, Python, RHEL	🗑️ 🔗
Red Hat Fuse	Red Hat Fuse stack	Maven, OpenJDK, RHEL	🗑️ 🔗
RHEL 8 Beta - Java 1.8	RHEL 8 Beta Java stack with OpenJDK 1.8 an...	Maven, OpenJDK, RHEL	🗑️ 🔗
Spring Boot	Spring Boot stack	Maven, OpenJDK, RHEL	🗑️ 🔗
Vert.x	Vert.x stack	Maven, OpenJDK, RHEL	🗑️ 🔗

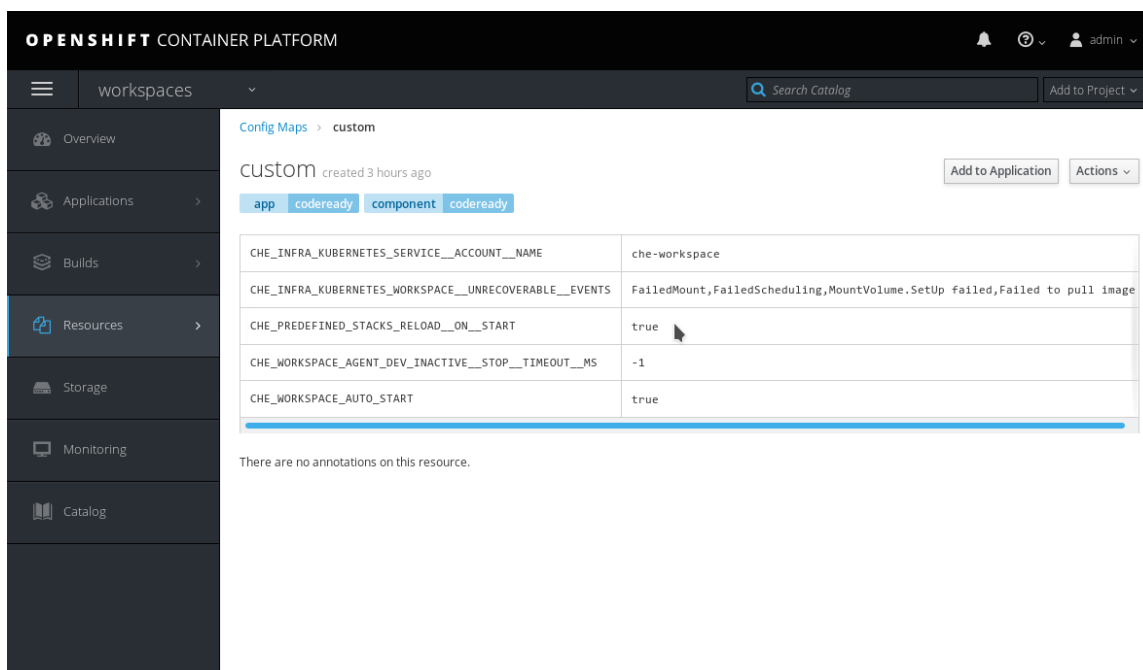
- Click and expand **DEV-MACHINE**.
- In the **Source** field, replace the default Red Hat Container Catalog image with a local image that you pushed to an internal registry.

Figure 4.2. Editing a stack

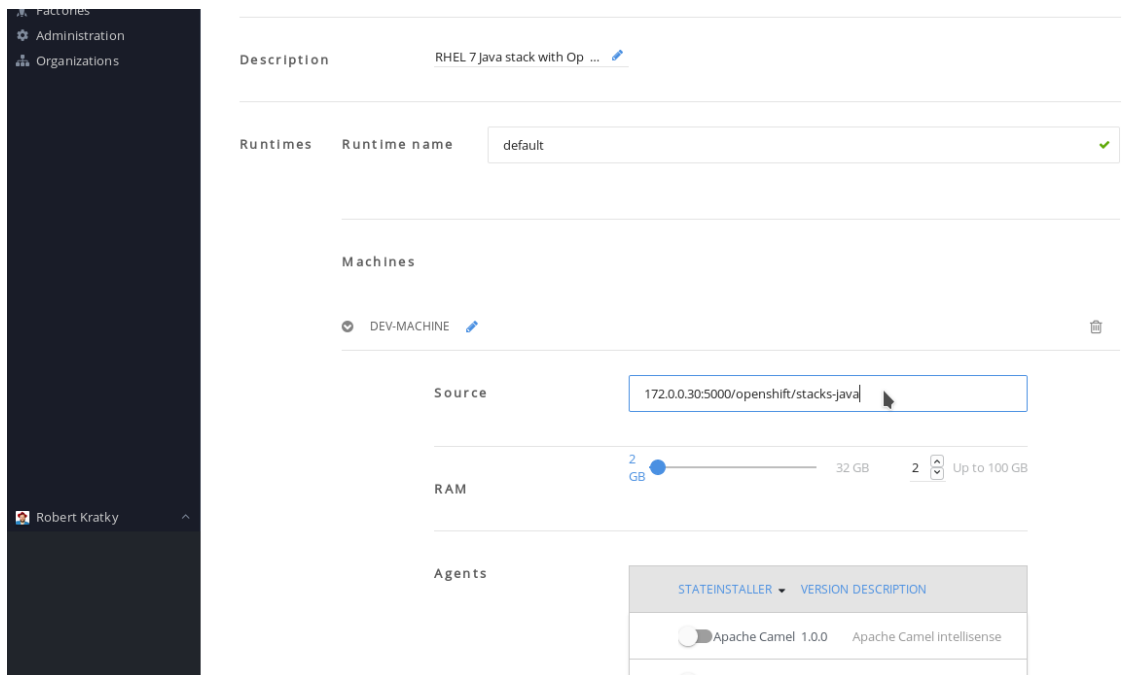


- f. Click **Save**.  
Repeat step 2 for all the stacks that you want to use.
3. By default, the CodeReady Workspaces server is configured to overwrite default stacks on application start. To ensure that the changes persist between restarts, edit the **custom** config map:
  - a. In the OpenShift Web Console, navigate to the **Resources > Config Maps** tab, and select the **custom** config map:

Figure 4.3. Navigate to the custom Config Map

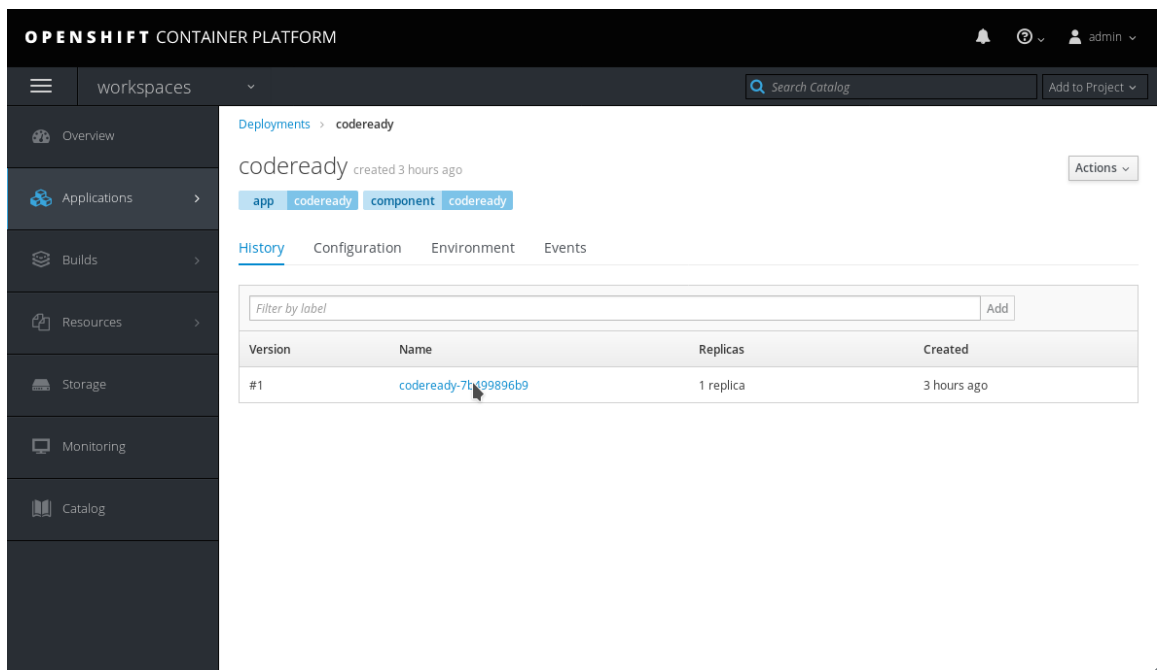


- b. From the **Actions** drop-down menu in the top-right corner, select **Edit**.
- c. Scroll to the **CHE\_PREDEFINED\_STACKS\_RELOAD\_\_ON\_\_START** key, and enter **false**:

Figure 4.4. Set the `CHE_PREDEFINED_STACKS_RELOAD_ON_START` key to false

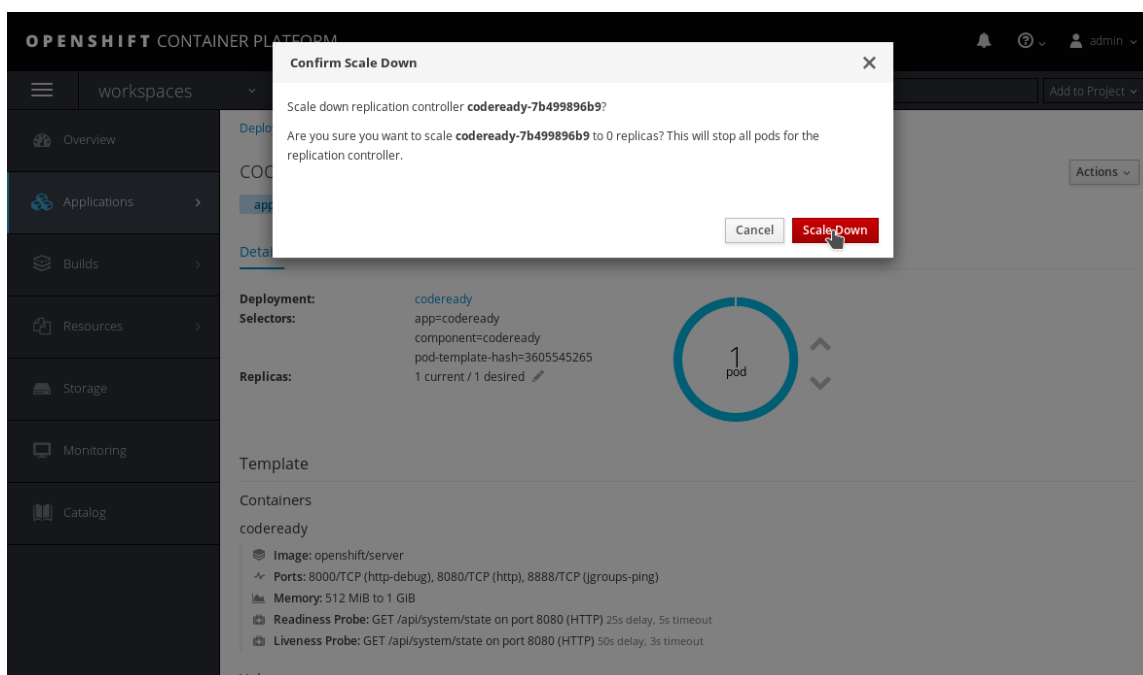
- d. Save the config map.
4. Forcibly start a new CodeReady Workspaces deployment, either by scaling the current deployment to **0** and then back to **1**, or deleting the CodeReady Workspaces server pod:
    - a. In the OpenShift Web Console, navigate to the **Applications > Deployments** tab, and select the **codeready** deployment:

Figure 4.5. Select the codeready deployment



- b. Click the name of the running **codeready** pod.
- c. Click the down arrow next to the number of pods, and confirm by clicking the **Scale Down** button:

Figure 4.6. Scale down the codeready deployment



- d. Scale the deployment back up by clicking the up arrow.

## 4.4. MAKING CODEREADY WORKSPACES IMAGES AVAILABLE FROM A LOCAL REGISTRY

This procedure is one of the possible ways for making CodeReady Workspaces container images that are required for deployment available from a local (internal) registry. First download the images using a machine connected to the Internet, then use some other method than network connection to transfer them to the restricted environment where you intend to deploy the application. Finally, upload (push) the images to the local registry that will be accessible during the deployment phase.

### Prerequisites

- Machine with access to the Internet that can be used to pull and save container images from **registry.redhat.io**.
- Means of transferring the downloaded images to the restricted environment where you intend to deploy CodeReady Workspaces.
- An account with the rights to push to a local registry within the restricted environment. The local registry must be accessible for the deployment.

At minimum, the following images need to be made locally available:


- **registry.redhat.io/codeready-workspaces/server-rhel8:1.2**
- **registry.redhat.io/codeready-workspaces/server-operator-rhel8:1.2**
- **registry.redhat.io/rhscpl/postgresql-96-rhel7:1-40**
- **registry.redhat.io/redhat-ss0-7/ss073-openshift:1.0-11**
- **registry.redhat.io/ubi8-minimal:8.0-127**

**NOTE**

You also need to follow the steps for image import for all stack images that you will want to use within CodeReady Workspaces.

**Procedure**

Use a tool for container management, such as Podman, to both download the container images and subsequently push them to a local registry within the restricted environment.

**NOTE**

The **podman** tool is available from the *podman* package in Red Hat Enterprise Linux starting with version 7.6. On earlier versions of Red Hat Enterprise Linux, the **docker** tool can be used with the same command syntax as suggested below for **podman**.

**Steps to perform on a machine with connection to the Internet**

1. Pull the required images from the **registry.redhat.io**. For example, for the **codeready-workspaces/server-rhel8** image, run:

```
# podman pull registry.redhat.io/codeready-workspaces/server-rhel8:1.2
```

Repeat the command for all the images you need.

2. Save all the pulled images to a **tar** file in your current directory on the disk. For example, for the **codeready-workspaces/server-rhel8** image, run:

```
# podman save --output codeready-server.tar \
registry.redhat.io/codeready-workspaces/server-rhel8:1.2
```

Repeat the command for all the images you need.

Transfer the saved **tar** image files to a machine connected to the restricted environment.

**Steps to perform on a machine within the restricted environment**

1. Load all the required images to the local container repository from which they can be uploaded to OpenShift. For example, for the **codeready-workspaces/server-rhel8** image, run:

```
# podman load --input codeready-server.tar
```

Repeat the command for all the images you need.

2. Optionally, check that the images have been successfully loaded to your local container repository. For example, to check for the **codeready** images, run:

```
# podman images */codeready/*
REPOSITORY                                TAG
registry.redhat.io/codeready-workspaces/server-rhel8  1.2
registry.redhat.io/codeready-workspaces/server-operator-rhel8  1.2
```

3. Log in to the instance of OpenShift Container Platform where you intend to deploy CodeReady Workspaces as a user with the **cluster-admin** role. For example, to log in as the user **admin**, run:

-

```
$ oc login --username <admin> --password <password>
```

In the above command, substitute **<admin>** for the username of the account you intend to use to deploy CodeReady Workspaces and **<password>** for the associated password.

4. Log in with **podman** to the local OpenShift registry that CodeReady Workspaces will be deployed from:

```
# podman login --username <admin> \
  --password $(oc whoami --show-token) 172.0.0.30:5000
```

In the above command:

- Substitute **<admin>** for the username of the account you used to log in to OpenShift.
  - Substitute **172.0.0.30:5000** for the URL of the local OpenShift registry to which you are logging in.
5. Tag the required images to prepare them to be pushed to the local OpenShift registry. For example, for the **codeready-workspaces/server-rhel8** image, run:

```
# podman tag registry.redhat.io/codeready-workspaces/server-rhel8:1.2 \
  172.0.0.30:5000/openshift/codeready-server:1.2
```

Repeat the command for all the images you need.

In the above command, substitute **172.0.0.30:5000** for the URL of the local OpenShift registry to which you will be pushing the images.

6. Optionally, check that the images have been successfully tagged for pushing to the local OpenShift registry. For example, to check for the **codeready** images, run:

```
# podman images */openshift/codeready-*
REPOSITORY                                TAG
172.0.0.30:5000/openshift/codeready-operator 1.2
172.0.0.30:5000/openshift/codeready-server  1.2
```

7. Push the required images to the local OpenShift registry. For example, for the **codeready-workspaces/server-rhel8** image, run:

```
# podman push 172.0.0.30:5000/openshift/codeready-server:1.2
```

Repeat the command for all the images you need.

In the above command, substitute **172.0.0.30:5000** for the URL of the local OpenShift registry to which you will be pushing the images.

8. Switch to the **openshift** project:

```
$ oc project openshift
```

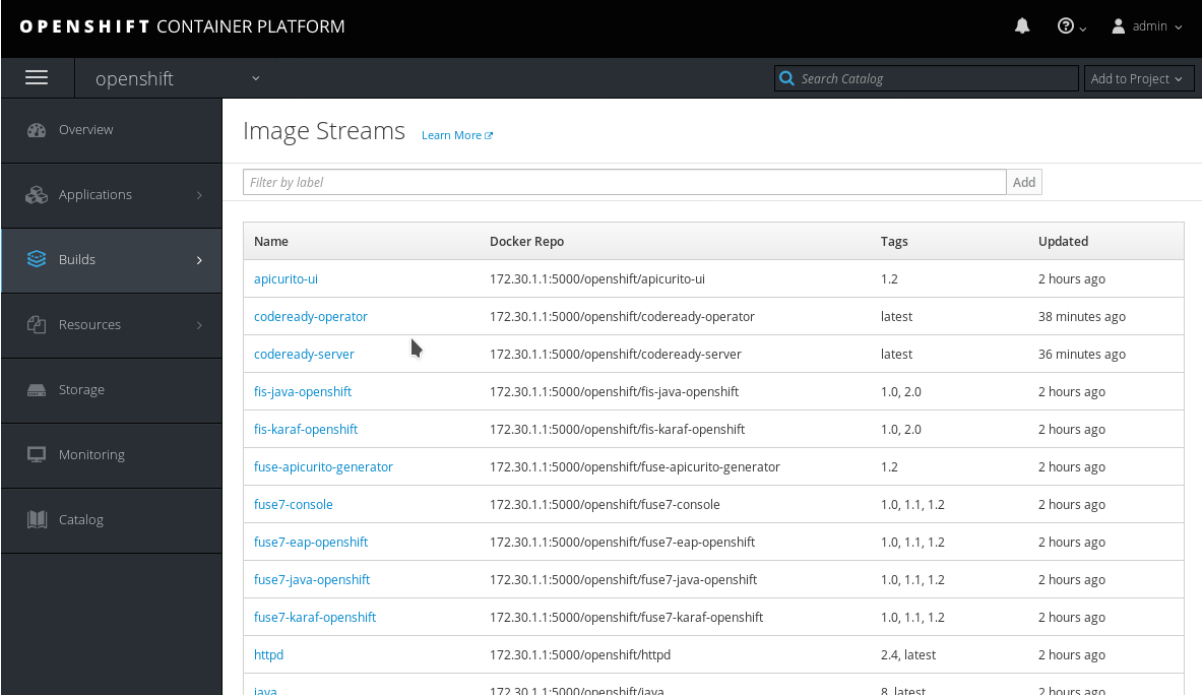
9. Optionally, check that the images have been successfully pushed to the local OpenShift registry. For example, to check for the **codeready** images, run:

```
$ oc get imagestream codeready-*
```

```
NAME                DOCKER REPO                                TAGS
codeready-operator  172.0.0.30:5000/openshift/codeready-operator  1.2
codeready-server    172.0.0.30:5000/openshift/codeready-server    1.2
```

You can also verify that the images have been successfully pushed in the OpenShift Console. Navigate to the **Builds > Images** tab, and look for image streams available in the **openshift** project:

Figure 4.7. Confirming images have been pushed to theopenshift project



Name	Docker Repo	Tags	Updated
apicurito-ui	172.30.1.1:5000/openshift/apicurito-ui	1.2	2 hours ago
codeready-operator	172.30.1.1:5000/openshift/codeready-operator	latest	38 minutes ago
codeready-server	172.30.1.1:5000/openshift/codeready-server	latest	36 minutes ago
fis-java-openshift	172.30.1.1:5000/openshift/fis-java-openshift	1.0, 2.0	2 hours ago
fis-karaf-openshift	172.30.1.1:5000/openshift/fis-karaf-openshift	1.0, 2.0	2 hours ago
fuse-apicurito-generator	172.30.1.1:5000/openshift/fuse-apicurito-generator	1.2	2 hours ago
fuse7-console	172.30.1.1:5000/openshift/fuse7-console	1.0, 1.1, 1.2	2 hours ago
fuse7-eap-openshift	172.30.1.1:5000/openshift/fuse7-eap-openshift	1.0, 1.1, 1.2	2 hours ago
fuse7-java-openshift	172.30.1.1:5000/openshift/fuse7-java-openshift	1.0, 1.1, 1.2	2 hours ago
fuse7-karaf-openshift	172.30.1.1:5000/openshift/fuse7-karaf-openshift	1.0, 1.1, 1.2	2 hours ago
httpd	172.30.1.1:5000/openshift/httpd	2.4, latest	2 hours ago
lava	172.30.1.1:5000/openshift/lava	8, latest	2 hours ago

The required CodeReady Workspaces container images are now available for use in your restricted environment.


### Additional resources

- Continue by [Section 4.2, “Running the CodeReady Workspaces deployment script in a restricted environment”](#)
- See [Importing Tag and Image Metadata](#) in the *OpenShift Container Platform Developer Guide* for detailed information about importing images.



## CHAPTER 5. INSTALLING CODEREADY WORKSPACES ON OPENSIFT DEDICATED

CodeReady Workspaces deployment requires **cluster-admin** OpenShift privileges. To install CodeReady Workspaces on OpenShift Dedicated cluster, a special procedure needs to be followed.



### IMPORTANT

- It is not possible to use OpenShift OAuth for CodeReady Workspaces authentication when deployed on OpenShift Dedicated using the method described in this section.
- If your OpenShift Dedicated cluster is compatible with 4.1, the installation of CodeReady Workspaces 1.2 will be handled from the Operator Hub. This installation will not require the **cluster-admin** rights. For detailed steps to install CodeReady Workspaces from the Operator Hub, see [Chapter 3, Installing CodeReady Workspaces on OpenShift 4 from OperatorHub](#).

### Prerequisites

- OpenShift Dedicated cluster with recommended resources. See [Chapter 2, Installing CodeReady Workspaces on OpenShift v3](#) for detailed information about cluster sizing.

### Procedure

To install CodeReady Workspaces on OpenShift Dedicated, you need to use the deployment script distributed with CodeReady Workspaces 1.0.2 and manually specify CodeReady Workspaces container images from version 1.2.

1. Download and unpack the deployment script from CodeReady Workspaces 1.0.2 from [developers.redhat.com](#). See [Section 2.2, “Downloading the CodeReady Workspaces deployment script”](#) for detailed instructions. Ensure that you are downloading the 1.0.2 version (direct link: [codeready-workspaces-1.0.2.GA-operator-installer.tar.gz](#)).
2. Run the deployment script and specify:
  - 1.2 **codeready-workspaces/server-rhel8** image
  - 1.0 **codeready-workspaces/server-operator** image

```

$ ./deploy.sh --deploy \
  --operator-image=registry.redhat.io/codeready-workspaces/server-operator:1.0 \
  --server-image=registry.redhat.io/codeready-workspaces/server-rhel8:1.2

```

See [Running the CodeReady Workspaces deployment script](#) for a description of other available parameters that can be used with the deployment script to customize the CodeReady Workspaces deployment.

### Additional resources

- [Section 2.6, “CodeReady Workspaces deployment script parameters”](#)
- [Chapter 4, Installing CodeReady Workspaces in restricted environments](#)

## CHAPTER 6. UPGRADING CODEREADY WORKSPACES

CodeReady Workspaces 1.2 introduces an operator that uses a controller to watch custom resources. This chapter lists steps to upgrade CodeReady Workspaces 1.1 to CodeReady Workspaces 1.2. For steps to upgrade from CodeReady Workspaces 1.0 to CodeReady Workspaces 1.1, see [https://access.redhat.com/documentation/en-us/red\\_hat\\_codeready\\_workspaces/1.1/html/administration\\_guide/upgrading-codeready-workspaces](https://access.redhat.com/documentation/en-us/red_hat_codeready_workspaces/1.1/html/administration_guide/upgrading-codeready-workspaces).

### 6.1. UPGRADING CODEREADY WORKSPACES USING A MIGRATION SCRIPT

The upgrade process using the **migrate\_\*.sh** scripts is an automation of the manual steps listed in the following [manual](#) section. The downloaded **codeready-workspaces-operator-installer** archive contains the following two **migrate\_\*.sh** scripts:

- **migrate\_1.0\_to\_1.1.sh**: Use to upgrade from CodeReady Workspaces 1.0 to CodeReady Workspaces 1.1
- **migrate\_1.1\_to\_1.2.sh**: Use to upgrade from CodeReady Workspaces 1.1 to CodeReady Workspaces 1.2

#### Prerequisites

- You must have a running instance of Red Hat OpenShift Container Platform 3.11 or OpenShift Dedicated 3.11.
- The OpenShift Origin Client Tools 3.11, **oc**, is in the path. See [Section 2.1, “Downloading the Red Hat OpenShift Origin Client Tools”](#).
- You must be logged in to the OpenShift instance using the following command:


```
$ oc login https://<your-IP-address>:8443 --username <your-username> --password <your-password>
```

- You must have obtained the Red Hat login credentials. For details, see <https://access.redhat.com/RegistryAuthentication#getting-a-red-hat-login-2>.

#### Procedure

1. To upgrade from CodeReady Workspaces 1.1 to CodeReady Workspaces 1.2, run the following command (the default project namespace is **workspaces**):

```
$ ./migrate_1.1_to_1.2.sh -n=<OpenShift-project-name>
```



#### IMPORTANT

Substitute *<OpenShift-project-name>* for the project namespace that you are using.

### 6.2. RUNNING THE CODEREADY WORKSPACES MIGRATION SCRIPT IN RESTRICTED ENVIRONMENTS

The upgrade process using the **migrate\_\*.sh** scripts is an automation of the manual steps listed in the following [manual](#) section. The downloaded **codeready-workspaces-operator-installer** archive contains the following two **migrate\_\*.sh** scripts:

- **migrate\_1.0\_to\_1.1.sh**: Use to upgrade from CodeReady Workspaces 1.0 to CodeReady Workspaces 1.1
- **migrate\_1.1\_to\_1.2.sh**: Use to upgrade from CodeReady Workspaces 1.1 to CodeReady Workspaces 1.2

### Prerequisites

- A local (internal) container registry that can be accessed by the OpenShift instance where CodeReady Workspaces is to be deployed.
- The following images downloaded from [registry.redhat.io](https://registry.redhat.io):
  - **registry.redhat.io/codeready-workspaces/server-rhel8:1.2**
  - **registry.redhat.io/codeready-workspaces/server-operator-rhel8:1.2**
  - **registry.redhat.io/rhscl/postgresql-96-rhel7:1-40**
  - **registry.redhat.io/redhat-sso-7/sso73-openshift:1.0-11**
  - **registry.redhat.io/ubi8-minimal:8.0-127**  
Download any of the following container images, with pre-configured stacks for creating workspaces, as required. Note that for most of the images the container name now contains a **-rhel8** suffix, but the Node.js 8 container continues to be based on RHEL 7.
  - **registry.redhat.io/codeready-workspaces/stacks-java-rhel8:1.2**
  - **registry.redhat.io/codeready-workspaces/stacks-node-rhel8:1.2**
  - **registry.redhat.io/codeready-workspaces/stacks-php-rhel8:1.2**
  - **registry.redhat.io/codeready-workspaces/stacks-python-rhel8:1.2**
  - **registry.redhat.io/codeready-workspaces/stacks-dotnet-rhel8:1.2**
  - **registry.redhat.io/codeready-workspaces/stacks-golang-rhel8:1.2**
  - **registry.redhat.io/codeready-workspaces/stacks-java-rhel8:1.2**
  - **registry.redhat.io/codeready-workspaces/stacks-cpp-rhel8:1.2**
  - **registry.redhat.io/codeready-workspaces/stacks-node:1.2**
- To obtain the above images, you must have Red Hat login credentials. For details, see <https://access.redhat.com/RegistryAuthentication#getting-a-red-hat-login-2>.
- The OpenShift Origin Client Tools 3.11, **oc**, is in the path. See [Section 2.1, “Downloading the Red Hat OpenShift Origin Client Tools”](#).
- To perform the upgrade procedure, you must have a running instance of Red Hat OpenShift Container Platform 3.11 or OpenShift Dedicated 3.11, and must be logged in to the OpenShift instance using the following command:
  -

```
$ oc login https://<your-IP-address>:8443 --username <your-username> --password <your-password>
```

## Procedure

Use a tool for container management, such as Podman, to push the images to a local registry within the restricted environment.



### NOTE

The podman tool is available from the podman package in Red Hat Enterprise Linux starting with version 7.6. On earlier versions of Red Hat Enterprise Linux, the docker tool can be used with the same command syntax as suggested below for podman.

1. Pull the required images from **registry.redhat.io**. For example, to pull the **codeready-workspaces/server-rhel8** image, run:

```
# podman pull registry.redhat.io/codeready-workspaces/server-rhel8:1.2
```

Repeat the command for all the images that you want to pull.



### NOTE

After pulling images, it may be necessary to re-name and re-tag them so they are accessible from your local registry. For example:

```
# podman tag registry.redhat.io/codeready-workspaces/server-rhel8:1.2
registry.yourcompany.com:5000/crw-server-rhel8:1.2
```

1. To upgrade from CodeReady Workspaces 1.1 to CodeReady Workspaces 1.2, run a command like the following, using the correct registry and image names used in the previous steps (The default project namespace is **workspaces**. Your registry, container names, and tag versions may vary.)

```
$. /migrate_1.1_to_1.2.sh --project=<workspaces> \
--server-image=registry.yourcompany.com:5000/crw-server-rhel8:1.2 \
--operator-image=registry.yourcompany.com:5000/crw-server-operator-rhel8:1.2 \
--identity-provider-image=registry.yourcompany.com:5000/redhat-ss073-openshift:1.0-11 \
--postgres-image=registry.yourcompany.com:5000/rhscl-postgresql-96-rhel7:1-40 \
--pvc-jobs-image=registry.yourcompany.com:5000/rhel8-ubi8-minimal:8.0-127 \
--verbose
```

## 6.3. UPGRADING CODEREADY WORKSPACES MANUALLY



### NOTE

This section walks you through the manual upgrade steps from CodeReady Workspaces 1.1 to CodeReady Workspaces 1.2. For steps to manually upgrade from CodeReady Workspaces 1.0 to CodeReady Workspaces 1.1, see [https://access.redhat.com/documentation/en-us/red\\_hat\\_codeready\\_workspaces/1.1/html/administration\\_guide/upgrading-codeready-workspaces#upgrading-codeready-workspaces-manually](https://access.redhat.com/documentation/en-us/red_hat_codeready_workspaces/1.1/html/administration_guide/upgrading-codeready-workspaces#upgrading-codeready-workspaces-manually).

## Prerequisites

- You must have a running instance of Red Hat OpenShift Container Platform 3.11 or OpenShift Dedicated 3.11.
- The OpenShift Origin Client Tools 3.11, **oc**, is in the path. See [Section 2.1, “Downloading the Red Hat OpenShift Origin Client Tools”](#).
- You must be logged in to the OpenShift instance using the following command:

```
$ oc login https://<your-IP-address>:8443 --username <your-username> --password <your-password>
```

- You must have obtained the Red Hat login credentials. For details, see <https://access.redhat.com/RegistryAuthentication#getting-a-red-hat-login-2>.

## Procedure

1. Log in to the **registry.redhat.io** registry. For details, see <https://access.redhat.com/RegistryAuthentication#using-authentication-3>.
2. To keep your login secret for the **registry.redhat.io** registry in a separate file, run the following command:

```
$ docker --config /tmp/CRW.docker.config.json login https://registry.redhat.io
```

If you do not keep the secret in a separate file, the secret is stored in the `~/.docker/config.json` file and all your secrets are imported to OpenShift in the following step.

3. To add your secret to OpenShift, run the following commands:

```
$ oc create secret generic registryredhatio --type=kubernetes.io/dockerconfigjson --from-file=.dockerconfigjson=/tmp/CRW.docker.config.json && \
oc secrets link default registryredhatio --for=pull && \
oc secrets link builder registryredhatio
```


If the secret is added successfully, the following query shows your new secret:

```
$ oc get secret registryredhatio
```

If unsuccessful, see <https://access.redhat.com/RegistryAuthentication#allowing-pods-to-reference-images-from-other-secured-registries-9>.

4. Select your CodeReady Workspaces deployment (the default project namespace is **workspaces**):

```
$ oc project -n=<OpenShift-project-name>
```



### IMPORTANT

Substitute `<OpenShift-project-name>` for the project namespace that you are using.

5. Obtain the **DB\_PASSWORD** from the running CodeReady Workspaces instance:

```
$ DB_PASSWORD=$(oc get deployment keycloak -o=jsonpath=
{' .spec.template.spec.containers[0].env[?(@.name=="DB_PASSWORD")].value'})
```

6. Create the JSON patch:

```
$ PATCH_JSON=$(cat << EOF
{
  "spec": {
    "database": {
      "postgresImage": "registry.redhat.io/rhsc/postgresql-96-rhel7:1-40"
    },
    "auth": {
      "identityProviderImage": "registry.redhat.io/redhat-sso-7/sso73-openshift:1.0-11",
      "identityProviderPostgresPassword":"${DB_PASSWORD}"
    },
    "server": {
      "chelImage":"registry.redhat.io/codeready-workspaces/server-rhel8",
      "chelImageTag":"1.2"
    }
  }
}
EOF
)
```

7. Patch the running cluster:

```
$ oc patch checluster codeready -p "${PATCH_JSON}" --type merge
```

8. Scale down the running deployments:

```
$ oc scale deployment/codeready --replicas=0 && \
oc scale deployment/keycloak --replicas=0
```

9. Update the operator image and wait until the update is complete:

```
$ oc set image deployment/codeready-operator *=registry.redhat.io/codeready-
workspaces/server-operator-rhel8:1.2
```

10. After the operator is redeployed, scale up the keycloak deployment (RH SSO) and wait until the update is complete:

```
$ oc scale deployment/keycloak --replicas=1
```

11. After keycloak is redeployed, scale up the CodeReady Workspaces deployment (Che Server) and wait until the update is complete:

```
$ oc scale deployment/codeready --replicas=1
```

12. After CodeReady Workspaces is redeployed, set the new image for PostgreSQL:

```
$ oc set image deployment/postgres "**=registry.redhat.io/rhsc/postgresql-96-rhel7:1-40"
```

13. Scale down the PostgreSQL deployment and then scale it back up. Wait until the update is complete:

```
$ oc scale deployment/postgres --replicas=0 && \  
oc scale deployment/postgres --replicas=1
```

## CHAPTER 7. UNINSTALLING CODEREADY WORKSPACES

There is no dedicated function in the **deploy.sh** script to uninstall CodeReady Workspaces.

However, you can delete a custom resource, which deletes all the associated objects.

Note that in all the commands in this section, **<OpenShift-project-name>** is an OpenShift project with deployed CodeReady Workspaces (**workspaces** is the OpenShift project by default).

### Procedure

1. To delete CodeReady Workspaces and its components:

```
$ oc delete checluster/codeready -n=<OpenShift-project-name>
$ oc delete namespace -n=<OpenShift-project-name>
```

Use selectors to delete particular deployments and associated objects.

2. To remove all CodeReady Workspaces server-related objects:

```
$ oc delete all -l=app=che -n=<OpenShift-project-name>
```

3. To remove all Keycloak-related objects:

```
$ oc delete all -l=app=keycloak -n=<OpenShift-project-name>
```

4. To remove all PostgreSQL-related objects:

```
$ oc delete all -l=app=postgres -n=<OpenShift-project-name>
```

PVCs, service accounts, and role bindings should be deleted separately because the **oc delete all** command does not delete them.

5. To delete CodeReady Workspaces server PVC, ServiceAccount, and RoleBinding:

```
$ oc delete sa -l=app=che -n=<OpenShift-project-name>
$ oc delete rolebinding -l=app=che -n=<OpenShift-project-name>
$ oc delete pvc -l=app=che -n=<OpenShift-project-name>
```

6. To delete Keycloak and PostgreSQL PVCs:

```
$ oc delete pvc -l=app=keycloak -n=<OpenShift-project-name>
$ oc delete pvc -l=app=postgres -n=<OpenShift-project-name>
```



## CHAPTER 8. VIEWING CODEREADY WORKSPACES OPERATION LOGS

After the CodeReady Workspaces pods are created, you can view the operation logs of the application in the terminal or through the OpenShift console.

### 8.1. VIEWING CODEREADY WORKSPACES OPERATION LOGS IN THE TERMINAL

To view the operation logs on the terminal, run the following commands:

1. To obtain the names of the pods you must switch to project where CodeReady Workspaces is installed:

```
$ oc get pods -n=<OpenShift-project-name>
```

This command shows the pods that have been created:

NAME	READY	STATUS	RESTARTS	AGE
codeready-6b4876f56c-qdlll	1/1	Running	0	24m
codeready-operator-56bc9599cc-pkqkn	1/1	Running	0	25m
keycloak-666c5f9f4b-zz88z	1/1	Running	0	24m
postgres-96875bcbd-tfxr4	1/1	Running	0	25m

2. To view the operation log for a specific pod, run:

```
$ oc logs <log-name>
```

The output of this command for the **codeready-6b4876f56c-qdlll** pod (as an example) is as follows:

```
2019-03-29 12:22:17,710[ost-startStop-1] [INFO] [unknown.jul.logger 49] -
UndertowServer reused (pre-existing).
-----
GMS: address=codeready-6b4876f56c-qdlll-42917, cluster=WorkspaceStateCache, physical
address=fe80:0:0:0:42:acff:fe11:c%eth0:7801
-----
2019-03-29 12:22:21,784[ost-startStop-1] [INFO] [unknown.jul.logger 49] -
namespace [workspaces] set; clustering enabled
2019-03-29 12:22:21,785[ost-startStop-1] [INFO] [unknown.jul.logger 49] - Starting
UndertowServer on port 8888 for channel address: codeready-6b4876f56c-qdlll-34344
2019-03-29 12:22:21,786[ost-startStop-1] [INFO] [unknown.jul.logger 49] -
UndertowServer reused (pre-existing).
-----
GMS: address=codeready-6b4876f56c-qdlll-34344, cluster=WorkspaceLocks, physical
address=fe80:0:0:0:42:acff:fe11:c%eth0:7802
-----
2019-03-29 12:22:25,822[ost-startStop-1] [INFO] [o.e.c.a.w.s.WorkspaceRuntimes 166] -
Configured factories for environments: '[openshift, kubernetes, dockerimage,
```

3. For operation logs of the other pods, run:

- For the **codeready-operator-56bc9599cc-pkqkn** pod: **oc logs codeready-operator-**

**56bc9599cc-pkqkn**

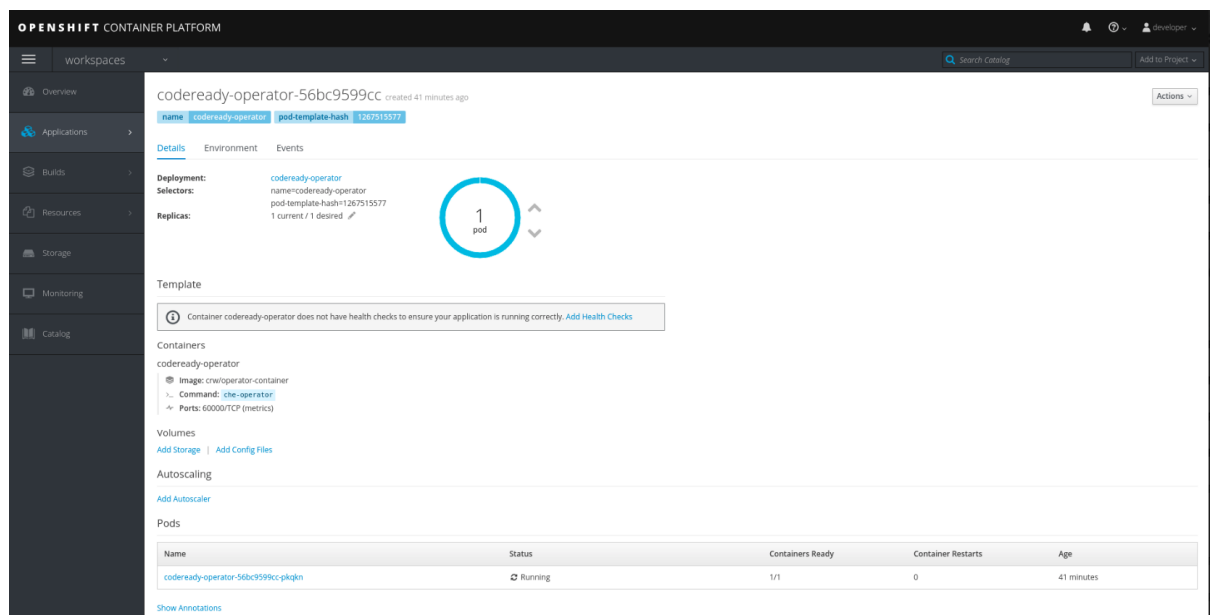
- For the `keycloak-666c5f9f4b-zz88z` pod: **oc logs keycloak-666c5f9f4b-zz88z**
- For the `postgres-96875bcbd-tfxr4` pod: **oc logs postgres-96875bcbd-tfxr4**

## 8.2. VIEWING CODEREADY WORKSPACES OPERATION LOGS IN THE OPENSIFT CONSOLE

To view the operation logs in the OpenShift console, take the following steps:

1. Navigate to the OpenShift web console.
2. In the **My Projects** pane, click **workspaces**.
3. In the **Overview** tab, click the application that you want to view the logs for (example: **codeready-operator, #1**).
4. In the **Deployments > <application-name>** window, click the name of the pod.
5. Scroll down to the **Pods** section and click the `<pod-name>`.
6. Click **Logs**

**Figure 8.1. Clicking View Log**



7. Click **Follow** to follow the log.

```
time="2019-03-29T12:22:00Z" level=info msg="Updating codeready CR with status: Che API: Unavailable"
time="2019-03-29T12:22:00Z" level=info msg="Custom resource codeready updated"
time="2019-03-29T12:22:00Z" level=info msg="Waiting for deployment codeready. Default timeout: 420 seconds"
time="2019-03-29T12:22:31Z" level=info msg="Deployment 'codeready' successfully scaled to 1"
time="2019-03-29T12:22:31Z" level=info msg="Updating codeready CR with Keycloak URL status: http://keycloak-workspaces.192.168.42.114.nip.io "
time="2019-03-29T12:22:31Z" level=info msg="Custom resource codeready updated"
```

```
time="2019-03-29T12:22:31Z" level=info msg="Updating codeready CR with status:
CodeReady Workspaces server: Available"
time="2019-03-29T12:22:31Z" level=info msg="Custom resource codeready updated"
time="2019-03-29T12:22:31Z" level=info msg="Updating codeready CR with CodeReady
Workspaces server URL: http://codeready-workspaces.192.168.42.114.nip.io "
time="2019-03-29T12:22:31Z" level=info msg="Custom resource codeready updated"
time="2019-03-29T12:22:31Z" level=info msg="CodeReady Workspaces is now available at:
http://codeready-workspaces.192.168.42.114.nip.io "
time="2019-03-29T12:22:31Z" level=info msg="Updating codeready CR with version: 1.1-52"
time="2019-03-29T12:22:31Z" level=info msg="Custom resource codeready updated"
```

## CHAPTER 9. USING THE CHE 7 IDE IN CODEREADY WORKSPACES



### IMPORTANT

CodeReady Workspaces 1.2 is based on upstream Che 6. The next version of Che, Che 7, is being worked on. To try this upcoming version of Che 7 in CodeReady Workspaces, follow the instructions in this section.

Che 7-based stacks are not included by default in CodeReady Workspaces. However, you can use a Che 7-based workspace configuration in CodeReady Workspaces. To use the Che 7-based stacks in CodeReady Workspaces, take the following steps:

1. In the **Dashboard**, click **Workspaces** and then click **Add Workspace**.
2. Select any stack from the list and click the dropdown icon next to **CREATE & OPEN**.
3. Click **Create & Proceed Editing**.
4. Click the **Config** tab and replace the default configuration content with the following content:

```
{
  "environments": {},
  "projects": [],
  "name": "che7",
  "attributes": {
    "editor": "org.eclipse.che.editor.theia:1.0.0",
    "plugins": "che-machine-exec-plugin:0.0.1"
  },
  "commands": [],
  "links": []
}
```

5. To add additional Che plugins (for example to add support for a particular language), add them in the workspace configuration attributes#plugins list. For a list of available plugins, see <https://github.com/eclipse/che-plugin-registry/>.
6. To add a runtime environment (for example Java runtime environment), add the following content in the **Config** tab:

```
"environments": {
  "default": {
    "machines": {
      "ws/dev": {
        "attributes": {
          "memoryLimitBytes": "536870912"
        },
        "servers": {},
        "volumes": {
          "projects": {
            "path": "/projects"
          }
        }
      },
      "installers": [],
```

```
    "env": {}
  }
},
"recipe": {
  "type": "kubernetes",
  "content": "kind: List\nitems:\n - \n apiVersion: v1\n kind: Pod\n metadata:\n  name:\n  ws\n spec:\n  containers:\n   - \n   image: 'eclipse/che-dev:nightly'\n   name: dev\n  resources:\n   limits:\n    memory: 512Mi\n",
  "contentType": "application/x-yaml"
}
}
```

7. To define it as the default environment, add the following content:

```
"defaultEnv": "default"
"environments": {
  "default": {...}
}
```

8. Click the **Open** button and run the Che 7-based workspace that you just created.

## CHAPTER 10. USING THE ANALYTICS PLUG-IN IN CODEREADY WORKSPACES

The Analytics plug-in provides insights about application dependencies: security vulnerabilities, license compatibility, and AI-based guidance for additional, alternative dependencies.

The Analytics plug-in is enabled by default in the Java and Node.js stacks in CodeReady Workspaces. When a user opens the **pom.xml** or the **package.json** files, the dependencies are analyzed. The editor shows warnings for available CVEs or issues with any dependency.

## CHAPTER 11. USING VERSION CONTROL

CodeReady Workspaces natively supports the Git version control system (VCS), which is provided by the [JGit library](#). Versioning functionality is available in the IDE and in the terminal.

The following sections describe how to connect and authenticate users to a remote Git repository. Any operations that involve authentication on the remote repository need to be done via the IDE interface unless authentication is configured separately for the workspace machine (terminal, commands).

Private repositories require a secure SSH connection. In order to connect to Git repositories over SSH, an SSH key pair needs to be generated. SSH keys are saved in user preferences, so the same key can be used in all workspaces.



### NOTE

HTTPS Git URLs can only be used to push changes when OAuth authentication is enabled. See [Enabling authentication with social accounts and brokering](#).

## 11.1. GENERATING AND UPLOADING SSH KEYS

SSH keys can be generated in the CodeReady Workspaces user interface.

1. Go to **Profile > Preferences > SSH > VCS** and click the **Generate Key** button.
2. When prompted to provide the host name for your repository, use the bare host name (do not include **www** or **https**) as in the example below.
3. Save the resulting key to your Git-hosting provider account.



### IMPORTANT

The host name is an actual host name of a VCS provider. Examples: **github.com**, **bitbucket.org**.

### 11.1.1. Using existing SSH keys

You can upload an existing public key instead of creating a new SSH key. When uploading a key, add the host name (using no **www** or **https** - as in the example below). Note that the **Public key > View** button is not be available when using this option because the public file should be generated already.

The screenshot shows the Eclipse IDE interface. The main editor displays the following Java code for `GreetingController.java`:

```

1  package org.eclipse.che.examples;
2
3  import org.springframework.web.servlet.ModelAndView;
4  import org.springframework.web.servlet.mvc.Controller;
5
6  import javax.servlet.http.HttpServletRequest;
7  import javax.servlet.http.HttpServletResponse;
8
9  public class GreetingController implements Controller
10 {
11
12     @Override
13     public ModelAndView handleRequest(HttpServletRequest request, HttpServletResponse
14     {
15         String userName = request.getParameter("user");
16         String result = "";
17         if (userName != null)
18         {
19             result = "Howdy, " + userName + "!";
20         }
21
22         ModelAndView view = new ModelAndView("hello_view");
23         view.addObject("greeting", result);
24         return view;
25     }
26 }

```

At the bottom, a terminal window shows the output of the `ls -la` command:

```

user@ff417a030d54:/projects$ ls -la
total 16
drwxr-xr-x. 3 user root 4096 Dec 26 08:48 .
drwxr-xr-x. 1 root root 4096 Dec 26 08:47 ..
drwxr-xr-x. 5 user root 4096 Dec 26 08:48 web-java-spring
user@ff417a030d54:/projects$

```

The following examples are specific to GitHub and GitLab, but a similar procedure can be used with all Git or SVN repository providers that use SSH authentication. See documentation provided by other providers for additional assistance.

## IMPORTANT

The default type of SSH key does not work in CodeReady Workspaces. The **git pull** command fails on authorization. To work around the issue, you need a different type of an SSH key.

To convert an existing SSH key to the **PEM** format, run the following command and substitute **<your\_key\_file>** for the path to your key file.

```
$ ssh-keygen -p -f <your_key_file> -m PEM
```

Or, to generate a **new PEM** SSH key instead of the default **RFC 4716/SSH2 key**, omit the **-p** flag. Substitute **<new\_key\_file>** for the path to the new file.

```
$ ssh-keygen -f <new_key_file> -m PEM
```

### Example 11.1. GitHub example

When not using GitHub OAuth, you can manually upload a key. To add the associated public key to a repository or account on **github.com**:

1. Click your user icon (top right).



2. Go to **Settings > SSH and GPG keys** and click the **New SSH key** button.
3. Enter a title and paste the public key copied from CodeReady Workspaces to the **Key** text field.

## SSH keys / Add new

### Title

### Key

```
ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQDTzatzHms/00I51NikNT7iODwVKBScjdqM1pS/WaGGkOLWX0y9qSwK
H/cYFX4O+T1O10mqSbyJi5XUeT+Q/Twbc4DweDMdb89BIRnpOBwadoWiV79DuHhJFfeMe0B9v4edzpf7b4g3eA
rhUYjE/Pn45ZCZbi32r95kJxXOKkhpRc3RkwwXjZ+Rf65ugr5m09RNavsgz+yPLn42geHBtuXCD296Aqq8UWZQVG
FcGHjJVs7FbkOZEvs7uCJEKLAK1Dxbv/D2ssgc+AMxMlfG9cwjRdNsBf0GzbUHxY3zu7IDMeLdH3fn1CG00stLS+K
I4gHSAUbFvROMoZVo2xtTt41V eclipse@che
```

### Example 11.2. GitLab example

To add the associated public key to a Git repository or account on **gitlab.com**:

1. Click your user icon (top right).
2. Go to **Settings > SSH Keys**.
3. Enter a title and paste the public key copied from CodeReady Workspaces to the **Key** text field.

## Add an SSH key





Before you can add an SSH key you need to [generate it](#).

### Key

```
ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQBTzatzHms/00l51NikNT7jODwVKBScdqM1pS/WaGGkOL
WX0y9qSwKH/cYFX4O+T1O10mqSbyJi5XUeT+Q/Twbc4DweDMdb89BIRnpOBwadoWiV79DuHh
JFfeMe0B9v4edzfp7b4g3eArhUYjE/Pn45ZCZbi32r95kJxXQKkhpRc3RkwwXjZ+Rf65ugr5m09RNav
sqz+yPLn42geHBtuXCD296Aqq8UWZQVGFcGHjJV57FbkOZEvs7uCJEKLAK1Dxbv/D2ssgc+AMxMI
fG9cwjRdNsBf0GzbUHxY3zu7lDMeLdH3fn1CG00stLS+KI4gHSAUbFvROmoZVo2xtTt41V
eclipse@che
```

### Title

### Your SSH keys (2)

 <a href="#">codenvy.io</a> b7:c5:d4:cf:76:62:81:ee:a4:5e:4a:ab:c7:e4:c8:da last used: n/a	created a year ago	
 <a href="#">che</a> 9e:4d:6d:6f:ea:77:19:c1:70:b9:52:38:d2:d5:e9:ed last used: n/a	created a year ago	

## 11.2. CONFIGURING GITHUB OAUTH

OAuth for Github allows users to import projects using SSH addresses (**git@**), push to repositories, and use the pull request panel. To enable automatic SSH key upload to GitHub for users:

1. On **github.com**, click your user icon (top right).
2. Go to **Settings > Developer settings > OAuth Apps**
3. Click the **Register a new application** button.
4. In the **Application name** field, enter, for example, **CodeReady Workspaces**.
5. In the **Homepage URL** field, enter **http://\${CHE\_HOST}:\${CHE\_PORT}**.
6. In the **Authorization callback URL** field, enter **http://\${CHE\_HOST}:\${CHE\_PORT}/api/oauth/callback**.

### Application name

Something users will recognize and trust

### Homepage URL

The full URL to your application homepage

### Application description

This is displayed to all users of your application

### Authorization callback URL

Your application's callback URL. Read our [OAuth documentation](#) for more information.

[Update application](#)[Delete application](#)

7. On OpenShift, update the deployment configuration.

```
CHE_OAUTH_GITHUB_CLIENTID=<your-github-client-id>
CHE_OAUTH_GITHUB_CLIENTSECRET=<your-github-secret>
CHE_OAUTH_GITHUB_AUTHURI=https://github.com/login/oauth/authorize
CHE_OAUTH_GITHUB_TOKENURI=https://github.com/login/oauth/access_token
CHE_OAUTH_GITHUB_REDIRECTURIS=http://{CHE_HOST}:{CHE_PORT}/api/oauth/callback
```



#### NOTE

- Substitute all occurrences of **`{CHE_HOST}`** and **`{CHE_PORT}`** with the URL and port of your CodeReady Workspaces installation.
- Substitute **`<your-github-client-id>`** and **`<your-github-secret>`** with your GitHub client ID and secret.

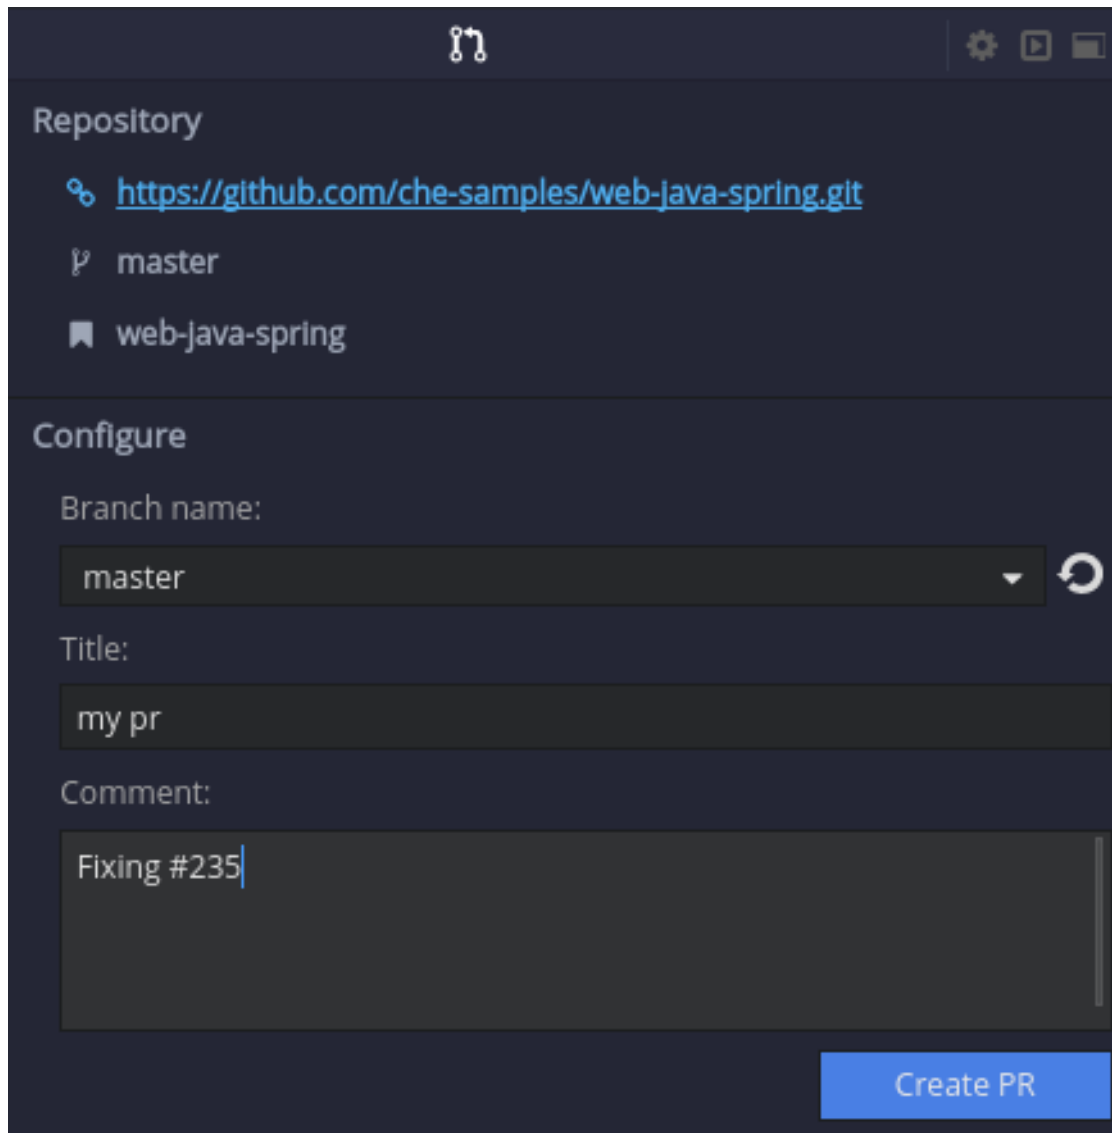
Once OAuth is configured, SSH keys are generated and uploaded automatically to GitHub by a user in the IDE in **Profile > Preferences > SSH > VCS** by clicking the *Octocat* icon. You can connect to your GitHub account and choose repositories to clone, rather than having to manually type (or paste) GitHub project URLs.

## 11.3. CONFIGURING GITLAB OAUTH

OAuth integration with GitLab is not supported. Although GitLab supports OAuth for clone operations, pushes are not supported. A feature request to add support exists in the GitLab issue management system: [Allow runners to push via their CI token](#) .

## 11.4. SUBMITTING PULL REQUESTS USING THE BUILT-IN PULL REQUEST PANEL

CodeReady Workspaces provides a **Pull Request** panel to simplify the creation of pull requests for GitHub, BitBucket, and Microsoft VSTS (with Git) repositories.



The screenshot shows a dark-themed interface for creating a pull request. At the top, there is a Git logo and three icons: a gear, a play button, and a folder. Below this, the 'Repository' section contains a link icon followed by the URL `https://github.com/che-samples/web-java-spring.git`, a branch icon followed by 'master', and a folder icon followed by 'web-java-spring'. The 'Configure' section has three fields: 'Branch name:' with a dropdown menu showing 'master' and a refresh icon; 'Title:' with a text input field containing 'my pr'; and 'Comment:' with a text area containing 'Fixing #235'. At the bottom right, there is a blue button labeled 'Create PR'.

## 11.5. SAVING COMMITTER NAME AND EMAIL

Committer name and email are set in **Profile > Preferences > Git > Committer**. Once set, every commit will include this information.

## 11.6. INTERACTING WITH GIT FROM A WORKSPACE

After importing a repository, you can perform the most common Git operations using interactive menus or as terminal commands.



## NOTE

Terminal Git commands require their own authentication setup. This means that keys generated in the IDE work only when Git is used through the IDE menus. Git installed in a terminal is a different Git system. You can generate keys in `~/ssh` there as well.

The screenshot shows an IDE interface with a dark theme. The top part displays a Maven POM file with XML content. The bottom part shows a terminal window with a shell prompt.

```

1 <?xml version="1.0" encoding="UTF-8"?>
2 <project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
3   xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/maven-v4_0_0.xsd">
4   <modelVersion>4.0.0</modelVersion>
5   <groupId>org.eclipse.che.examples</groupId>
6   <artifactId>web-java-spring</artifactId>
7   <packaging>war</packaging>
8   <version>1.0-SNAPSHOT</version>
9   <name>SpringDemo</name>
10  <properties>
11    <maven.compiler.source>1.6</maven.compiler.source>
12    <maven.compiler.target>1.6</maven.compiler.target>
13  </properties>
14  <dependencies>
15    <dependency>
16      <groupId>javax.servlet</groupId>
17      <artifactId>servlet-api</artifactId>
18      <version>2.5</version>
19      <scope>provided</scope>
20    </dependency>
21  </dependencies>

```

The terminal window shows the following prompt:

```

user@64a00a7a13d7:~/projects/web-java-spring$


```

Use keyboard shortcuts to access the most frequently used Git functionality faster:

Commit	<b>Alt+C</b>
Push to remote	<b>Alt+Shift+C</b>
Pull from remote	<b>Alt+P</b>
Work with branches	<b>Ctrl+B</b>
Compare current changes with the latest repository version	<b>Ctrl+Alt+D</b>

## 11.7. GIT STATUS HIGHLIGHTING IN THE PROJECT TREE AND EDITOR

Files in project explorer and editor tabs can be colored according to their Git status:



- Green: new files that are staged in index
- Blue: files that contain changes
- Yellow: files that are not staged in index

The editor displays change markers according to file edits:

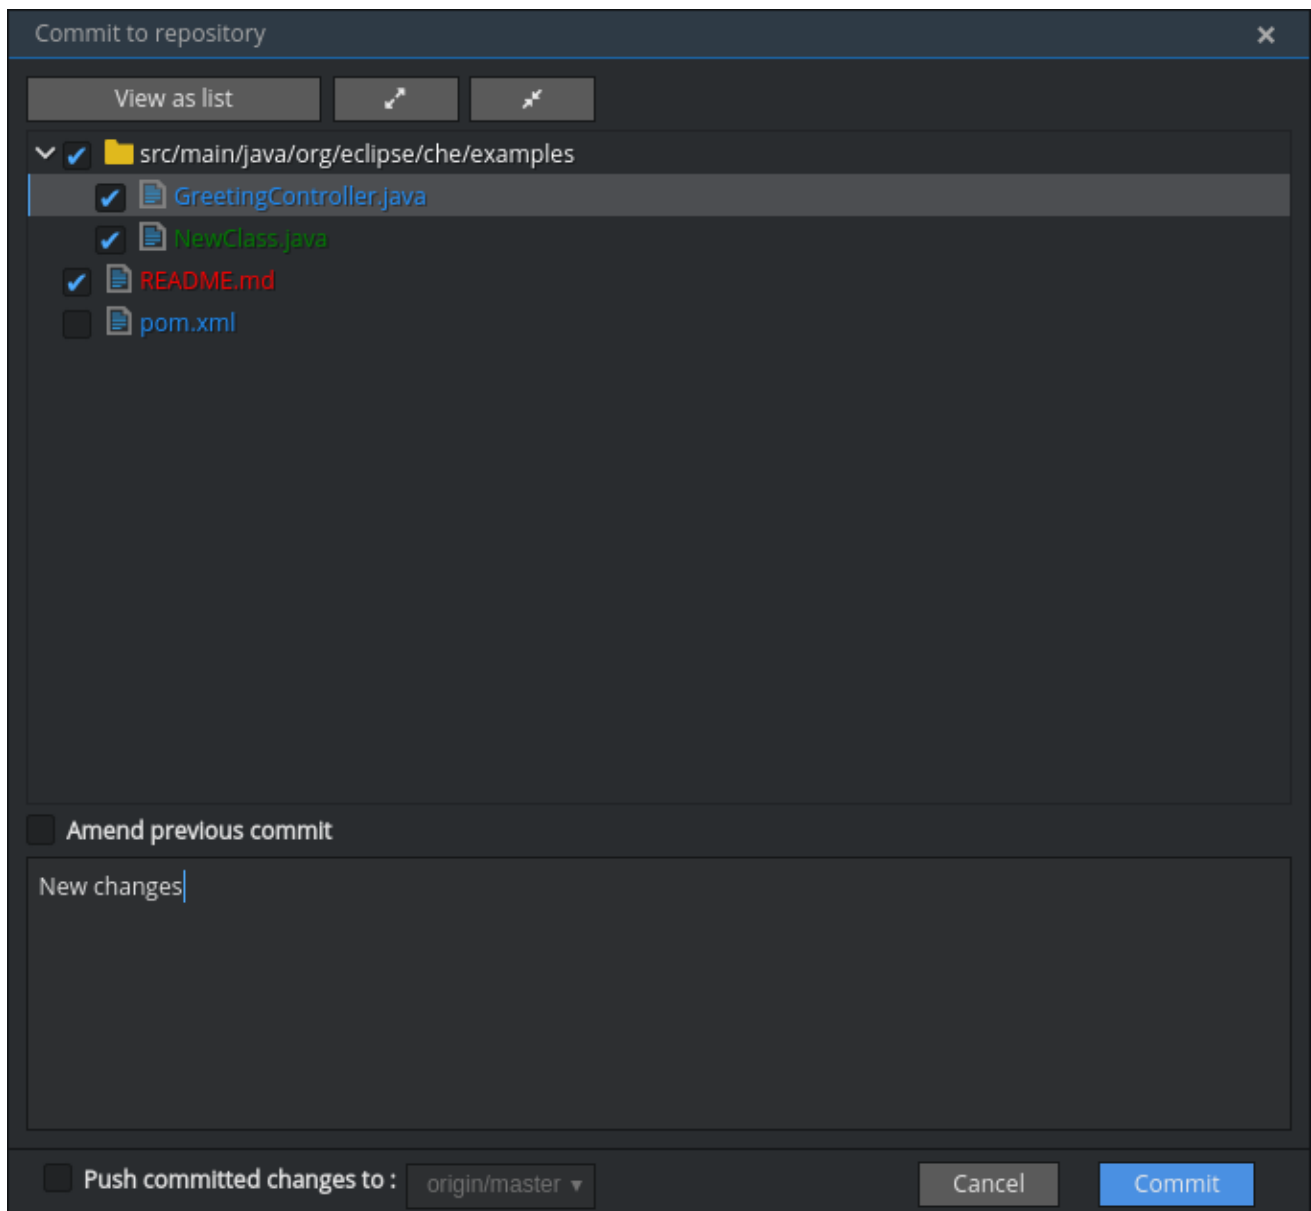


- Yellow marker: modified line(s)
- Green marker: new line(s)
- White triangle: removed line(s)

## 11.8. PERFORMING GIT OPERATIONS

### 11.8.1. Committing

Commit your changes by navigating to **Git > Commit...** in the main menu, or use the **Alt+C** shortcut.

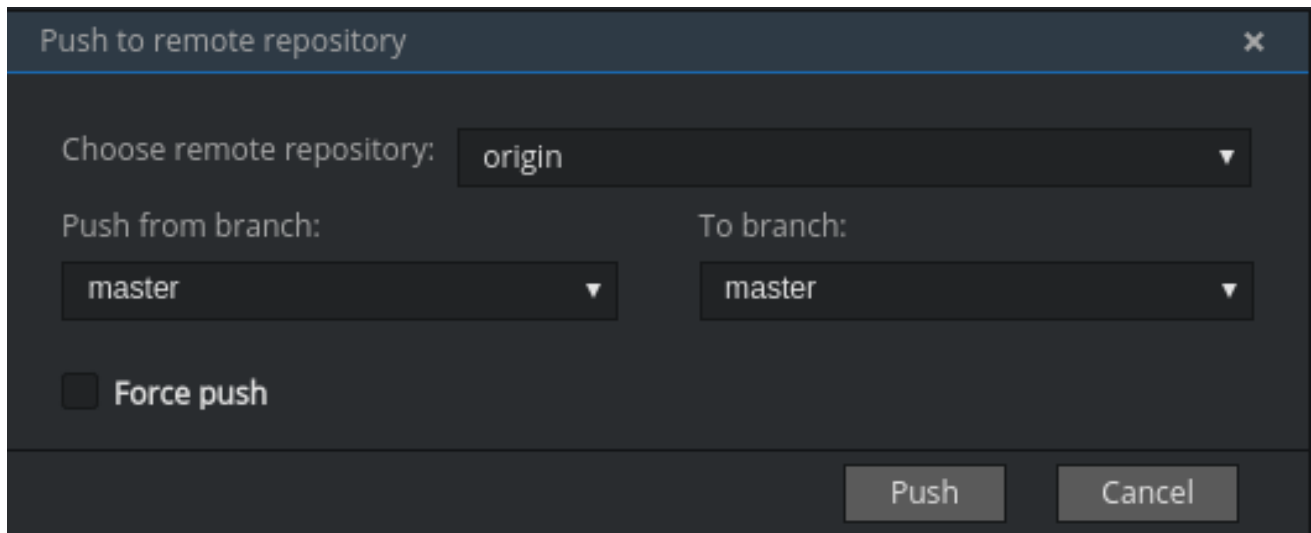


1. Select files that will be added to index and committed. All files in the selected package or folder in the project explorer are checked by default.
2. Type your commit message. Optionally, you can select **Amend previous commit** to modify the previous commit (for more details, see Git commit documentation).
3. To push your commit to a remote repository by checking the **Push committed changes to** check-box and select a remote branch.
4. Click **Commit** to proceed (the **Commit** button is active when at least one file is selected and a commit message is present, or **Amend previous commit** is checked).

Behavior for files in the list view is the same as in the **Compare** window (see [Reviewing changed files](#) section). Double-clicking a file opens diff window with it.

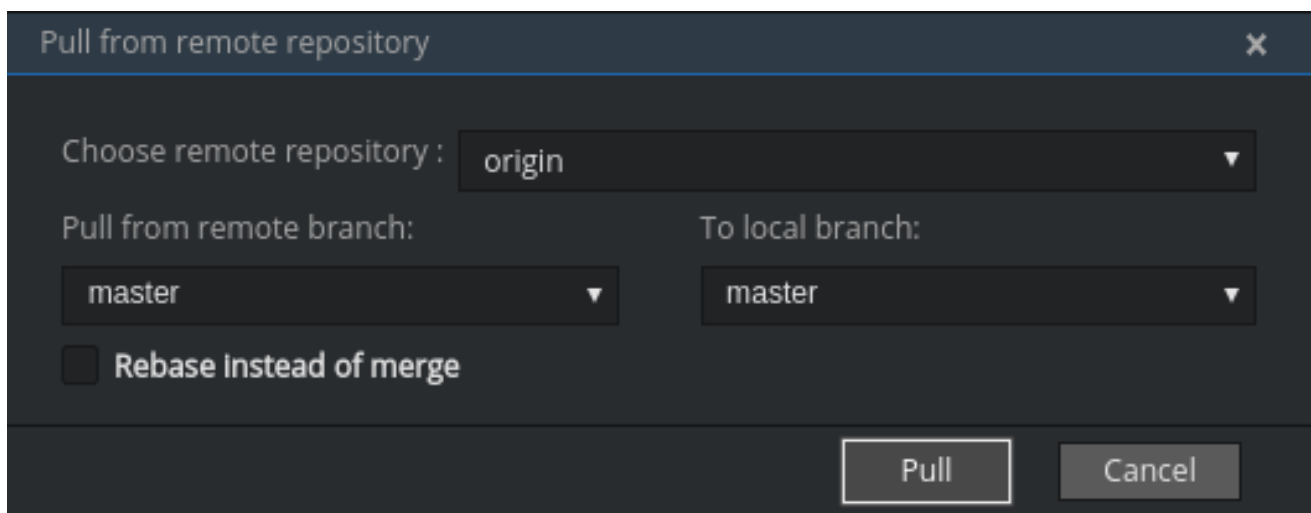
### 11.8.2. Pushing and pulling

Push your commits by navigating to **Git > Remotes... > Push** in the main menu, or use the **Alt+Shift+C** shortcut.



1. Choose the remote repository.
2. Choose the local and remote branch.
3. Optionally, you can force select **Force push**.

Get changes from a remote repository by navigating to **Git > Remotes... > Pull** in the main menu, or use the **Alt+P** shortcut.

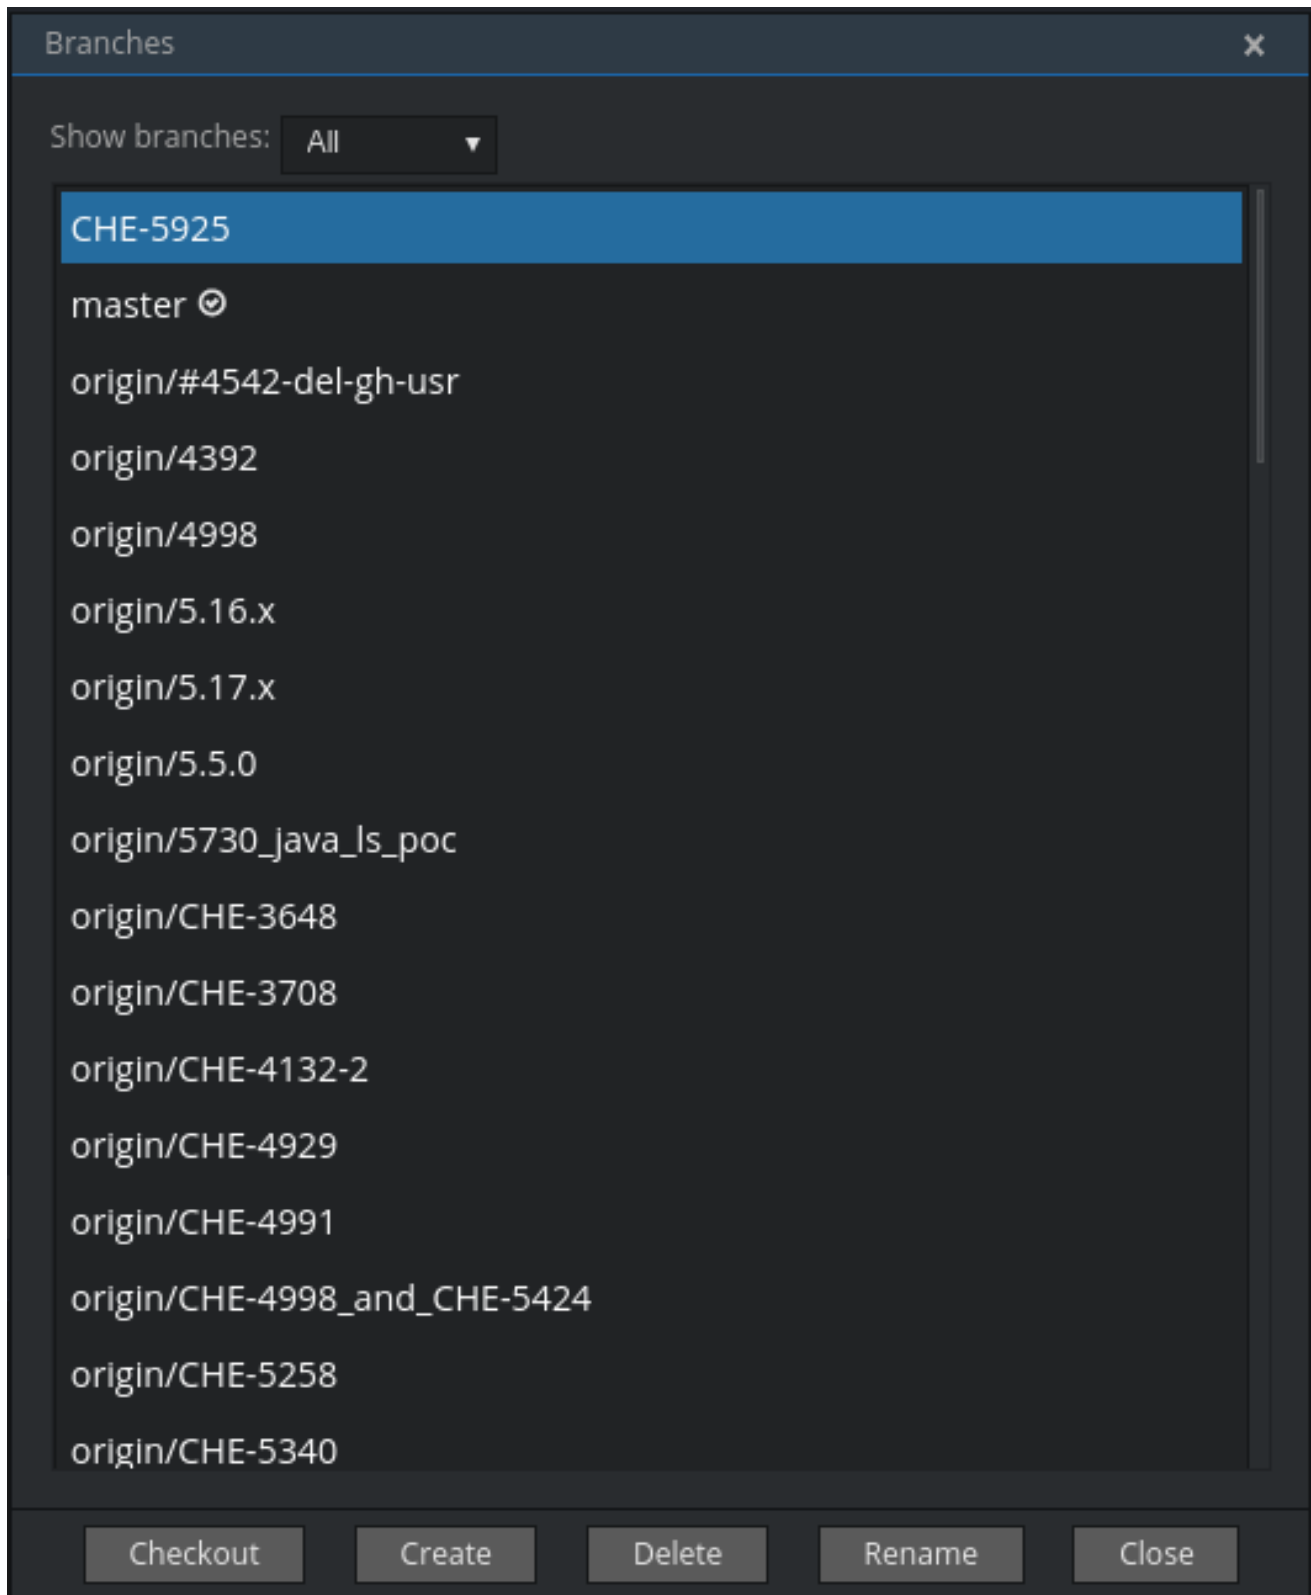


You can use **Rebase instead of merge** to keep your local commits on top (for more information, see [Git pull documentation](#)).

### 11.8.3. Managing branches

Manage your git branches by navigating to **Git > Branches...** in the main menu, or use the **Ctrl+B** shortcut.





You can filter the branches view by choosing to see only local or remote branches.


## 11.9. REVIEWING CHANGED FILES

The **Git Compare** window is used to show files that have changed.

To compare the current state of code to the latest local commit, navigate to **Git > Compare > Select-to-what** in the main menu, or use the **Ctrl+Alt+D** shortcut. Another way is to select an object in the project tree and choose **Git > Select-to-what** from the context menu of an item.


## Listing changed files

The **Git Compare** window shows changed files in the selected object in the project explorer. To see all changes, select a project folder. If only one file has changed, a diff window is shown instead of the compare window.



By default, affected files are listed as a tree.

The **Expand all directories** and **Collapse all directories** options help to get a better view. The **View as list** button switches the view of changed files to a list, where each file is shown with its full path. To return to the tree view, click **Group by directories**.



## Viewing diffs

To view a diff for a file, select the file and click **Compare**, or double-click the file name.

You can review changes between two states of code. To view the diff, go to **Git > Compare > Select-to-what** in main menu. If more than one file has changed, a list of the changed files is opened first. To select a file to compare, double-click it, or select a file, and then click **Compare**. Another way to open a diff is to select a file in the **Projects Explorer** and choose **Git > Select-to-what** from its context menu or directly from the context menu in the editor.

```

Review diff in: pom.xml (2/3)

Your version:
1 <?xml version="1.0" encoding="UTF-8"?>
2 <project xmlns="http://maven.apache.org/POM
3   xsi:schemaLocation="http://maven.ap
4   <modelVersion>4.0.0</modelVersion>
5   <groupId>org.eclipse.che.examples</grou
6   <artifactId>web-java-spring</artifactId
7   <packaging>war</packaging>
8   <version>1.1-SNAPSHOT</version>
9   <name>SpringDemo</name>
10  <properties>
11    <maven.compiler.source>1.6</maven.comp
12    <maven.compiler.target>1.6</maven.comp
13  </properties>
14  <dependencies>
15    <dependency>
16      <groupId>javax.servlet</groupId>
17      <artifactId>servlet-api</artifa
18      <version>2.5</version>
19      <scope>provided</scope>
20    </dependency>
21    <dependency>
22      <groupId>org.springframework</g
23      <artifactId>spring-webmvc</arti
24      <version>3.0.6.RELEASE</version>
25    </dependency>
26    <dependency>
27      <groupId>junit</groupId>
28      <artifactId>iunit</artifactId>

HEAD(Read-only)
7 <packaging>war</packaging>
8 <version>1.0-SNAPSHOT</version>
9 <name>SpringDemo</name>
10 <properties>
11   <maven.compiler.source>1.6</maven.comp
12   <maven.compiler.target>1.6</maven.comp
13 </properties>
14 <dependencies>
15   <dependency>
16     <groupId>javax.servlet</groupId>
17     <artifactId>servlet-api</artifactI
18     <version>2.5</version>
19     <scope>provided</scope>
20   </dependency>
21   <dependency>
22     <groupId>org.springframework</grou
23     <artifactId>spring-webmvc</artifac
24     <version>3.0.5.RELEASE</version>
25   </dependency>
26   <dependency>
27     <groupId>junit</groupId>
28     <artifactId>junit</artifactId>
29     <version>4.12</version>
30     <scope>test</scope>
31   </dependency>
32 </dependencies>
33 </project>
34

Previous Next Save Changes Refresh Close

```

Your changes are shown on the left, and the file being compared to is on the right. The left pane can be used for editing and fixing your changes.

To review multiple files, you can navigate between them using the **Previous** (or **Alt+.**) and **Next** (or **Alt+,**) buttons. The number of files for review is displayed in the title of the diff window.

The **Refresh** button updates the difference links between the two panes.

# CHAPTER 12. CODEREADY WORKSPACES ADMINISTRATION GUIDE

## 12.1. RAM PREREQUISITES

You must have at least 5 GB of RAM to run CodeReady Workspaces. The Keycloak authorization server and PostgreSQL database require the extra RAM. CodeReady Workspaces uses RAM in this distribution:

- CodeReady Workspaces server: approximately 750 MB
- Keycloak: approximately 1 GB
- PostgreSQL: approximately 515 MB
- Workspaces: 2 GB of RAM per workspace. The total workspace RAM depends on the size of the workspace runtime(s) and the number of concurrent workspace pods.

### 12.1.1. Setting default workspace RAM limits

The default workspace RAM limit and the RAM allocation request can be configured by passing the **CHE\_WORKSPACE\_DEFAULT\_MEMORY\_LIMIT\_MB** and **CHE\_WORKSPACE\_DEFAULT\_MEMORY\_REQUEST\_MB** parameters to a CodeReady Workspaces deployment.

For example, use the following configuration to limit the amount of RAM used by workspaces to 2048 MB and to request the allocation of 1024 MB of RAM:

```
$ oc set env dc/che CHE_WORKSPACE_DEFAULT_MEMORY_LIMIT_MB=2048 \
    CHE_WORKSPACE_DEFAULT_MEMORY_REQUEST_MB=1024
```



#### NOTE

- The user can override the default values when creating a workspace.
- A RAM request greater than the RAM limit is ignored.

## 12.2. REQUIREMENTS FOR RESOURCE ALLOCATION AND QUOTAS

Workspace pods are created in the account of the user who deploys CodeReady Workspaces. The user needs enough quota for RAM, CPU, and storage to create the pods.

## 12.3. SETTING UP THE PROJECT WORKSPACE

Workspace objects are created differently depending on the configuration. CodeReady Workspaces currently supports two different configurations:

- Single OpenShift project
- Multi OpenShift project

### 12.3.1. Setting up a single OpenShift project

To setup a single OpenShift project:

1. Define the service account used to create workspace objects with the **CHE\_OPENSHIFT\_SERVICEACCOUNTNAME** variable.
2. To ensure this service account is visible to the CodeReady Workspaces server, put the service and the CodeReady Workspaces server in the same namespace.
3. Give the service account permissions to create and edit OpenShift resources.
4. If the developer needs to create an object outside of the service accounts bound namespace, give the service account cluster-admin rights by running this command:

```
$ oc adm policy add-cluster-role-to-user self-provisioner system:serviceaccount:eclipse-che:che
```

In the command above, **eclipse-che** is the CodeReady Workspaces namespace.

### 12.3.2. Setting up a multi OpenShift project

1. To create workspace objects in different namespaces for each user, set the **NULL\_CHE\_INFRA\_OPENSHIFT\_PROJECT** variable to **NULL**.
2. To create resources on behalf of the currently logged-in user, use the user's OpenShift tokens.

## 12.4. HOW THE CODEREADY WORKSPACES SERVER USES PVCS AND PVS FOR STORAGE

CodeReady Workspaces server, Keycloak and PostgreSQL pods, and workspace pods use Persistent Volume Claims (PVCs), which are bound to the physical Persistent Volumes (PVs) with [ReadWriteOnce access mode](#). When the deployment YAML files run, they define the CodeReady Workspaces PVCs. You can configure [workspace PVC](#) access mode and claim size with CodeReady Workspaces deployment environment variables.



### IMPORTANT

When the **StorageClass** resource object in OpenShift is configured with **volumeBindingMode=WaitForFirstConsumer**, the workspaces in CodeReady Workspaces fail to start. To work around this issue, see [Workspaces fail to start with certain configurations of StorageClass](#) in the *CodeReady Workspaces Known Issues* document.

### 12.4.1. Storage requirements for CodeReady Workspaces infrastructure

- CodeReady Workspaces server: 1 GB to store logs and initial workspace stacks.
- Keycloak: 2 PVCs, 1 GB each to store logs and Keycloak data.
- PostgreSQL: 1 GB PVC to store database.

### 12.4.2. Storage strategies for workspaces

The workspace PVC strategy is configurable:

strategy	details	pros	cons
<b>unique (default)</b>	One PVC per workspace volume	Storage isolation	An undefined number of PVs is required
<b>common</b>	One PVC for all workspaces in one OpenShift Project  Sub-paths pre-created	Easy to manage and control storage	Workspaces must be in a separate OpenShift Project if PV does not support ReadWriteMany (RWX) access mode
<b>per-workspace</b>	One PVC for one workspace  Sub-paths pre-created	Easy to manage and control storage	Workspace containers must all be in one pod if PV does not support ReadWriteMany (RWX) access mode

### 12.4.3. Unique PVC strategy

To define the **unique** strategy, set **CHE\_INFRA\_KUBERNETES\_PVC\_STRATEGY** to **unique**.

Every workspace gets its own PVC, which means a workspace PVC is created when a workspace starts for the first time. A workspace PVC is deleted when a corresponding workspace is deleted.

### 12.4.4. Common PVC Strategy

#### 12.4.4.1. How the common PVC strategy works

All workspaces (within one OpenShift Project) use the same PVC to store data declared in their volumes (projects and workspace logs by default and whatever additional [volumes](#) that a user can define.)

A PV that is bound to PVC **che-claim-workspace** has the following structure:

```
pv0001
workspaceid1
workspaceid2
workspaceidn
che-logs projects <volume1> <volume2>
```

Volumes can be anything that a user defines as volumes for workspace machines. The volume name is equal to the directory name in **/\${PV}/\${ws-id}**.

When a workspace is deleted, a corresponding subdirectory (**/\${ws-id}**) is deleted in the PV directory.

#### 12.4.4.2. Enabling the common strategy

If you have already deployed CodeReady Workspaces with the **unique** strategy, set the **CHE\_INFRA\_KUBERNETES\_PVC\_STRATEGY** variable to **common** in **dc/che**.

If applying the **che-server-template.yaml** configuration, pass **-p CHE\_INFRA\_KUBERNETES\_PVC\_STRATEGY=common** to the **oc new-app** command.

### 12.4.4.3. Restrictions on using the common PVC strategy

When the **common** strategy is used and a workspace PVC access mode is ReadWriteOnce (RWO), only one OpenShift node can simultaneously use the PVC. If there are several nodes, you can use the **common** strategy, but the workspace PVC access mode is ReadWriteMany (RWM). Multiple nodes can use this PVC simultaneously.

To change the access mode for workspace PVCs, pass the **CHE\_INFRA\_KUBERNETES\_PVC\_ACCESS\_MODE=ReadWriteMany** environment variable to CodeReady Workspaces deployment either when initially deploying CodeReady Workspaces or through the CodeReady Workspaces deployment update.

Another restriction is that only pods in the same namespace can use the same PVC. The **CHE\_INFRA\_KUBERNETES\_PROJECT** environment variable should not be empty. It should be either the CodeReady Workspaces server namespace where objects can be created with the CodeReady Workspaces service account (SA) or a dedicated namespace where a token or a user name and password need to be used.

### 12.4.5. Per workspace PVC strategy

To define the unique strategy, set **CHE\_INFRA\_KUBERNETES\_PVC\_STRATEGY** to **per-workspace**.

#### 12.4.5.1. How the per-workspace PVC strategy works

The **per-workspace** strategy works similarly to the **common** PVC strategy. The only difference is that all workspace volumes (but not all workspaces) use the same PVC to store data (projects and workspace logs by default and any additional [volumes](#) that a user can define).

## 12.5. UPDATING YOUR CODEREADY WORKSPACES DEPLOYMENT

To update a CodeReady Workspaces deployment:

1. Change the image tag:

You can change the image tag in one of the following ways:

- On the command line, edit the image tag by running:

```
$ oc edit dc/che
```

- In the OpenShift web console, edit the **image:tag** line in the YAML file in **Deployments**
- Using the Docker service:

```
$ oc set image dc/che che=eclipse/che-server:${VERSION} --source=docker
```

2. Update Keycloak and PostgreSQL deployments (optional):

- Run the **eclipse/che-keycloak** command.
- Run the **eclipse/che-postgres** command.

You can get the list of available versions at [CodeReady Workspaces GitHub page](#).

3. Change the pull policy (optional):

To change the pull policy, do one of the following:



- Add **--set chelimagePullPolicy=IfNotPresent** to the CodeReady Workspaces deployment.
- Manually edit **dc/che** after deployment.

The default pull policy is **Always**. The default tag is **nightly**. This tag sets the image pull policy to **Always** and triggers a new deployment with a newer image, if available.

## 12.6. SCALABILITY

To run more workspaces, [add more nodes to your OpenShift cluster](#) . An error message is returned when the system is out of resources.

## 12.7. GDPR

To delete data or request the administrator to delete data, run this command with the user or administrator token:

```
$ curl -X DELETE http://che-server/api/user/{id}
```

## 12.8. DEBUG MODE

To run CodeReady Workspaces Server in debug mode, set the following environment variable in the CodeReady Workspaces deployment to **true** (default is **false**):

```
CHE_DEBUG_SERVER=true
```

## 12.9. PRIVATE DOCKER REGISTRIES

See [OpenShift documentation](#).

## 12.10. CODEREADY WORKSPACES SERVER LOGS

Logs are persisted in a PV .The PVC **che-data-volume** is [created](#) and bound to a PV after CodeReady Workspaces deploys to OpenShift.

To retrieve logs, do one of the following:

- Run the **oc get log dc/che** command.
- Run the **oc describe pvc che-data-claim** command to find the PV. Next, run the **oc describe pv \$pvName** command with the PV to get a local path with the logs directory. Be careful with permissions for that directory, since once changed, CodeReady Workspaces server will not be able to write to a file.
- In the OpenShift web console, select **Pods > che-pod > Logs**

It is also possible to configure CodeReady Workspaces master not to store logs, but produce JSON encoded logs to output instead. It may be used to collect logs by systems such as Logstash. To configure JSON logging instead of plain text environment variable **CHE\_LOGS\_APPENDERS\_IMPL** should have value **json**.

## 12.11. WORKSPACE LOGS

Workspace logs are stored in an PV bound to **che-claim-workspace** PVC. Workspace logs include logs from workspace agent, [bootstrapper](#) and other agents if applicable.

## 12.12. CODEREADY WORKSPACES MASTER STATES

The CodeReady Workspaces master has three possible states:

- **RUNNING**
- **PREPARING\_TO\_SHUTDOWN**
- **READY\_TO\_SHUTDOWN**

The **PREPARING\_TO\_SHUTDOWN** state means that no new workspace startups are allowed. This situation can cause two different results:

- If your infrastructure does not support workspace recovery, all running workspaces are forcibly stopped.
- If your infrastructure does support workspace recovery, any workspaces that are currently starting or stopping is allowed to finish that process. Running workspaces do not stop.

For those that did not stop, automatic fallback to the shutdown with full workspaces stopping will be performed.

If you want a full shutdown with workspaces stopped, you can request this by using the **shutdown=true** parameter. When preparation process is finished, the **READY\_TO\_SHUTDOWN** state is set which allows to stop current CodeReady Workspaces master instance.

## 12.13. WORKSPACE TERMINATION GRACE PERIOD

The default grace termination period of OpenShift workspace pods is **0**. This setting terminates pods almost instantly and significantly decreases the time required for stopping a workspace.

To increase the grace termination period, use the following environment variable:  
**CHE\_INFRA\_KUBERNETES\_POD\_TERMINATION\_GRACE\_PERIOD\_SEC.**



### IMPORTANT

If the **terminationGracePeriodSeconds** variable is explicitly set in the OpenShift recipe, the **CHE\_INFRA\_KUBERNETES\_POD\_TERMINATION\_GRACE\_PERIOD\_SEC** environment variable does not override the recipe.

## 12.14. AUTO-STOPPING A WORKSPACE WHEN ITS PODS ARE REMOVED

CodeReady Workspaces Server includes a job that automatically stops workspace runtimes if their pods have been terminated. Pods are terminated when, for example, users remove them from the OpenShift console, administrators terminate them to prevent misuse, or an infrastructure node crashes.

The job is disabled by default to avoid problems in configurations where CodeReady Workspaces Server cannot interact with the Kubernetes API without user intervention.

The job cannot function with the following CodeReady Workspaces Server configuration:

- CodeReady Workspaces Server communicates with the Kubernetes API using a token from the OAuth provider.

The job can function with the following CodeReady Workspaces Server configurations:

- Workspaces objects are created in the same namespace where CodeReady Workspaces Server is located.
- The **cluster-admin** service account token is mounted to the CodeReady Workspaces Server pod.

To enable the job, set the **CHE\_INFRA\_KUBERNETES\_RUNTIMES\_CONSISTENCY\_CHECK\_PERIOD\_MIN** environment variable to contain a value greater than **0**. The value is the time period in minutes between checks for runtimes without pods.

## 12.15. UPDATING CODEREADY WORKSPACES WITHOUT STOPPING ACTIVE WORKSPACES

The differences between a Recreate update and a Rolling update:

Recreate update	Rolling update
CodeReady Workspaces downtime	No CodeReady Workspaces downtime
-	New deployment starts in parallel and traffic is hot-switched

### 12.15.1. Performing a recreate update

To perform a recreate update:

- Ensure that the new master version is fully API compatible with the old workspace agent version.
- Set the deployment update strategy to Recreate
- Make POST request to the **/api/system/stop** api to start WS master suspend. This means that all new attempts to start workspaces will be refused, and all current starts and stops will be finished. Note that this method requires system admin credentials.
- Make periodical **GET** requests to the **/api/system/state** API, until it returns the **READY\_TO\_SHUTDOWN** state. Also, you can check for "System is ready to shutdown" in the server logs.
- Perform new deploy.

### 12.15.2. Performing a rolling update

To perform a rolling update:

- Ensure that the new master is fully API compatible with the old ws agent versions, as well as database compatibility. It is impossible to use database migrations on this update mode.

- Set the deployment update strategy set to Rolling.
- Ensure **terminationGracePeriodSeconds** deployment parameter has enough value (see details below).
- Press **Deploy** button or execute **oc rollout latest che** from cli client.

### 12.15.2.1. Known issues

- Workspaces may fallback to the stopped state when they are started five to thirty seconds before the network traffic are switched to the new pod. This happens when the bootstrappers use the CodeReady Workspaces server route URL for notifying the CodeReady Workspaces Server that bootstrapping is done. Since traffic is already switched to the new CodeReady Workspaces server, the old CodeReady Workspaces server cannot get the bootstrapper's report and fails to start after the waiting timeout is reached. If the old CodeReady Workspaces server is killed before this timeout, the workspaces can be stuck in the **STARTING** state. The **terminationGracePeriodSeconds** parameter must define enough time to cover the workspace start timeout, which is eight minutes plus some additional time. Typically, setting **terminationGracePeriodSeconds** to 540 sec is enough to cover all timeouts.
- Users may experience problems with websocket reconnections or missed events published by WebSocket connection when a workspace is **STARTED** but dashboard displays that it is **STARTING**. In this case, you need to reload the page to restore connections and the actual workspace states.

### 12.15.3. Updating with database migrations or API incompatibility

If new version of CodeReady Workspaces server contains some DB migrations, but there is still API compatibility between old and new version, recreate update type may be used, without stopping running workspaces.

API incompatible versions should be updated with full workspaces stop. It means that **/api/system/stop?shutdown=true** must be called prior to update.

## 12.16. DELETING DEPLOYMENTS

The fastest way to completely delete CodeReady Workspaces and its infrastructure components is to delete the project and namespace.

To delete CodeReady Workspaces and components:

```
$ oc delete namespace che
```

You can use selectors to delete particular deployments and associated objects.

To remove all CodeReady Workspaces server related objects:

```
$ oc delete all -l=app=che
```

To remove all Keycloak related objects:

```
$ oc delete all -l=app=keycloak
```

To remove all PostgreSQL-related objects:

```
$ oc delete all -l=app=postgres
```

PVCs, service accounts and role bindings should be deleted separately because **oc delete all** does not delete them.

To delete CodeReady Workspaces server PVC, ServiceAccount and RoleBinding:

```
$ oc delete sa -l=app=che
$ oc delete rolebinding -l=app=che
```

To delete Keycloak and PostgreSQL PVCs:

```
$ oc delete pvc -l=app=keycloak
$ oc delete pvc -l=app=postgres
```

## 12.17. MONITORING CODEREADY WORKSPACES MASTER SERVER

Master server emits metrics in Prometheus format by default on port **8087** of the CodeReady Workspaces server host (this can be customized by the **che.metrics.port** configuration property).

You can configure your own Prometheus deployment to scrape the metrics (as per convention, the metrics are published on the **<CHE\_HOST>:8087/metrics** endpoint).

The CodeReady Workspaces's Helm chart can optionally install Prometheus and Grafana servers preconfigured to collect the metrics of the CodeReady Workspaces server. When you set the **global.metricsEnabled** value to **true** when installing CodeReady Workspaces's Helm chart, Prometheus and Grafana servers are automatically deployed. The servers are accessible on **prometheus-<CHE\_NAMESPACE>.domain** or **grafana-<CHE\_NAMESPACE>.domain** domains respectively. The Grafana server is preconfigured with a sample dashboard showing the memory usage of the CodeReady Workspaces server. You can log in to the Grafana server using the predefined username **admin** with the default password **admin**.

## 12.18. CREATING WORKSPACE OBJECTS IN PERSONAL NAMESPACES

You can register the OpenShift server as an identity provider when CodeReady Workspaces is installed in multi-user mode. This allows you to create workspace objects in the OpenShift namespace of the user that is logged in CodeReady Workspaces through Keycloak.

To create a workspace object in the namespace of the user that is logged into CodeReady Workspaces:


- [Register](#), inside Keycloak, an OpenShift identity provider that points to the OpenShift console of the cluster.
- Configure CodeReady Workspaces to use the Keycloak identity provider to retrieve the OpenShift tokens of the CodeReady Workspaces users.

Every workspace action such as start or stop creates an OpenShift resource in the OpenShift user account. A notification message displays which allows you to link the Keycloak account to your OpenShift user account.

But for non-interactive workspace actions, such as workspace stop on idling or CodeReady Workspaces server shutdown, the dedicated OpenShift account configured for the Kubernetes infrastructure is used. See [Setting up the project workspace](#) for more information.

To install CodeReady Workspaces on OpenShift with this feature enabled, see [Section 2.3.2, “Deploying CodeReady Workspaces with a self-signed certificate and OpenShift OAuth”](#).

## 12.19. OPENSIFT IDENTITY PROVIDER REGISTRATION



### NOTE

Cluster-wide administrator rights are required to add an OAuth client.

To add the OpenShift identity provider:

1. Use the following settings in the Keycloak administration console:

[Identity Providers](#) > [Openshift v3](#)

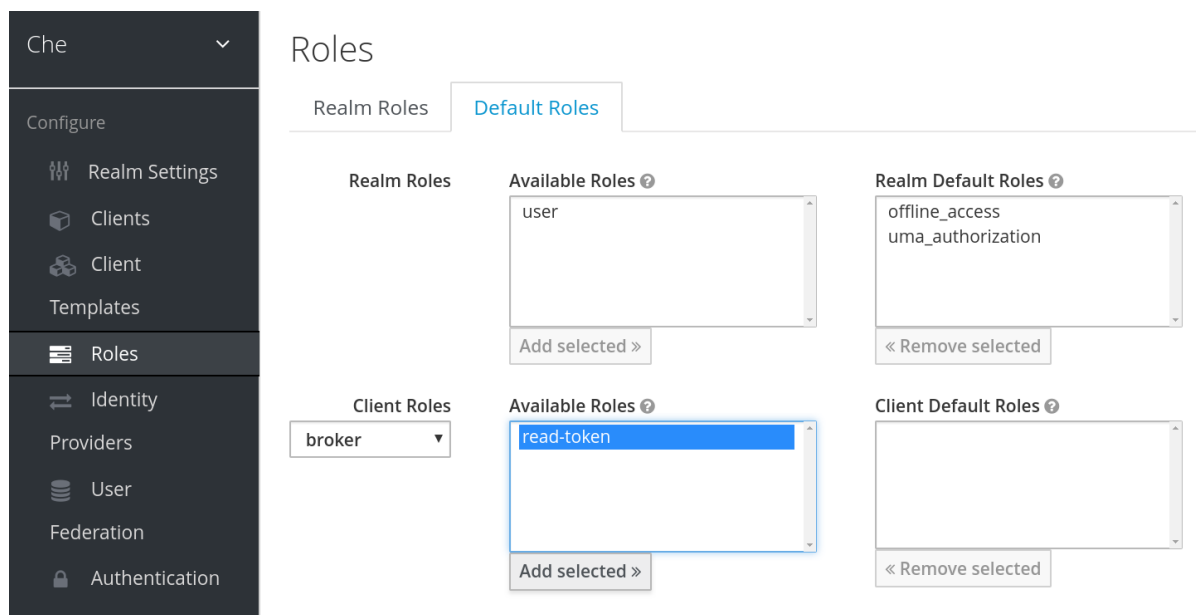
### Openshift v3

[Settings](#) [Mappers](#)

Redirect URI	<input type="text" value="http://keycloak-che.192.168.42.196.nip.io/auth/realms/che/broker/openshift-v3/endpoint"/>
* Alias	<input type="text" value="openshift-v3"/>
Display Name	<input type="text" value=""/>
* Client ID	<input type="text" value="kc-client"/>
* Client Secret	<input type="text" value="*****"/>
* Base URL	<input type="text" value="https://192.168.42.196:8443"/>
Default Scopes	<input type="text" value="user:full"/>
Store Tokens	<input checked="" type="checkbox"/>
Stored Tokens Readable	<input checked="" type="checkbox"/>
Enabled	<input checked="" type="checkbox"/>
Disable User Info	<input type="checkbox"/> OFF
Trust Email	<input type="checkbox"/> OFF
Account Linking Only	<input checked="" type="checkbox"/>
Hide on Login Page	<input checked="" type="checkbox"/>
GUI order	<input type="text" value=""/>
First Login Flow	<input type="text" value="first broker login"/>
Post Login Flow	<input type="text" value=""/>

The **Base URL** is the URL of the OpenShift console.

2. Add a default read-token role.



3. Declare the identity provider as an OAuth client inside OpenShift with the following commands:

```
$ oc create -f <(echo '
apiVersion: v1
kind: OAuthClient
metadata:
  name: kc-client
secret: "<value set for the 'Client Secret' field in step 1>"
redirectURIs:
  - "<value provided in the 'Redirect URI' field in step 1>"
grantMethod: prompt
')
```

See [Keycloak documentation](#) for more information on the Keycloak OpenShift identity provider.

## 12.20. CONFIGURING CODEREADY WORKSPACES

To configure CodeReady Workspaces deployment:

- Set the **CHE\_INFRA\_OPENSHIFT\_PROJECT** variable to **NULL** to ensure a new distinct OpenShift namespace is created for every workspace that is started.
- Set the **CHE\_INFRA\_OPENSHIFT\_OAUTH\_IDENTITY\_PROVIDER** variable to the alias of the OpenShift identity provider specified in step 1 of its [registration in Keycloak](#). The default value is **openshift-v3**.

## 12.21. PROVIDING THE OPENSIFT CERTIFICATE TO KEYCLOAK

If the certificate used by the OpenShift console is self-signed or is not trusted, then by default the Keycloak will not be able to contact the OpenShift console to retrieve linked tokens.

Keycloak cannot contact the OpenShift console to retrieve linked tokens when the certificate used by the OpenShift console is self-signed or is not trusted.

When the certificate is self-signed or is not trusted, use the **OPENSIFT\_IDENTITY\_PROVIDER\_CERTIFICATE** variable to pass the OpenShift console certificate to the Keycloak deployment. This will enable the Keycloak server to add the certificate to the list of

trusted certificates. The environment variable refers to a secret that contains the certificate.



## CHAPTER 13. MANAGING USERS

### 13.1. AUTHORIZATION AND USER MANAGEMENT

CodeReady Workspaces uses [Keycloak](#) to create, import, manage, delete, and authenticate users. Keycloak uses built-in authentication mechanisms and user storage. It can use third-party identity management systems to create and authenticate users. CodeReady Workspaces requires a Keycloak token when you request access to CodeReady Workspaces resources.

Local users and imported federation users must have an email address in their profile.

The default Keycloak credentials are **admin:admin**. You can use the **admin:admin** credentials when logging into CodeReady Workspaces for the first time. It has system privileges.

To find your Keycloak URL:

If CodeReady Workspaces is deployed on OpenShift:

- Go to the OpenShift web console and navigate to the **Keycloak** namespace.

### 13.2. CONFIGURING CODEREADY WORKSPACES TO WORK WITH KEYCLOAK

The deployment script ensures that Keycloak is properly configured when CodeReady Workspaces is deployed on OpenShift or installed on Docker. When the **che-public** client is created, the following fields are populated:

- **Valid Redirect URIs:** Use this URL to access CodeReady Workspaces.
- **Web Origins**

The following are common errors when configuring CodeReady Workspaces to work with Keycloak:

**Invalid redirectURI error:** occurs when you access CodeReady Workspaces at **myhost**, which is an alias, and your original **CHE\_HOST** is **1.1.1.1**. If this error occurs, go to the Keycloak administration console and ensure that the valid redirect URIs are configured.

**CORS error:** occurs when you have an invalid web origin

### 13.3. CONFIGURING KEYCLOAK TOKENS

A user token expires after 30 minutes by default.

You can change the following Keycloak token settings:

Che 

General

Login

Keys

Email


Themes

Cache

Tokens

Client Registration

Security Defenses

Revoke Refresh Token  OFFSSO Session Idle  SSO Session Max  Offline Session Idle  Access Token Lifespan  Access Token Lifespan For Implicit Flow  Client login timeout  Login timeout  Login action timeout  User-Initiated Action Lifespan  Default Admin-Initiated Action Lifespan   

## 13.4. SETTING UP USER FEDERATION

Keycloak federates external user databases and supports LDAP and Active Directory. You can test the connection and authenticate users before choosing a storage provider.

See the [User storage federation](#) page in Keycloak documentation to learn how to add a provider.

See the [LDAP and Active Directory](#) page in Keycloak documentation to specify multiple LDAP servers.

## 13.5. ENABLING AUTHENTICATION WITH SOCIAL ACCOUNTS AND BROKERING

Keycloak provides built-in support for GitHub, OpenShift, and most common social networks such as Facebook and Twitter. See Instructions to [enable Login with GitHub](#).

You can also enable the SSH key and upload it to the CodeReady Workspaces users' GitHub accounts.

To enable this feature when you register a GitHub identity provider:

1. Set scope to **repo,user,write:public\_key**.
2. Set store tokens and stored tokens readable to **ON**.

## GitHub

Settings

Mappers

Redirect URI ?	<input type="text" value="http://keycloak-eclipse-che.192.168.42.225.nip.io/auth/realms/che/broker/github/endpoint"/>
* Client ID ?	<input type="text" value="4c2b4d41fac8c7e3ec8c"/>
* Client Secret ?	<input type="password" value="....."/>
Default Scopes ?	<input type="text" value="repo,user,write:public_key"/>
Store Tokens ?	<input checked="" type="checkbox"/> ON
Stored Tokens Readable ?	<input checked="" type="checkbox"/> ON
Enabled ?	<input checked="" type="checkbox"/> ON
Disable User Info ?	<input type="checkbox"/> OFF
Trust Email ?	<input type="checkbox"/> OFF
Account Linking Only ?	<input type="checkbox"/> OFF
Hide on Login Page ?	<input type="checkbox"/> OFF
GUI order ?	<input type="text"/>
First Login Flow ?	<input type="text" value="first broker login"/>
Post Login Flow ?	<input type="text"/>
	<input type="button" value="Save"/> <input type="button" value="Cancel"/>

1. Add a default read-token role.

Che ▾

Configure

- ⚙️ Realm Settings
- 📦 Clients
- 👤 Client
- 📄 Templates
- ☰ Roles
- 🔗 Identity
- 🏠 Providers
- 👤 User
- 🌐 Federation
- 🔒 Authentication

## Roles

Realm Roles

Default Roles

Realm Roles

user

Available Roles ?

read-token

Realm Default Roles ?

offline\_access  
uma\_authorization

Client Roles

broker ▾

Available Roles ?

read-token

Client Default Roles ?

This is the default **delegated** OAuth service mode for CodeReady Workspaces. You can configure the OAuth service mode with the property **che.oauth.service\_mode**.

To use CodeReady Workspaces's OAuth Authenticator, set **che.oauth.service\_mode** to **embedded** and use Instructions for single-user mode.

See SSH key management for more information.

## 13.6. USING PROTOCOL-BASED PROVIDERS

Keycloak supports [SAML v2.0](#) and [OpenID Connect v1.0](#) protocols. You can connect your identity provider systems if they support these protocols.

## 13.7. MANAGING USERS

You can add, delete, and edit users in the user interface. See: [Keycloak User Management](#) for more information.

## 13.8. CONFIGURING SMTP AND EMAIL NOTIFICATIONS

CodeReady Workspaces does not provide any pre-configured MTP servers.

To enable SMTP servers in Keycloak:

1. Go to **che realm settings > Email**.
2. Specify the host, port, username, and password.

CodeReady Workspaces uses the default theme for email templates for registration, email confirmation, password recovery, and failed login.

## 13.9. CODEREADY WORKSPACES AUTHENTICATION

### 13.9.1. Authentication on CodeReady Workspaces Master

#### 13.9.1.1. OpenId

OpenId authentication on CodeReady Workspaces master, implies presence of an external OpenId provider and has 2 main steps:

1. Authenticate the user through the JWT token he brought or redirect him to login; (Authentication tokens should be send in a **Authorization** header. Also, in limited cases when it's not possible to use **Authorization** header, token can be send in **token** query parameter. An example of such exceptional case can be: OAuth authentication initialization, IDE shows javadoc in iframe where authentication must be initialized.)
2. Compose internal "subject" object which represents the current user inside of the CodeReady Workspaces master code.

At the time of writing the only supported/tested OpenId provider is Keycloak, so all examples/links will refer to this implementation.

The flow starts from the settings service where clients can find all the necessary URLs and properties of the OpenId provider such as **jwtks.endpoint**, **token.endpoint**, **logout.endpoint**, **realm.name**, **client\_id** etc returned. in JSON format.

Service class is `org.eclipse.che.multiuser.keycloak.server.KeycloakSettings`, and it is bound only in multi-user version of CodeReady Workspaces, so by its presence it is possible to detect if authentication enabled in current deployment or not.

Example output:

```
{
  "che.keycloak.token.endpoint": "http://172.19.20.9:5050/auth/realms/che/protocol/openid-connect/token",
  "che.keycloak.profile.endpoint": "http://172.19.20.9:5050/auth/realms/che/account",
  "che.keycloak.client_id": "che-public",
  "che.keycloak.auth_server_url": "http://172.19.20.9:5050/auth",
  "che.keycloak.password.endpoint": "http://172.19.20.9:5050/auth/realms/che/account/password",
  "che.keycloak.logout.endpoint": "http://172.19.20.9:5050/auth/realms/che/protocol/openid-connect/logout",
  "che.keycloak.realm": "che"
}
```

Also, this service allows to download JS client library to interact with provider. Service URL is `<che.host>:<che.port>/api/keycloak/settings` for retrieving settings JSON and `<che.host>:<che.port>/api/keycloak/OIDCKeycloak.js` for JS adapter library.

Next step is redirection of user to the appropriate provider's login page with all the necessary params like `client_id`, return redirection path etc. This can be basically done with any client library (JS or Java etc).

After user logged in on provider's side and client side code obtained and passed the JWT token, validation of it and creation of subject begins.

Verification of tokens signature occurs in the two main filters chain:

- `org.eclipse.che.multiuser.keycloak.server.KeycloakAuthenticationFilter` class. Token is extracted from **Authorization** header or **token** query param and tried to being parsed using public key retrieved from provider. In case of expired/invalid/malformed token, 403 error is sent to user. As noted above, usage of query parameter should be minimised as much as possible, since support of it may be limited/dropped at some point.

If validation was successful, token is passed to the

- `org.eclipse.che.multiuser.keycloak.server.KeycloakEnvironmentInitializationFilter` filter in the parsed form. This filter simply extracts data from JWT token claims, creates user in the local DB if it is not yet present, and constructs subject object and sets it into per-request **EnvironmentContext** object which is statically accessible everywhere.

If the request was made using machine token only (e.g. from ws agent) then it is only one auth filter in the chain:

- `org.eclipse.che.multiuser.machine.authentication.server.MachineLoginFilter` - finds userId given token belongs to, than retrieves user instance and sets principal to the session.

Master-to-master requests are performed using `org.eclipse.che.multiuser.keycloak.server.KeycloakHttpRequestFactory` which signs every request with the current subject token obtained from `EnvironmentContext`.

#### 13.9.1.1.1. User Profile

Since keycloak may store user specific information (first/last name, phone number, job title etc), there is special implementation of the ProfileDao which can provide this data to consumers inside CodeReady Workspaces. Implementation is read-only, so no create/update operations are possible. Class is `org.eclipse.che.multiuser.keycloak.server.dao.KeycloakProfileDao`.

#### 13.9.1.1.2. Obtaining Token From Keycloak

For the clients which cannot run JS or other type clients (like CLI or selenium tests), auth token may be requested directly from Keycloak. The simplest way to obtain Keycloak auth token, is to perform request to the token endpoint with username and password credentials. This request can be schematically described as following cURL request:

```
curl
  --data "grant_type=password&client_id=<client_name>&username=<username>&password=
  <password>"
  http://<keycloak_host>:5050/auth/realms/<realm_name>/protocol/openid-connect/token
```

Since the two main CodeReady Workspaces clients (IDE and Dashboard) utilizes native Keycloak js library, they're using a customized Keycloak login page and somewhat more complicated authentication mechanism using **grant\_type=authorization\_code**. It's a two step authentication: first step is login and obtaining authorization code, and second step is obtaining token using this code.

Example of correct token response:

```
{
  "access_token": "eyJhb...<rest of JWT token here>",
  "expires_in": 300,
  "refresh_expires_in": 1800,
  "refresh_token": "Nj0C...<rest of refresh token here>",
  "token_type": "bearer",
  "not-before-policy": 0,
  "session_state": "14de1b98-8065-43e1-9536-43e7472250c9"
}
```

#### 13.9.1.2. Other authentication implementations

If you want to adapt authentication implementation other than Keycloak, the following steps must be done:

- Write own or refactor existing info service which will provide list of configured provider endpoints to the clients;
- Write single or chain of filters to validate tokens, create user in CodeReady Workspaces DB and compose the Subject object;
- If the new auth provider supports OpenId protocol, OIDC JS client available at settings endpoint can be used as well since it is maximally decoupled of specific implementations.
- If the selected provider stores some additional data about user (first/last name, job title etc), it is a good idea to write an provider-specific ProfileDao implementation which will provide such kind of information.

#### 13.9.1.3. OAuth

OAuth authentication part has 2 main flows - internal and external based on Keycloak brokering

mechanism. So, there are 2 main OAuth API implementations - `org.eclipse.che.security.oauth.EmbeddedOAuthAPI` and `org.eclipse.che.multiuser.keycloak.server.oauth2.DelegatedOAuthAPI`.

They can be switched using `che.oauth.service_mode=<embedded|delegated>` configuration property.

Also, there is support of OAuth1 protocol can be found at `org.eclipse.che.security.oauth1` package.

The main REST endpoint in the OAuth API is `org.eclipse.che.security.oauth.OAuthAuthenticationService`, which contains `authenticate` method to start OAuth authentication flow, `callback` method to process callbacks from provider, `token` to retrieve current user's OAuth token, etc.

Those methods refer to the currently activated embedded/delegated OAuthAPI which is doing all the undercover stuff (finds appropriate authenticator, initializes the login process, user forwarding etc).

### 13.9.2. Authentication on CodeReady Workspaces Agents

Machines may contain services that must be protected with authentication, e.g. agents like workspace agent and terminal. For this purpose, machine authentication mechanism should be used. Machine tokens were introduced to avoid passing the Keycloak tokens to the machine side (which can be potentially insecure). Another reason is that Keycloak tokens may have relatively small lifetime and require periodical renewal/refresh which is hard to manage and keep in sync with same user session tokens on clients etc.

As agents cannot be queried using Keycloak token, there is only Machine Token option. Machine token can be also passed in header or query parameter.

#### 13.9.2.1. Machine JWT Token

Machine token is [JWT](#) that contains the following information in its claim:

- `uid` - id of user who owns this token
- `uname` - name of user who owns this token
- `wsid` - id of a workspace which can be queried with this token

Each user is provided with unique personal token for each workspace.

The structure of token and the signature are different to Keycloak and have the following view:

```
# Header
{
  "alg": "RS512",
  "kind": "machine_token"
}
# Payload
{
  "wsid": "workspacekrh99xjenek3h571",
  "uid": "b07e3a58-ed50-4a6e-be17-fcf49ff8b242",
  "uname": "john",
  "jti": "06c73349-2242-45f8-a94c-722e081bb6fd"
}
# Signature
```

```
{
  "value": "RSASHA512(base64UrlEncode(header) + . + base64UrlEncode(payload))"
}
```

The algorithm that is used for signing machine tokens is **SHA-512** and it's not configurable for now. Also, there is no public service that distributes the public part of the key pair with which the token was signed. But in each machine, there must be environment variables that contains key value. So, agents can verify machine JWT token using the following environment variables:

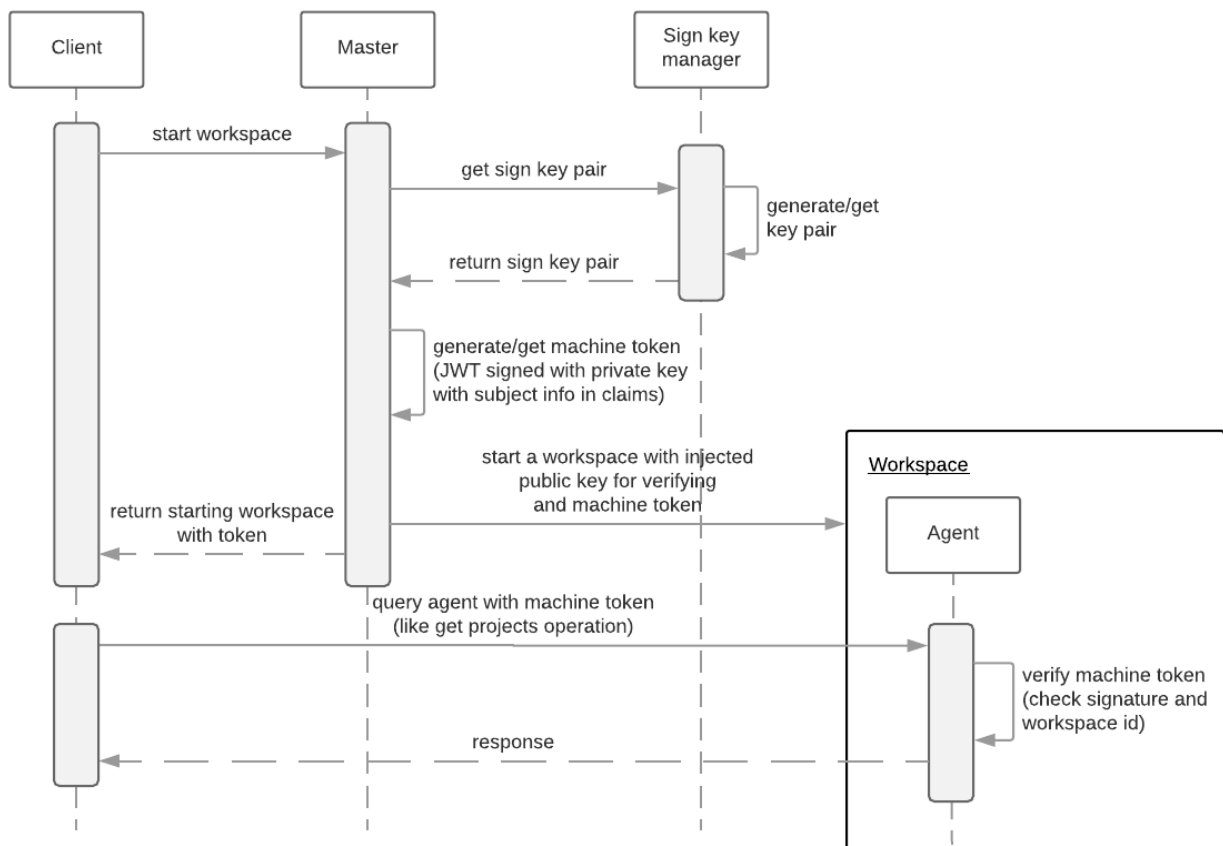
- **CHE\_MACHINE\_AUTH\_SIGNATURE\_\_ALGORITHM** - contains information about the algorithm which the token was signed
- **CHE\_MACHINE\_AUTH\_SIGNATURE\_\_PUBLIC\_\_KEY** - contains public key value encoded in Base64

It's all that is needed for verifying machine token inside of machine. To make sure that specified token is related to current workspace, it is needed to fetch **wsid** from JWT token claims and compare it with **CHE\_WORKSPACE\_ID** environment variable.

Also, if agents need to query CodeReady Workspaces Master they can use machine token provided in **CHE\_MACHINE\_TOKEN** environment, actually it is token of user who starts a workspace.

### 13.9.2.2. Authentication schema

The way how CodeReady Workspaces master interacts with agents with enabled authentication mechanism is the following:



Machine token verification on agents is done by the following components:





<http://172.19.20.180:32777/api/docs/swagger.json>

machinet9854958jdsdhakdh48f

Explore

## Eclipse Che

[project](#)[Show/Hide](#) | [List Operations](#) | [Expand Operations](#)[project-type](#)[Show/Hide](#) | [List Operations](#) | [Expand Operations](#)

## 13.10. PERMISSIONS

### 13.10.1. Overview

Permissions are used to control the actions of users and establish a security model. You can control resources managed by CodeReady Workspaces and allow certain actions by assigning permissions to users.

Permissions can be applied to the following:

- Workspace
- Organization
- Stack
- System

### 13.10.2. Workspace permissions

The user who creates a workspace is the workspace owner. The workspace owner has the following permissions by default: **read**, **use**, **run**, **configure**, **setPermissions**, and **delete**. Workspace owners invite users into the workspace and control workspace permissions for each user.

The following permissions are associated with workspaces:

Permission	Description
read	Allows reading the workspace configuration.
use	Allows using a workspace and interacting with it.
run	Allows starting and stopping a workspace.
configure	Allows defining and changing the workspace configuration.
setPermissions	Allows updating the workspace permissions for other users.
delete	Allows deleting the workspace.

### 13.10.3. Organization permissions

An organization is a named set of users.

The following permissions are applicable to organizations:

Permission	Description
update	Allows editing of the organization settings and information.
delete	Allows deleting an organization.
manageSuborganizations	Allows creating and managing sub-organizations.
manageResources	Allows redistribution of an organization's resources and defining the resource limits.
manageWorkspaces	Allows creating and managing all the organization's workspaces.
setPermissions	Allows adding and removing users and updating their permissions.

### 13.10.4. System permissions

System permissions control aspects of the whole CodeReady Workspaces installation.

The following permissions are applicable to the organization:

Permission	Description
manageSystem	Allows control of the system, workspaces, and organizations.
setPermissions	Allows updating the permissions for users on the system.
manageUsers	Allows creating and managing users.
monitorSystem	Allows for accessing endpoints used for monitoring the state of the server.

All system permissions will be granted to the administration user configured with the ``CHE_SYSTEM_ADMIN_NAME`` property (the default is **admin**). This happens at CodeReady Workspaces Server start. If the user is not present in the CodeReady Workspaces user database, it happens after the user's login.

### 13.10.5. manageSystem permission

Users with the **manageSystem** permission have access to the following services:

Path	HTTP Method	Description
<b>/resource/free/</b>	GET	Get free resource limits
<b>/resource/free/{accountId}</b>	GET	Get free resource limits for given account
<b>/resource/free/{accountId}</b>	POST	Edit free resource limit for given account
<b>/resource/free/{accountId}</b>	DELETE	Remove free resource limit for given account
<b>/installer/</b>	POST	Add installer to the registry
<b>/installer/{key}</b>	PUT	Update installer in the registry
<b>/installer/{key}</b>	DELETE	Remove installer from the registry
<b>/logger/</b>	GET	Get logging configurations in CodeReady Workspaces Server
<b>/logger/{name}</b>	GET	Get configurations of logger by its name in CodeReady Workspaces Server
<b>/logger/{name}</b>	PUT	Create logging in CodeReady Workspaces Server
<b>/logger/{name}</b>	POST	Edit logging in CodeReady Workspaces Server
<b>/resource/{accountId}/details</b>	GET	Get detailed information about resources for given account
<b>/system/stop</b>	POST	Shutdown all system services, prepare CodeReady Workspaces to stop
<b>/stacks</b>	All methods	All Stack service methods

### 13.10.6. monitorSystem permission

Users with the **monitorSystem** permission have access to the following services:

Path	HTTP Method	Description
------	-------------	-------------

Path	HTTP Method	Description
<b>/activity</b>	GET	Get workspaces in certain state for a certain amount of time

### 13.10.7. Super-privileged mode

The **manageSystem** permission can be extended to provide a super-privileged mode. This allows the user to perform advanced actions on any resources managed by the system. You can read and stop any workspaces with the **manageSystem** permission and assign permissions to users as needed.

The super-privileged mode is disabled by default. You can change to super-privileged mode by configuring the `CHE_SYSTEM_SUPER_PRIVILEGED_MODE` variable to **true** in the **che.env** file.

List of services that are enabled for users with **manageSystems** permissions and with super-privileged mode on:

Path	HTTP Method	Description
<b>/workspace/namespace/{namespace:.*}</b>	GET	Get all workspaces for given namespace.
<b>/workspace/{id}</b>	DELETE	Stop workspace
<b>/workspace/{key:.*}</b>	GET	Get workspace by key
<b>/organization/resource/{suborganizationId}/cap</b>	GET	Get resource cap for given organization
<b>/organization/resource/{suborganizationId}/cap</b>	POST	Set resource cap for given organization
<b>/organization/{parent}/organizations</b>	GET	Get child organizations
<b>/organization</b>	GET	Get user's organizations

### 13.10.8. Stack permissions

A stack is a runtime configuration for a workspace. See [stack definition](#) for more information on stacks.

The following permissions are applicable to stacks:

Permission	Description
search	Allows searching of the stacks.
read	Allows reading of the stack configuration.

Permission	Description
update	Allows updating of the stack configuration.
delete	Allows deleting of the stack.
setPermissions	Allows managing permissions for the stack.

### 13.10.9. Permissions API

All permissions can be managed using the provided REST API. The APIs are documented using Swagger at [\[\*\*{host}/swagger/#!/permissions\*\*\]](#).


### 13.10.10. Listing permissions

To list the permissions that apply to a specific resources, run this command:

```
$ GET /permissions
```

The **domain** values are:

Domain
<b>system</b>
<b>organization</b>
<b>workspace</b>
<b>stack</b>



#### NOTE

**domain** is optional. In this case, the API returns all possible permissions for all domains.

### 13.10.11. Listing permissions for a user

To list the permissions that apply to a user, run this command:

```
$ GET /permissions/{domain}
```

The **domain** values are:

Domain
<b>system</b>

Domain
<b>organization</b>
<b>workspace</b>
<b>stack</b>

### 13.10.12. Listing permissions for all users



#### NOTE

You must have sufficient permissions to see this information.

To list the permissions that apply to all users, run this command:

```
GET /permissions/{domain}/all
```

The **domain** values are:

Domain
<b>system</b>
<b>organization</b>
<b>workspace</b>
<b>stack</b>

### 13.10.13. Assigning permissions

To assign permissions to a resource, run this command:

**POST /permissions**

The **domain** values are:

Domain
<b>system</b>
<b>organization</b>
<b>workspace</b>

Domain
<b>stack</b>

The following is a message **body** that requests permissions for a user with a **userID** to a workspace with a **workspaceID**:

```
{
  "actions": [
    "read",
    "use",
    "run",
    "configure",
    "setPermissions"
  ],
  "userId": "userID",
  "domainId": "workspace",
  "instanceId": "workspaceID"
}
```

The **instanceId** parameter corresponds to the ID of the resource that retrieves the permission for all users. The **userId** parameter corresponds to the ID of the user that has been granted certain permissions.

### 13.10.14. Sharing permissions

A user with **setPermissions** privileges can share a workspace and grant **read, use, run, configure, or setPermissions** privileges to other users.

To share workspace permissions:

- Select a workspace in the user dashboard, navigate to the **Share** tab and enter emails of users. Use commas or space as separators if there are multiple emails.

## 13.11. ORGANIZATIONS

### 13.11.1. Organizations in CodeReady Workspaces

Organizations allow administrators to group CodeReady Workspaces users and allocate resources. The system administrator controls and allocates resources and permissions within the administrator dashboard.

#### 13.11.1.1. Roles in an organization

A user can have the following roles in an organization:

##### Members

Create workspaces, manage their own workspaces, and use any workspaces they have permissions for.

##### Administrators



Manage the organization, members, resources, and sub-organization, and can edit settings.

### System Administrators

Create root organizations, manages resources, members and sub-organizations. System administrators have more permissions than the administrators and members.

#### 13.11.1.2. Root organizations and sub-organizations

The top-level organizations are called root organizations. Multiple root organizations can be created. Any organization can have zero to a set number of sub-organizations. Only the system administrator can create root organizations and manage the resources of the root organization.

#### 13.11.1.3. Creating an organization

Only the system administrator can create root organizations. An administrator can create sub-organizations.

To create an organization:

1. Click the menu in the left sidebar. A new page displays all the organizations in your system.
2. Click on the upper-left button to create a new organization.

#### 13.11.1.4. Displaying the list of organizations

The **Organization** page displays a list of all the organizations. The list contains the following information for each organization: number of members, total RAM, available RAM, and number of sub-organizations.

#### 13.11.1.5. Adding members to organizations

To add members to an organization:

1. Click the **Add** button to add a member. A new pop-up window displays. You can change the role of a member or remove them from the organization at any time.
2. Enter the new member name.



#### NOTE

Users with the green checkmark beside their name already have a CodeReady Workspaces account and can be added to the organization. Users without a checkmark do not have an account and cannot be added into the organization.

#### 13.11.1.6. Workspaces in organizations

A workspace is created inside of an organization and uses the resources of the organization. The workspace creator chooses the organization on the **Workspace Creation** page.

#### 13.11.1.7. Setting email notifications

To send notifications from the CodeReady Workspaces server when a user joins or leaves an organization, you can do either of the following:

- Configure the SMTP server in the **che.env** file.

- For OpenShift, add environment variables to the deployment.

CodeReady Workspaces does not have a built-in SMTP server by default. You may use any mail server.

For example, to send a notification email to your Gmail account, set the following environment variables:

```
CHE_MAIL_PORT=465
CHE_MAIL_HOST=smtp.gmail.com
CHE_MAIL_SMTP_STARTTLS_ENABLE=true
CHE_MAIL_SMTP_AUTH=true
CHE_MAIL_SMTP_AUTH_USERNAME=no-reply@gmail.com
CHE_MAIL_SMTP_AUTH_PASSWORD=password
```

### 13.11.1.8. Creating sub-organizations

To create a sub-organization:

- On the **Organization Details** page, select the **Sub-Organizations** tab.
- Click the **Add Sub-Organization** button.

The steps to create a sub-organization are the same as that for creating an organization. Use them to create the organization.


### 13.11.1.9. Adding members to sub-organizations

You can only add members of the parent organization as members of the sub-organization.

### 13.11.1.10. Organization and sub-organization administration

The settings of the organization are visible to all members of the organization. Only the CodeReady Workspaces system administrator can modify the settings.

### 13.11.1.11. Renaming an organization or sub-organization



#### NOTE

Only an CodeReady Workspaces system administrator and administrator of the organization can rename an organization or sub-organization.

To rename an organization:

1. Click the **Name** field to edit the name of the organization. The save mode appears.
2. Click the **Save** button to update the name.

The name of the organization or sub-organization must follow these rules:


- Only alphanumeric characters and a single dash (-) can be used.
- Spaces cannot be used.
- Each organization name must be unique within the CodeReady Workspaces installation.

- Each sub-organization name must be unique within an organization.

#### 13.11.1.12. Leaving an organization or sub-organization

To leave an organization, members need to contact the administrator of the organization or the system administrator of CodeReady Workspaces.

#### 13.11.1.13. Deleting an organization or sub-organization



##### IMPORTANT

- Only system administrators or administrators of the organization can delete an organization or sub-organization.
- This action cannot be reverted, and all workspaces created under the organization will be deleted.
- All members of the organization will receive an email notification to inform them about the deletion of the organization.

To delete an organization or a sub-organization:

- Click the **Delete** button.

#### 13.11.1.14. Allocating resources for organizations

Workspaces use the resources of the organization that are allocated by the system administrator. The resources for sub-organizations are taken from the parent organization. Administrators control the portion of resources, of the parent organization, that are available to the sub-organization.

#### 13.11.1.15. Managing limits




##### NOTE

Managing limits is restricted to the CodeReady Workspaces system administrator and administrator of the organization.

The system configuration defines the default limits. The administrator of the organization manages only the limits of its sub-organizations. No resource limits apply to the organization by default. The following are the limits defined by the system administrator:

- **Workspace Cap:** The maximum number of workspaces that can exist in the organization.
- **Running Workspace Cap:** The maximum number of workspaces that can run simultaneously in the organization.
- **Workspace RAM Cap:** The maximum amount of RAM that a workspace can use in GB.

#### 13.11.1.16. Updating organization and sub-organization member roles

**NOTE**

Updating the members of an organization or sub-organization is restricted to the CodeReady Workspaces system administrator and administrator of the organization.

To edit the role of an organization member:

1. Click the **Edit** button in the **Actions** column. Update the role of the selected member in the pop-up window.
2. Click **Save** to confirm the update.

### 13.11.17. Removing members from an organization and sub-organization

**NOTE**

Removing the members of an organization or sub-organization is restricted to the CodeReady Workspaces system administrator and administrator of the organization.

To remove a member:

1. Click the **Delete** button in the **Actions** column. In the confirmation pop-up window, confirm the deletion.

To remove multiple members:

1. Select the check boxes to select multiple members from the organization.
2. Click the **Delete** button that appears in the header of the table. The members that are removed from the organization will receive an email notification.

## 13.12. RESOURCE MANAGEMENT

### 13.12.1. Overview

The Resource API manages the resources that are utilized by CodeReady Workspaces users. The CodeReady Workspaces administrators set the limits on the amount of resources available for each resource type and each user.

There are two kinds of accounts used in CodeReady Workspaces:

- *personal* - This account belongs to a user. Only one user can utilize resources provided to the account.
- *organizational* - This account belongs to an [organization](#). This type of account allows each member of the organization to use resources. Resources are distributed between an organization and sub-organizations.

Resource usage mostly refers to resources used by workspaces and runtimes in the development flow.

CodeReady Workspaces supports the following types of resources:

- **RAM** - Amount of RAM which can be used by running workspaces at the same time.

- **Timeout** - Period of time that is used to control idling of user workspaces.
- **Runtime** - Number of workspaces that users can run at the same time.
- **Workspace** - Number of workspaces that users can have at the same time.

### 13.12.2. Resource API

#### Total resources

**GET resource/{accountId}**: Gets the list of total resources an account can use;

#### Used resources

**GET resource/{accountId}/used**: Gets the resources used by an account;

#### Available resources

**GET resource/{accountId}/available**: Gets the resources that are available and not used by an account. If no resources are used, the available resources equal total resources. If resources are used, the available resources equals total resources minus used resources.

#### Resource details

**GET resource/{accountId}/details**: Gets detailed information about the resources available for an account. The detailed information includes: resource providers, resource-usage start time, and resource-usage end time.

For more information about the response objects and required parameters, see the Swagger page at [\\${che-host}/swagger/#/resource](#).

### 13.12.3. Distributing resources

The following are ways to distribute resources to an account:

- CodeReady Workspaces administrator specifies default free resources limit for account by configuration.
- CodeReady Workspaces administrator overrides the default free resources limit for account by resource-free API.

### 13.12.4. Configuring workspaces and resources

The CodeReady Workspaces administrator can limit how workspaces are created and the resources that these workspaces consume. Detailed information about each property can be found in the [che.env](#) file.

Property name	Default Value	Unit	Description
<b>CHE_LIMITS_USER_WORKSPACES_COUNT</b>	-1	item	maximum number of workspaces that the CodeReady Workspaces user can create

Property name	Default Value	Unit	Description
<b>CHE_LIMITS_USER_WORKSPACES_RUN_COUNT</b>	-1	item	maximum number of simultaneously running workspaces for a CodeReady Workspaces user
<b>CHE_LIMITS_USER_WORKSPACES_RAM</b>	-1	memory	maximum amount of RAM that workspaces use
<b>CHE_LIMITS_ORGANIZATION_WORKSPACES_COUNT</b>	-1	item	maximum number of workspaces that members of an organization can create
<b>CHE_LIMITS_ORGANIZATION_WORKSPACES_RUN_COUNT</b>	-1	item	maximum number of workspaces that members of an organization can simultaneously run
<b>CHE_LIMITS_ORGANIZATION_WORKSPACES_RAM</b>	-1	memory	maximum amount of RAM that workspaces from all organizations can simultaneously use
<b>CHE_LIMITS_WORKSPACE_IDLE_TIMEOUT</b>	-1	milliseconds	maximum number of workspaces that can stay inactive before they are idled
<b>CHE_LIMITS_WORKSPACE_ENV_RAM</b>	16gb	memory	maximum amount of RAM that workspace environment can use simultaneously

### 13.12.5. Unit formats

The unit has the following formats:

- **-1**: An unlimited value. Any operation, aggregation, and deduction of resources will return **-1**.
- **memory**: A plain or fixed-point integer measured in bytes. Memory uses one of the following suffixes:

Suffix name	Description
<b>k / kb / kib</b>	kilo bytes <b>1k = 1024b</b>
<b>m / mb / mib</b>	mega bytes <b>1m = 1024k</b>
<b>g / gb / gib</b>	giga bytes <b>1g = 1024m</b>
<b>t / tb / tib</b>	terra bytes <b>1t = 1024g</b>
<b>p / pb / pib</b>	peta bytes <b>1p = 1024t</b>

- **item** - An integer describing the number of objects.
- **millisecond** - An integer describing the time frame in milliseconds.

### 13.12.6. Resource-free API

The Resource-free API manages the free resources that are provided by the system configuration. You can override resources for an account.

#### Free Resources

**GET resource/free:** Gets the resources that are available.

**GET resource/free/{accountId}:** Gets the resources that are available for this account.

#### Set Free Resources

**POST resource/free:** Sets the maximum amount of resources available for the user organization account. This number overrides the Che configuration. It will be used in all further operations with resources.

#### Remove Free Resources

**DELETE resource/free/{accountId}:** Deletes the number of resources available for the user and organization account. The system configuration defines the amount of resources available.

For more information on response objects and required parameters, see the Swagger page at **{che-host}/swagger/#/resource-free**.

### 13.12.7. Organization Resource API

#### Distributed Organization Resources

**GET organization/resource/{organizationId}:** Gets the resources that the parent organization provides to the sub-organization.

#### Sub-Organization Resources Cap

**GET organization/resource/{suborganizationId}/cap:** Gets the maximum amount of resources that are available for a sub-organization; By default, sub-organizations can use all the resources of the parent organization.

## Set Sub-Organization Resources Cap

**POST organization/resource/{suborganizationId}/cap:** Sets the maximum amount of resources for a sub-organization. This limits the usage of shared resources by the sub-organization.

See the Swagger page at **{che-host}/swagger/#/organization-resource** for more detailed specification of response objects and required parameters.



## CHAPTER 14. ADMINISTERING WORKSPACES

### 14.1. WORKSPACE

A workspace is usually termed as a local directory with projects and meta-information that the integrated development environment (IDE) uses to configure projects. In CodeReady Workspaces, a workspace is the developer environment. The developer environment contains Docker containers, Kubernetes pods, and a virtual machine or localhost. Environment variables and storage volumes are part of the workspace. The developer environment also contains projects, project commands, and resource allocation attributes.

### 14.2. ENVIRONMENT

The workspace runtime environment is a set of machines where each machine is defined by a recipe. The environment is healthy when all the machines successfully start and the installers execute jobs. The environment is defined by a recipe that can have different types. The environment and infrastructure validate a recipe.

### 14.3. MACHINE

The runtime environment has a minimum of one machine that can be a Docker-formatted container or a Kubernetes pod. You can create multi-machine environments with as many machines as your project infrastructure requires. Each machine has a configuration and start policy. Machine crashes and start failures are signs of an unhealthy environment. Machines communicate by using the internal network, **service:port**.

### 14.4. RECIPE

A workspace environment is defined by a recipe. The recipe can be one of the following:

- single container image
- Dockerfile
- Docker Compose file
- Kubernetes list of objects with multiple pods and services

### 14.5. BOOTSTRAPPER

The bootstrapper starts the installer script after the first process is executed in the machine following the **CMD** or **ENTRYPOINT**. The role of the bootstrapper is to start the installer scripts with a set of parameters and a configuration file. The bootstrapper is a small binary compiled from Go code.

### 14.6. INSTALLER

The purpose of the installer is to install software and services, start servers, and activate agents. The workspace agent, executive agent, and terminal servers are important to the IDE and workspace. The language servers, SSH installer, and other servers bring new functionality to a workspace. The bootstrapper executes installer scripts that prepare the environment and checks for dependencies. See an [example of an installer script](#) that prepares the environment and installs the C# language server.

## 14.7. VOLUME

A volume is a fixed amount of storage that is used to persist workspace data. Workspace projects are automatically mounted into a host file system by default. A user can define extra volumes for each machine in the environment. Docker volumes, Kubernetes persistent volumes (PVs), and persistent volumes claims (PVCs) are examples of volumes.

## 14.8. ENVIRONMENT VARIABLES

The environment variables are propagated into each individual machine. Depending on the infrastructure, environment variables are propagated to Docker containers or Kubernetes pods.

## 14.9. WHAT IS NEXT?

- [Create and start your first workspace](#) .
- Learn how to define [volumes](#) and [environment variables](#).

## 14.10. MANAGING WORKSPACES


### 14.10.1. Creating workspaces

Use the [stacks](#) in the **Dashboard** to create a workspace. Images and configuration in these stacks are certified both for Docker and OpenShift. These stacks are used in daily functional testing.

#### 14.10.1.1. Creating a workspace from stacks in the dashboard

To create a workspace from stacks in the **Dashboard**, take the following steps:

1. In the **Dashboard**, in the left panel, click **Stacks**.
2. Click the **Duplicate stack** icon for the stack that you want to create a clone of. A page titled after the selected stack opens.
3. Edit the fields that you want to edit.
4. Click **Save**.




### 14.10.1.2. Duplicating an existing stack

Create a stack and then use the resulting stack to create a workspace.

To create a copy of an existing stack, take the following steps:

1. In the **Dashboard**, in the left panel, click **Stacks**.
2. Click the **Duplicate stack** icon for the stack that you want to clone.
3. Edit the **Name** field.
4. In the **Machines** field, edit the **Source** field.
5. Click **Save**. The **Stack is successfully updated** message confirms that the stack is updated.
6. In the **Dashboard**, click **Workspaces > Add Workspace**.
7. In the **SELECT STACK** section, scroll through the list to locate the stack that you created in the preceding steps.
8. Click **Create** to create the workspace based on this stack.

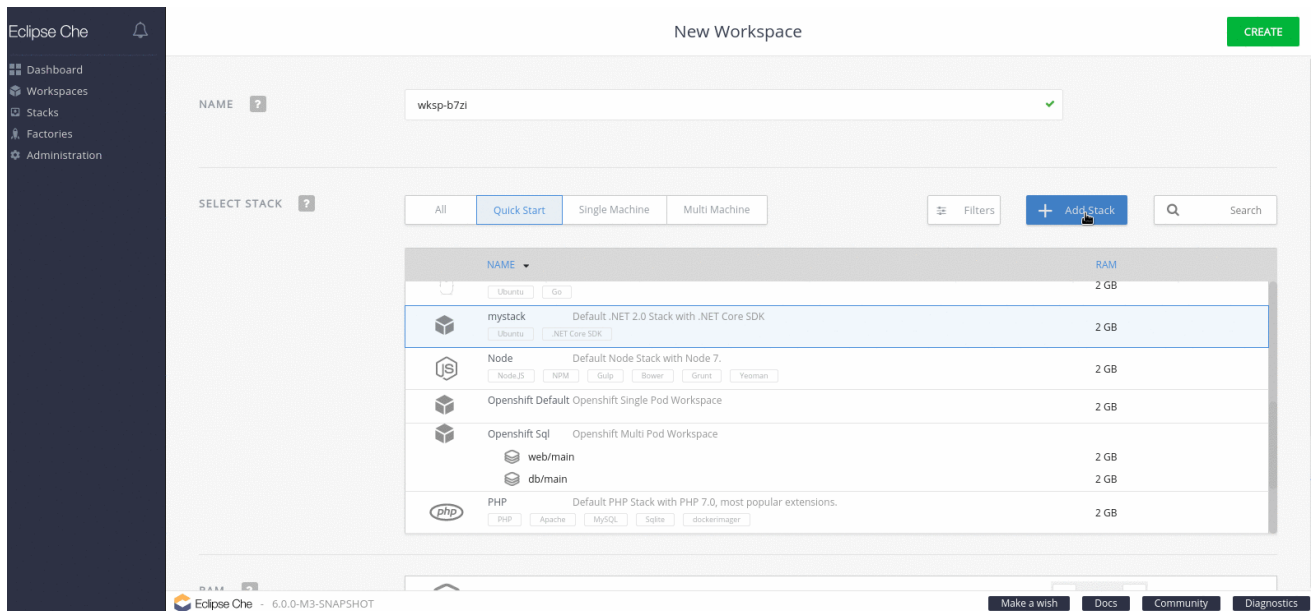


### 14.10.1.3. Creating a custom stack from a custom recipe

Author a custom [recipe](#) and then create a stack. Use the resulting stack to create a workspace.

To create a custom stack from a custom recipe, take the following steps:

1. In the **Dashboard**, click **Workspaces** > **Add Workspace**.
2. From the **SELECT STACK** field, select the required stack.
3. Click **Add Stack**.
4. In the **Create Stack** dialog box, click **Yes** to confirm that you want to create the stack from a recipe.
5. In the **Build stack from recipe** window, type the recipe name from which you want to build this stack (example: **FROM: eclipse/new-stack**).
6. Click **OK**.
7. In the **Name** field, type a name for the stack.
8. In the **Runtimes** > **Machines** > **Recipe** section, click **Show** to ensure that the stack is being created using the recipe that you set in the preceding steps (**eclipse/new-stack**, in this case).
9. Click **Save**.



### 14.10.2. Starting workspaces

You can start a workspace in one of the following ways:

- The workspace starts automatically after it is created in the user's **Dashboard**.
- In the user's **Dashboard**, use the **Run** or **Open** buttons in the **Workspace Details** view.
- Click a workspace name from the recent workspaces displayed in the left panel.
- Use the REST API.

The workspace may take time to start depending on factors like network conditions, container image availability, and configured installers attempting to install additional tools and software in the runtime. Track the progress of the workspace start operation in the **Workspace Start** tab. The tabs for each machine in the workspace environment stream logs from the installers (terminal, **exec** agent, **ws** agent, and language servers if any).

### 14.10.3. Managing a workspace

After a workspace is created or started, you can modify it by adding projects, installers, environment variables, and volumes.




#### IMPORTANT

To edit a raw workspace configuration, back up the working configuration to avoid breaking your workspace.

Change the configuration of a running workspace and saving it restarts the workspace. To learn more about workspace configuration, see:

- [Creating projects in workspaces](#)
- [Installers](#)
- [Environment variables](#)



## 14.11. COMMANDS AND IDE MACROS

A command is:

- A set of instructions that is injected into the workspace machine for execution
- A goal to organize commands for your workflow
- A context to scope the command to particular projects
- A preview URL to expose the URL of a running server

Commands are saved in the configuration storage of your workspace and are part of any workspace export operation.

### 14.11.1. Command goals

A developer executes a command to achieve a particular step in the development workflow. CodeReady Workspaces provides the ability to organize commands per goal:

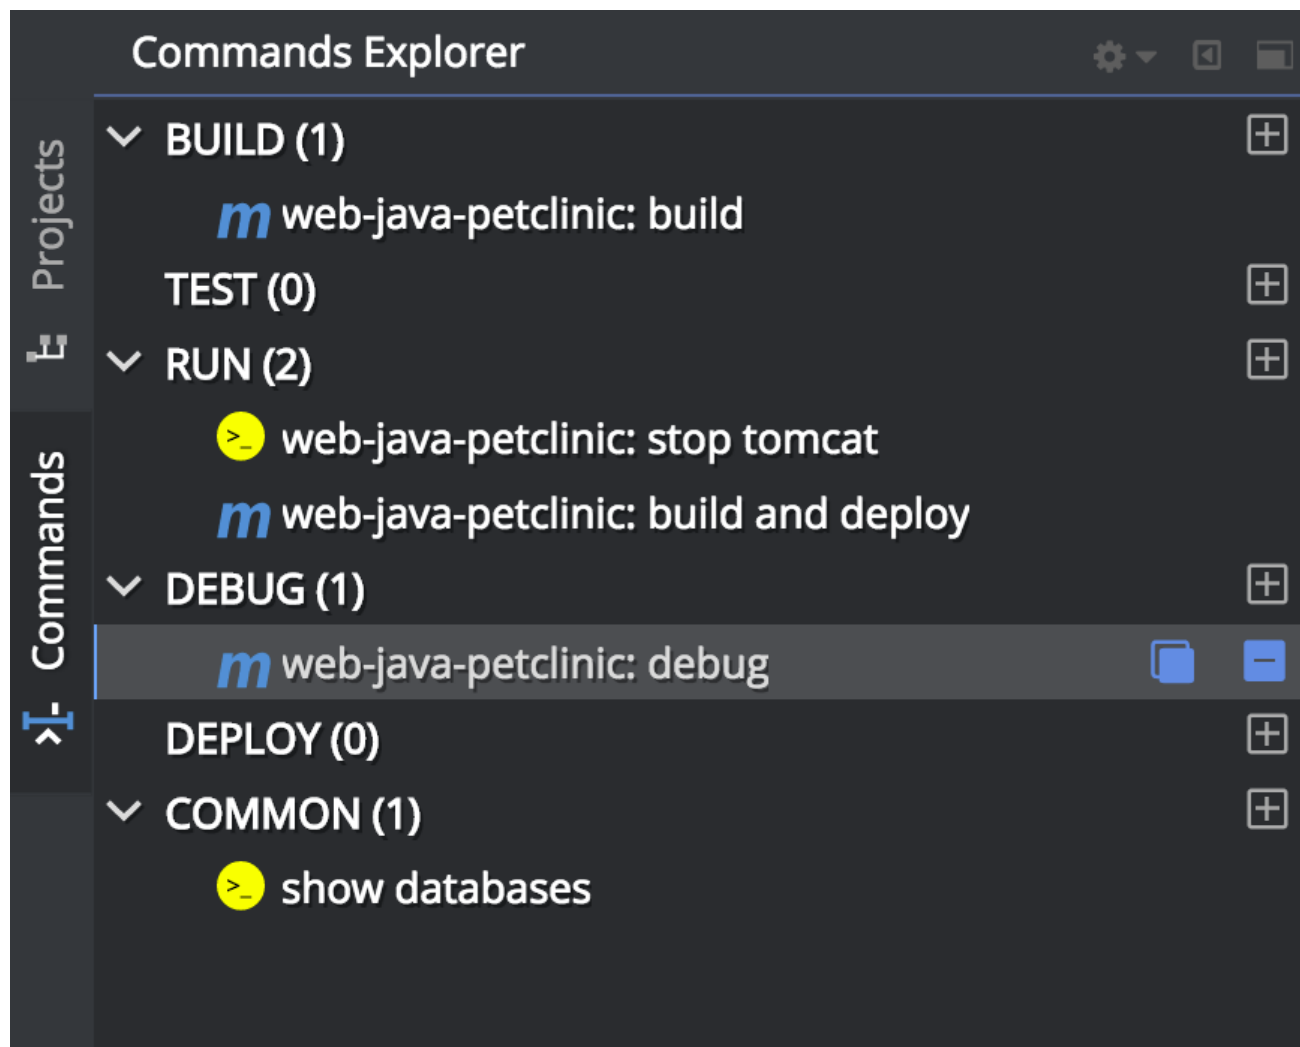
- *Build*: Commands that build the projects in a workspace.
- *Test*: Commands related to test execution.
- *Run*: Commands that run projects in a workspace.
- *Debug*: Commands used to start a debugging session.
- *Deploy*: Commands that are used to deploy projects in a workspace on specific servers or services.
- *Common*: Commands used for general purpose.

### 14.11.2. Command context


The context of a command defines the projects that the command can be used with. For example: a **maven build** command is relevant only if the project uses Maven.

### 14.11.3. Managing commands

Workspace commands are available in the **Commands Explorer** view. This view is accessible from the left pane. The commands are organized in this view by goal.



To create new commands, use the + icon next to a goal. Alternatively, you can select a command from the tree to edit, duplicate, or delete it.



The command editor is a distinct tab in the existing editor pane. Double-click the tab to use the full-screen edit mode. The editor can be split vertically or horizontally to display multiple editors at once.

- **Name:** A unique command name in your workspace. The name is not restricted to CamelCase.
- **Instructions:** To learn more about instructions and macros, see the [Macros](#) section.
- **Goal:** Use the dropdown list to change the goal of the command.
- **Context:** By default, the command is available with all the workspace projects. You can scope the command to be available only for selected projects.
- **Preview:** To learn more about servers, see the [Servers](#) section.

CodeReady Workspaces provides macros that can be used within a command or preview URL to reference workspace objects. To learn more, see the [Macros](#) section.

#### 14.11.4. Macros list

When editing a command, you can access all the macros that can be used in the command's instructions or in the preview URL. To display a complete list of macros, click on the **Macros** link.



Macro	Description
<code>\${current.class.fqn}</code>	The fully qualified name of the Java class currently active in the editor
<code>\${current.project.path}</code>	Absolute path to the project or module currently selected in the project explorer tree
<code>\${current.project.relpath}</code>	The path to the currently selected project relative to <code>/projects</code> . Effectively removes the <code>/projects</code> path from any project reference
<code>\${editor.current.file.name}</code>	Currently selected file in editor
<code>\${editor.current.file.path}</code>	Absolute path to the selected file in editor
<code>\${editor.current.file.relpath}</code>	Path relative to the <code>/projects</code> folder to the selected file in editor
<code>\${editor.current.project.name}</code>	Project name of the file currently selected in editor
<code>\${editor.current.project.type}</code>	Project type of the file currently selected in editor
<code>\${explorer.current.file.name}</code>	Currently selected file in project tree
<code>\${explorer.current.file.parent.path}</code>	Absolute path to the selected file's parent in project tree
<code>\${explorer.current.file.path}</code>	Absolute path to the selected file in project tree
<code>\${explorer.current.file.relpath}</code>	Path relative to the <code>/projects</code> folder in project tree
<code>\${explorer.current.project.name}</code>	Project name of the file currently selected in explorer
<code>\${explorer.current.project.type}</code>	Project type of the file currently selected in explorer
<code>\$(java.main.class)</code>	Path to the Main class
<code>\$(machine.dev.hostname)</code>	Dev-machine's host name

### 14.11.5. Auto-completing macros

You can auto-complete all macros used in the editor. To activate the auto-complete feature, press **Ctrl+Space**. This displays a menu listing all the possible macros based on what is typed.

The screenshot shows an IDE window with a terminal titled "web-java-petclinic: build and deploy". The command line contains the following text:

```
1 mvn -f /projects/web-java-petclinic clean install -DskipTests
2 cp /projects/web-java-petclinic/target/*.war $TOMCAT_HOME/webapps/ROOT.war
3 $TOMCAT_HOME/bin/catalina.sh run 2>&1
4 {current.}
```

A dropdown menu titled "Proposals:" is visible, listing the following macros:

- `$(current.class.fqn)`
- `$(current.project.path)`
- `$(current.project.relpath)`
- `$(editor.current.file.name)`
- `$(editor.current.file.path)`
- `$(editor.current.file.relpath)`
- `$(editor.current.project.name)`
- `$(editor.current.project.type)`
- `$(explorer.current.file.name)`
- `$(explorer.current.file.parent.path)`

A tooltip for the first macro, `$(current.class.fqn)`, is displayed, stating: "The fully qualified name of the Java class currently active in the editor".

### 14.11.6. Using commands

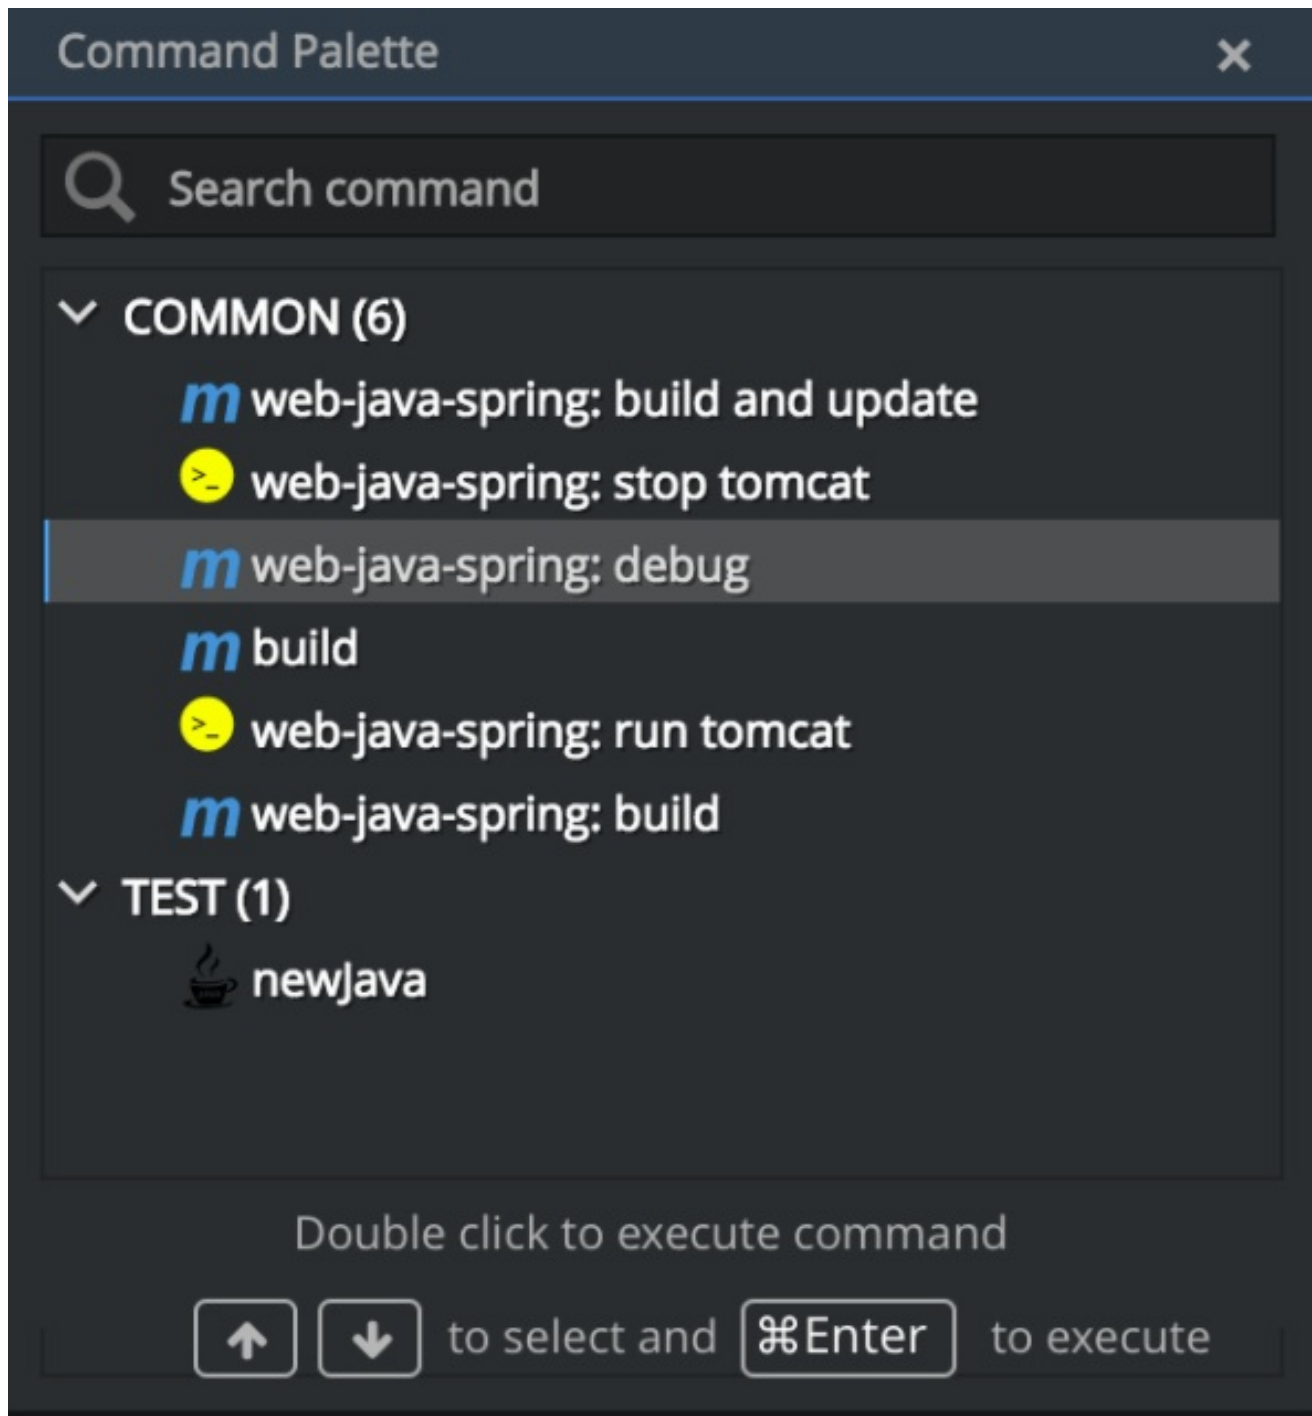
You can use commands from the following widgets:

- **Command Palette**

- Command toolbar
- Contextual menu in the **Project Explorer** view

#### 14.11.6.1. Using the Command Palette

Use the **Shift+F10** keys to open the **Command Palette**. Use the arrow keys to navigate through the list of commands. Use **Enter** to execute a selected command.



#### 14.11.6.2. Using the command toolbar

Using the command toolbar, you can execute the most common **Run** and **Debug** goals. The toolbar provides access to all the executed commands and previews.



#### 14.11.6.2.1. Using the Run and Debug buttons

To trigger commands directly from the **Run** and **Debug** buttons, you must define the commands for the **Run** and **Debug** goals.


You are asked to choose the default command associated with a button when you use this button for the first time and this button has multiple commands defined for a goal. Consecutive clicks of this button trigger the previously selected command.

To select a command to execute from a goal, click and hold the button. This command becomes the default command associated with the button.

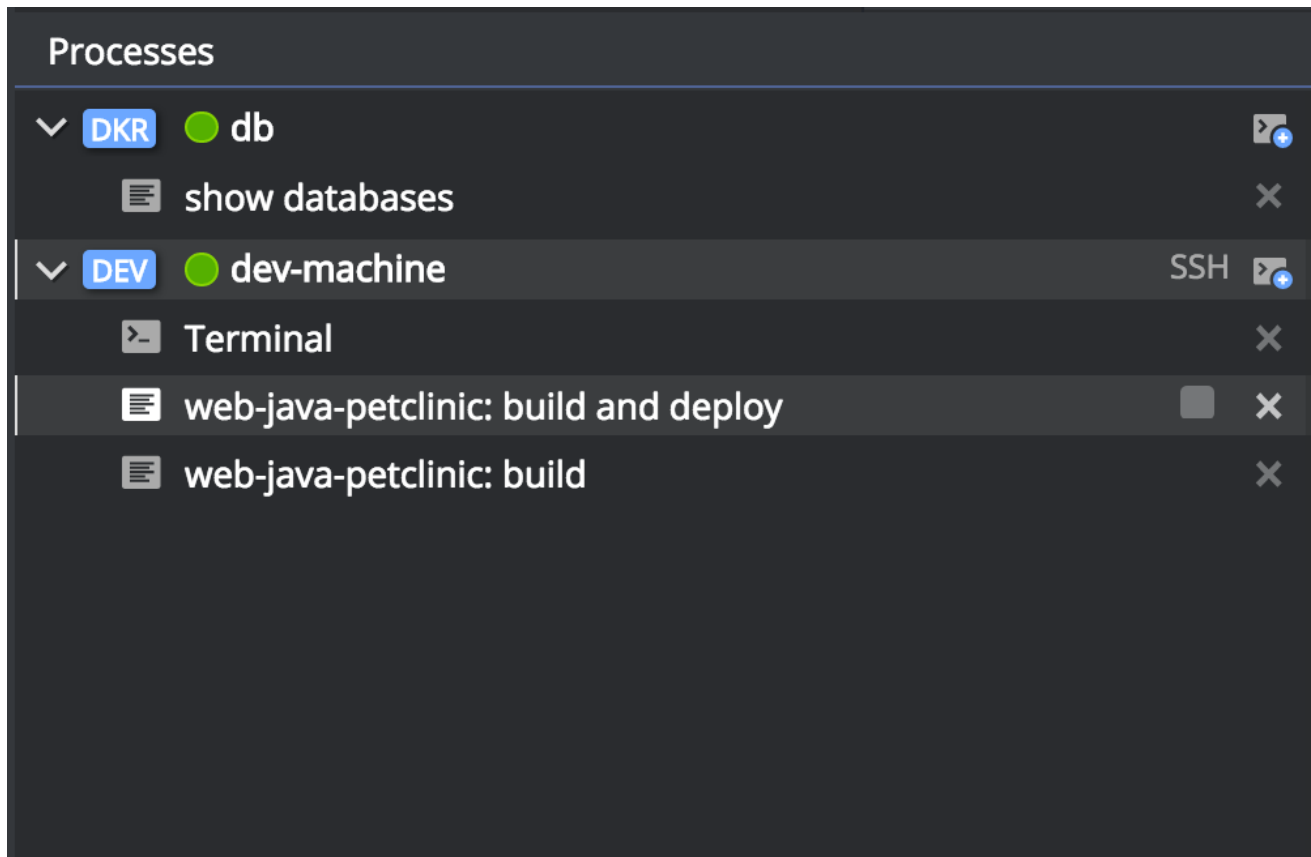
#### 14.11.6.2.2. Using the command controller

Use the command controller to view the state of the workspace and the last command that was executed. You can see the time span since a command was started and decide if it should be stopped or relaunched.

To view the list of all the previously executed commands when multiple commands are executed, click on the widget.



To clean the list, remove the process of the command from the list of processes.



#### 14.11.6.2.3. Using the Preview button

For commands used to start servers, define the preview URL to access the running server. To learn more, see the [Preview URLs](#) section.

The **Preview** button provides quick access to all of the servers that are running in the machines of the workspace.

#### 14.11.7. Authoring command instructions

A command may contain a single instruction or a succession of commands. For example:

```
cd /projects/spring
mvn clean install
```

- 1 Each command starts on a new line

```
cd /projects/spring; mvn clean install
```

- 1 A succession of several commands where ; stands for a new line

```
cd /projects/spring && mvn clean install
```

- 1 A succession of several commands where execution of a subsequent command depends on the execution of the preceding one. If the **/projects/spring** directory is absent, the **mvn clean install** command is not executed.

To check for conditions, use loops and other **bash** syntax:

```
mvn -f ${current.project.path} clean install 1
if [[ $? -eq 0 ]]; then
  cp /projects/kitchensink/target/*.war /home/user/wildfly-
10.0.0.Beta2/standalone/deployments/ROOT.war
  echo "BUILD ARTIFACT SUCCESSFULLY DEPLOYED..."
else
  echo "FAILED TO DEPLOY NEW ARTIFACT DUE TO BUILD FAILURE..."
fi
```

**1** Copy build artifact only if build is successful.

### 14.11.8. Macros

CodeReady Workspaces provides macros that can be used within a command or preview URL to reference workspace objects. Macros are translated into real values only when used in the IDE.



#### NOTE

You cannot use macros in commands that are launched from the server side.

The following table lists the macros and their descriptions.

Macro	Details
<code>\${current.project.path}</code>	Absolute path to the project or module currently selected in the <b>Project Explorer</b> tree.
<code>\${current.project.eldest.parent.path}</code>	Absolute path to a project root (for example, in the Maven-multi-module project).
<code>\${current.class.fqn}</code>	The fully qualified <b>package.class</b> name of the Java class currently active in the editor panel.
<code>\${current.project.relpath}</code>	The path to the currently selected project relative to <b>/projects</b> . Effectively removes the <b>/projects</b> path from any project reference.
<code>\${editor.current.file.name}</code>	Currently selected file in the editor.
<code>\${editor.current.file.basename}</code>	Currently selected file in the editor without extension.
<code>\${editor.current.file.path}</code>	Absolute path to the selected file in the editor.
<code>\${editor.current.file.relpath}</code>	Path relative to the <b>/projects</b> directory to the selected file in editor.

Macro	Details
<code>\${editor.current.project.name}</code>	Project name of the file currently selected in the editor.
<code>\${editor.current.project.type}</code>	Project type of the file currently selected in the editor.
<code>\${explorer.current.file.name}</code>	Currently selected file in the project tree.
<code>\${explorer.current.file.basename}</code>	Currently selected file in the project tree without extension.
<code>\${explorer.current.file.path}</code>	Absolute path to the selected file in the project tree.
<code>\${explorer.current.file.relpath}</code>	Path relative to the <b>/projects</b> directory in the project tree.
<code>\${explorer.current.project.name}</code>	Project name of the file currently selected in the explorer.
<code>\${java.main.class}</code>	Path to the main class.
<code>\${machine.dev.hostname}</code>	Current machine host name.
<code>\${project.java.classpath}</code>	Project classpath.
<code>\${project.java.output.dir}</code>	Path to the Java project output directory.
<code>\${project.java.sourcepath}</code>	Path to the Java project source directory.
<code>\${explorer.current.project.type}</code>	Project type of the file currently selected in the explorer.
<code>\${server.&lt;serverName&gt;}</code>	Returns protocol, hostname, and port of an internal server. <b>&lt;port&gt;</b> is defined by the same internal port of the internal service that you have exposed in your workspace recipe.

- Returns the hostname and port of a service or application you launch inside a machine.
- The hostname resolves to the hostname or the IP address of the workspace machine. This name varies depending on where Docker is running and whether it is embedded within a VM.
- The port returns the Docker ephemeral port that you can give to your external clients to connect to your internal service. Docker uses ephemeral port mapping to expose a range of ports that your clients may use to connect to your internal service. This port mapping is dynamic. In case of OpenShift, a route is returned.

Macro	Details
<code>\${workspace.name}</code>	Returns the name of the workspace.
<code>\${workspace.namespace}</code>	Workspace namespace (defaults to <b>che</b> in single-user CodeReady Workspaces).

### 14.11.9. Environment variables

The workspace machine has a set of system environment variables that have been exported. They are reachable from your command scripts using the **bash** syntax.

```
export ❶
$TOMCAT_HOME/bin/catalina.sh run ❷
```

- ❶ List of all the available machine system environment variables.
- ❷ Reference an environment variable, where **\$TOMCAT\_HOME** points to the **/home/user/tomcat8** directory.

## 14.12. STACKS

### 14.12.1. Stack overview

A stack is a [workspace configuration](#) template. Stacks are used to create workspaces in the **User Dashboard**. The stack includes meta-information such as scope, tags, components, description, name, and identification. You can filter stacks by machine type and scope. The type is either single machine or multi machine. You can also search for a stack by keyword. Stacks are displayed in the **User Dashboard** on the **Create a workspace** page.

See the [Creating and starting workspaces user guide](#) for more information.

### 14.12.2. Importing community supported stacks and applications

CodeReady Workspaces includes some stacks and sample applications that are pre-configured and tested. Stacks that are contributed by the CodeReady Workspaces community are not tested. Community stacks are located in the [community stacks GitHub repository](#).

Each directory has **\${technology}-stack.json** and **\${technology}-samples.json**.

To import a stack, follow these steps:

1. Copy the content of the JSON files.
2. Go to **\${CHE\_HOST}/swagger/#!/stack/createStack**.
3. Paste the content of the JSON file to the **body** field.
4. Click the **Try it out** button.

You can choose a different name or ID when there is a conflict with the stack ID or name.

For a multi-user setup, you can make your stack available for a particular user or all users in the system. See [stack sharing](#) for more information.

To import sample applications, move **\*-stacks.json** files to:

- **`\${LOCAL\_STORAGE}/instance/data/templates** for Docker infrastructure.
- **`\${mount-path-of-che-data-volume}/templates** for OpenShift and Kubernetes infrastructure. You need administrator privileges to get the host path and to access the host directory. Also, the new JSON files have the same permissions as the original **samples.json** file.

You can find Dockerfiles for all stacks in the [CodeReady Workspaces Dockerfiles repository](#).

### 14.12.3. Sharing stacks and system stacks

You can share stacks with selected users or with all users in the system if you have system privileges. You share stacks by making REST calls.

To share stacks with users:

- Log in as administrator
- Go to **/swagger/#!/stack/searchStacks** to get a list of all stacks. You may filter search by tags.
- Find your stack by name and get its ID.
- The next API to use is: **/swagger/#!/permissions**
- Find the below POST method:

**permissions** Show/Hide | List Operations | Expand Operations

**GET** /permissions Get all supported domains or only requested if domain parameter specified

**OPTIONS** /permissions

**POST** /permissions Store given permissions

**Parameters**

Parameter	Value	Description	Parameter Type	Data Type
<b>body</b>	<pre>{   "userId": "*",   "domainId": "stack",   "instanceId": "stacksuebxdwzzijkn",   "actions": [     "read",     "search"   ] }</pre>	The permissions to store	body	Model   Model Schema

```
{
  "instanceId": "string",
  "domainId": "string",
  "userId": "string",
  "actions": [
    "string"
  ]
}
```

Click to set as parameter value

Parameter content type:

**Response Messages**

HTTP Status Code	Reason	Response Model	Headers
200	The permissions successfully stored		

- Use the following JSON file and replace **`\${STACK\_ID}** with an actual ID:



```
{
  "userId": "*",
  "domainId": "stack",
  "instanceId": "${STACK_ID}",
  "actions": [
    "read",
    "search"
  ]
}
```

If you get 204, all the users in the system see the stack. To share a stack with a particular user, get the user's ID and use it instead of \* in the JSON file.

The administrator can remove pre-configured stacks and replace them with custom stacks. The administrator can also remove permissions from stacks. You can create stacks either in the user dashboard or by using any REST client. You can use Swagger (`$CHE_HOST:$CHE_PORT/swagger`) to bundle with CodeReady Workspaces.

#### 14.12.4. Loading stacks

Stacks are loaded from a JSON file that is packaged into resources of a special component that is deployed with the workspace master. This JSON file is not exposed to users. You can perform stack management using REST APIs in the **User Dashboard**.

When a user first starts CodeReady Workspaces, stacks are loaded from a JSON file only when the database is initialized. This is the default policy that can be changed. To keep getting stack updates with the new CodeReady Workspaces stacks, set **CHE\_PREDEFINED\_STACKS\_RELOADONSTART=true** in **che.env**. When set to **true**, **stacks.json** is used to update CodeReady Workspaces database each time the CodeReady Workspaces server starts. This means CodeReady Workspaces gets all the stacks in **stacks.json** and uploads the stacks to a database. This allows you to keep existing custom stacks and get stack updates from new CodeReady Workspaces releases. New and edited stacks that have fixes in the stack definition are merged in with the other stacks.

Name conflicts are possible. A name conflict happens when a new CodeReady Workspaces version provides a stack with a name that already exists in the database.

#### 14.12.5. Creating stacks in CodeReady Workspaces

Every stack is based on a container image. The image is used in an OpenShift deployment when a workspace is started. The resulting container in a pod is used both as build and runtime for a user application. It is also used in CodeReady Workspaces agents that are installers that activate the terminal, workspace agent, and language servers.

Because all agents have their dependencies, the underlying images must have those dependencies available.

- A custom stack requires Java 8 and the **oc** client.
  - For details to download Red Hat OpenJDK, see [Red Hat OpenJDK download page](#).
  - For details to download the **oc** client, see [Section 2.1, "Downloading the Red Hat OpenShift Origin Client Tools"](#).
- If you want to add the Analytics language server to the stack, Node.js 10 and Python 3 are also required.

Agents are injected in the running containers. Therefore, the current container user must have write access to the `~/che` directory. This is also a requirement for an image that can be used in a workspace stack definition.

### 14.12.5.1. Choosing a base image

For a base image of your custom stack, you can:

- Inherit an image from one of the certified images
- Use an existing Dockerfile or a container image

#### 14.12.5.1.1. Using a certified CodeReady Workspaces image

To create a custom image and to satisfy all the CodeReady Workspaces agent dependencies, inherit the image from one of the certified CodeReady Workspaces images that are used in the provided stacks.

#### Example Dockerfile using a certified CodeReady Workspaces image

```
FROM registry.redhat.io/codeready-workspaces/stacks-java-rhel8
RUN sudo dnf install some-software
```

Table 14.1. List of certified CodeReady Workspaces images

Language	Image
Java	<a href="https://registry.redhat.io/codeready-workspaces/stacks-java-rhel8">registry.redhat.io/codeready-workspaces/stacks-java-rhel8</a>
C++	<a href="https://registry.redhat.io/codeready-workspaces/stacks-cpp-rhel8">registry.redhat.io/codeready-workspaces/stacks-cpp-rhel8</a>
.NET	<a href="https://registry.redhat.io/codeready-workspaces/stacks-dotnet-rhel8">registry.redhat.io/codeready-workspaces/stacks-dotnet-rhel8</a>
Golang	<a href="https://registry.redhat.io/codeready-workspaces/stacks-golang-rhel8">registry.redhat.io/codeready-workspaces/stacks-golang-rhel8</a>
Node.js	<a href="https://registry.redhat.io/codeready-workspaces/stacks-node-rhel8">registry.redhat.io/codeready-workspaces/stacks-node-rhel8</a>
PHP	<a href="https://registry.redhat.io/codeready-workspaces/stacks-php-rhel8">registry.redhat.io/codeready-workspaces/stacks-php-rhel8</a>
Python	<a href="https://registry.redhat.io/codeready-workspaces/stacks-python-rhel8">registry.redhat.io/codeready-workspaces/stacks-python-rhel8</a>



#### NOTE

The images are available in the **recipe** blocks of the [stacks.json](#) file in the [Red Hat CodeReady Workspaces repository](#).

#### 14.12.5.1.2. Using a custom image

To use your own Dockerfile or a container image for a custom stack, ensure that you modify the Dockerfile, so that the image meets the following requirements:

- Includes OpenJDK 1.8 or newer: Even for Node.js or PHP images, Java is required (the workspace agent is a Tomcat server that needs Java).
- Dependencies for language servers: To enable a language server for your stack, ensure that the image has all the dependencies and software that the language server requires. To view the install scripts that agents use, see the particular page in the [Eclipse Che repository](#). For example, a JSON language server requires [Node.js](#).
- Installed language servers: Use the following commands to install the language server:

```
curl -s ${AGENT_BINARIES_URI} | tar xzf - -C ${LS_DIR}

touch ${LS_LAUNCHER}
chmod +x ${LS_LAUNCHER}
echo "nodejs ${LS_DIR}/vscode-json-server/out/jsonServerMain.js --stdio" >
${LS_LAUNCHER}
```

- Write access to the `/home/jboss/` and `/projects/` directories and to the `/etc/passwd` and `/etc/group` files: The user's home directory must be writable for any user. By default, all containers in OpenShift are run with arbitrary users that do not have sudo privileges and write access to most of the directories in the container. To give users sudo privileges and write access, see, for example, lines 76 and 77 of the [stacks-java-rhel8](#) Dockerfile (select the **Dockerfile** tab to see the code). Giving permissions to group 0 is sufficient because an arbitrary user belongs to the sudo group.
- Non-terminating **CMD** instruction: CodeReady Workspaces workspace master creates a deployment and waits for a pod to acquire a `RUNNING` state. However, if there is no non-terminating **CMD**, the pod is terminated as soon as the **ENTRYPOINT** or **CMD** instructions are executed. Therefore, a non-terminating **CMD** instruction is added to all images (for example, **CMD tail -f /dev/null**).

### Example Dockerfile using a custom image

```
FROM <myregistry/myimage>
USER root
RUN dnf install -y java-1.8.0-openjdk-devel [some-software]
CMD tail -f /dev/null
```

## 14.12.5.2. Building a custom stack

### 14.12.5.2.1. Building a Docker image

#### Procedure

To build a Docker image, see the [docker build](#) documentation.

### 14.12.5.2.2. Uploading an image to the registry

#### Procedure

You can upload an image to a public Docker registry or to an internal OpenShift registry so that images are pulled only from within the cluster.

### 14.12.5.2.3. Creating a custom stack

## Procedure

For detailed steps to create a custom stack, see the [Duplicate an existing stack](#) and the [Creating a custom stack](#) sections in the [Eclipse Che documentation chapter on managing workspaces](#).



### NOTE

When duplicating an existing stack, ensure to use your custom image and add or remove the agents as required by your stack.

### 14.12.5.3. Sharing stacks

#### Procedure

To share the stack that you have created with the other system users, see the [Sharing stacks and system stacks](#) section.

## 14.13. RECIPES

### 14.13.1. Supported Recipe Formats

Depending on the infrastructure, CodeReady Workspaces supports the following default recipe formats:

Infrastructure	Docker-formatted container image	Dockerfile	Composefile	Kubernetes YAML
Docker	Supported	Supported	Supported	Not supported
OpenShift	Supported	Not supported	Not supported	Supported

### 14.13.2. Docker-formatted container image requirements and limitations

The Docker-formatted container image recipe pulls an image from a Docker registry or uses the local image. The recipe then runs the image and creates a pod that references this image in the container specification. The following are Docker-formatted container image requirements and limitations for a workspace machine:

1. Use a non-terminating **CMD** or **ENTRYPOINT**. For a custom image, use, for example, **tail -f /dev/null** as one of the main processes.
2. For OpenShift only:
  - Do not use any processes and operations with **sudo** in **CMD**. See [Enable SSH and sudo](#) for more information.
  - Use [CodeReady Workspaces base stacks](#). You can also build your own image, but inherit from one of the base stacks.

### 14.13.3. Dockerfile definition and limitations

A Dockerfile is a set of instructions that Docker performs to build an image. After you provide a Dockerfile for your workspace machine, CodeReady Workspaces initiates a Docker build and runs the resulting image. The following are the limitations:

1. The **COPY** and **ADD** instructions fail because there is no context in **docker build**.
2. To avoid long build times with long Dockerfiles, build your image locally, push it to DockerHub, and then use the pushed image as a Docker-formatted container image recipe type. The start timeout for a workspace is five minutes.

#### 14.13.4. Running multi-container workspaces using Compose files

You can run multi-container workspaces using Compose files on Docker. The following syntax is not supported: **Local "build.context" and "build.dockerfile"**.

Because workspaces can be distributed, you cannot have host-local build and Dockerfile contexts. You can remotely host these aspects in a Git repository. CodeReady Workspaces sources the Compose file from the remote system and uses it as the build context.

You can run into a failure when the Dockerfile or build context requires you to **ADD** or **COPY** other files into the image. The local workspace generator cannot access these remote files.

##### 14.13.4.1. Accessing remote files

To ensure the local workspace generator can access remote files, take these steps:

1. Pre-package the build context or Dockerfile into an image.
2. Push that image into a registry.
3. Reference the pre-built image in your Compose file.

The following is an example of a remote context that works:

```
build:
  ## remote context will work
  context: https://github.com/eclipse/che-dockerfiles.git#master:recipes/stack-base/ubuntu

  ## local context will not work
  context: ./my/local/filesystem
```

##### 14.13.4.2. Using private repositories

To use private repositories in a remote build context:

1. Set up the SSH keys on your host machine.
2. Add the remote repository hostname or IP to the list of known hosts.

The following is an example of a YAML file using a private repository:

```
## The following will use master branch and build in recipes/stack-base/ubuntu folder
build:
  context: git@github.com:eclipse/che-dockerfiles.git#master:recipes/stack-base/ubuntu
```

##### 14.13.4.3. Configuring privileged access

The **privileged** Compose option does not support securing the underlying host system.

To configure the CodeReady Workspaces server to give all containers privileged access, set the **CHE\_PROPERTY\_machine\_docker\_privilege\_\_mode** variable to **true**.



### IMPORTANT

Setting the **CHE\_PROPERTY\_machine\_docker\_privilege\_\_mode** variable to **true** makes the host system vulnerable and gives all containers access to the host system.

#### 14.13.4.4. Special considerations when using Compose files

##### Build images

When a Compose file includes both build instructions and a build image, the build instructions override the build image, if it exists.

##### Container names

The **container\_name** is skipped during execution. Instead, CodeReady Workspaces generates container names based on its own internal patterns. Avoid naming conflicts. Many developers can be running the same Compose file on the same workspace node at the same time.

The following is an example of a YAML file using a container name:

```
container_name: my_container
```

##### Volumes

To define volumes for workspace machines, see [Volumes](#). Volume instructions in a Compose file are not supported.

##### Networks

CodeReady Workspaces does not support Compose networks. The use of aliases is supported by the **links** command.

The following is an example of a YAML file using networks:

```
## Not supported
networks:
  internal:
    aliases: ['my.alias']
## Not supported
networks:
  internal:
    driver: bridge
```

##### Hostname

Hostname is not supported. The machine's name is used for the hostname. You can use **links** aliases syntax to add additional hostnames to a machine.

##### Binding ports

Binding ports to the host system is not supported to ensure that containers do not use already assigned host ports. Users can work around this limitation by adding [servers](#) to machines.

## Environment file

The **env\_file** Compose option is not supported. Environment variables can be manually entered in the Compose file or machine configuration. See [Environment variables](#) for more information.

### 14.13.5. Kubernetes YAML limitations and restrictions

When a workspace is starting, CodeReady Workspaces creates various Kubernetes resources to support the IDE and development tools. Workspaces primarily consist of a [Deployment](#) which runs a [Kubernetes pod](#). The following are limitations and restrictions:

1. CodeReady Workspaces allows users specify Pods, Deployments, ConfigMaps, and Services in recipes
  - If a Pod is specified, it will be wrapped in a simple Deployment when running the workspace
2. Other object kinds will be ignored (PVC and route) or a workspace fails to start with an exception from Kubernetes.
3. CodeReady Workspaces performs some minimal validation of Kubernetes YAML, but invalid yaml in a recipe can cause workspaces to fail to run (e.g. referring to a non-existent configmap)
4. You cannot use volumes in the container and pod definition. See [Volumes](#) for information about persisting and sharing data between pods.

The following is an example of a custom recipe with two containers, a simple config map, one deployment, and a service that is bound to port 8081:

```
kind: List
items:
-
  apiVersion: apps/v1
  kind: Deployment
  metadata:
    name: my-deployment
  spec:
    replicas: 1
    selector:
      matchLabels:
        my-workspace-pod: dev
  template:
    metadata:
      name: dev-pod
    labels:
      my-workspace-pod: dev
    spec:
      containers:
      -
        image: eclipse/ubuntu_jdk8:latest
        name: main
        ports:
        -
          containerPort: 8081
          protocol: TCP
        env:
        -
```

```

      name: MY_ENV_VAR
      valueFrom:
        configMapKeyRef:
          name: my-configmap
          key: my-key
    -
      image: eclipse/ubuntu_jdk8:latest
      name: main1
  -
    apiVersion: v1
    kind: ConfigMap
    metadata:
      name: my-configmap
    data:
      my-key: my-value
  -
    kind: Service
    apiVersion: v1
    metadata:
      name: my-service
    spec:
      selector:
        name: app
      ports:
      - protocol: TCP
        port: 8081
        targetPort: 8081

```

As a bare minimum, a Kubernetes YAML recipe must contain at least one Pod or Deployment, in which the main dev machine is run.

You can also specify multiple containers within the workspace pod. CodeReady Workspaces treats those containers as workspace machines. These containers can have machine names defined in annotations.

**PodName/Container Name** is the default naming pattern for a machine.

The following is an example of using annotations:

```

kind: List
items:
-
  apiVersion: v1
  kind: Pod
  metadata:
    name: any123123
    annotations:
      org.eclipse.che.container.main.machine_name: myMachine
      org.eclipse.che.container.main1.machine_name: myMachine1
  spec:
    containers:
    -
      image: rhche/spring-boot:latest
      name: main
      ports:
      -
        containerPort: 8080
        protocol: TCP

```



```
resources: {}  
  
-  
  image: rhche/spring-boot:latest  
  name: main1  
  ports:  
  -  
    containerPort: 8080  
    protocol: TCP  
  resources: {}
```

## 14.14. SERVERS

### 14.14.1. What are servers?

A server defines the protocol port of a process that runs in a machine. It has a name, path, and attributes. The path defines the base path of the service that is used by the server. Attributes are optional and can be used to tune the server or for identification. You can add a server when you need to access a process in your workspace machine.

To add a server, use the **User Dashboard** or edit the workspace machine configuration YAML file.

The following is an example of the YAML file:

```
"node": {  
  "port": "3000",  
  "protocol": "http",  
  "path": "/",  
  "attributes": {}  
}
```

The following is an example of the **User Dashboard**:

The screenshot shows the 'Servers' tab in the Red Hat CodeReady Workspaces interface. A modal dialog titled 'Add a new server' is open, allowing the user to add a new server. The dialog has three input fields: 'Reference' with the value 'node', 'Port' with the value '3000' and a green checkmark, and 'Protocol' with the value 'http' and a green checkmark. At the bottom of the dialog are 'Add' and 'Close' buttons. In the background, a table lists existing servers:

KEY	PORT	PROTOCOL	RUNTIME
<input type="checkbox"/> codeserver	9876	http	
<input type="checkbox"/> tomcat8			
<input type="checkbox"/> tomcat8-debug			



## NOTE

If your workspace is running, saving a new server restarts the workspace.

### 14.14.2. Preview URLs


Adding the server with port 3000 does not mean you can use this port to access a server. Each server is assigned a URL when a workspace is running.

- On Docker, port 3000 is published to a random port from the ephemeral port range from 32768 to 65535. The server URLs change every time you start a workspace.
- On OpenShift, a route bound to a service is created. Routes are persistent URLs.

### 14.14.3. Getting preview URLs

In this example, you added a server with port 3000 and started a workspace. The following are ways to get the server's preview URL:

- Use a [macro](#) command.
- In the IDE, Click the **+** icon in the bottom panel under the editor.



- In the **User Dashboard**, click **Workspaces > YourWorkspace > Servers**

You can also see internal server URLs. Internal servers launch when the workspace container or pod is available.

#### 14.14.4. Exposing internal servers

To access a port internally within a workspace, expose it internally, but do not make it publicly accessible. For example, a database server is exposed only for the web application and because of security concerns, it should not be accessible publicly. The database server is exposed as internal.

To expose a server as internal, add the corresponding attribute into the server configuration YAML file:

```
"db": {
  "port": "3200",
  "protocol": "tcp",
  "attributes": {
    "internal": "true"
  }
}
```

The application is able to fetch the database URL from the workspace runtime and access the database. The database URL is not accessible publicly from the browser.

#### 14.14.5. Exposing secure servers

Secure servers are exposed publicly but access is restricted only for users who have permissions to the workspace. The authentication proxy is set up as the exposed server and the machine token is required to request it.

To expose a server as secure, add the corresponding attributes into the server configuration YAML file:

```
"tooling": {
  "port": "4921",
  "protocol": "http",
  "attributes": {
    "secure": "true",
    "unsecuredPaths": "/liveness",
    "cookiesAuthEnabled": "true"
  }
}
```

The following describes the attributes:

##### **secure**


Indicates whether the server is exposed as secure. The default value is **false**.

### **unsecuredPaths**

Configures the secure servers. It contains a comma-separated list of URLs that are considered non-secure on the given server and can be accessible without a token. It may be needed when the server provides any public APIs. The API endpoint for health checks is an example.

### **cookiesAuthEnabled**

Indicates whether cookies should be searched for a token. By default, it is disabled. You can enable this attribute if you are sure that servers cannot be attacked by Cross-Site Request Forgery (CSRF) or have special protection from it.



#### **NOTE**

This is in the beta phase and disabled by default. See [Secure servers](#) for information on how to enable secure servers.

## 14.15. INSTALLERS

### 14.15.1. What are installers?

Installers are scripts that are added into machines in a runtime. Once running, installers:

1. Prepare the environment and download dependencies for a particular software or tool.
2. Install chosen software and dependencies.
3. Launch software and tools with particular arguments and modes that provide extra functionality for a workspace.

Installers are typically language servers and tools that provide features such as SSH access to a workspace machine. You can find a complete list of available installers in the [Workspace details > Installers](#) tab.

The following is an example of installers:

```
"installers": [
  "org.eclipse.che.exec",
  "org.eclipse.che.terminal",
  "org.eclipse.che.ws-agent",
  "org.eclipse.che.ssh"
]
```

### 14.15.2. How installers work

Installers are saved in a configuration file that a [bootstrapper](#) uses to execute jobs. An installer script works exactly the same way as other shell scripts in Linux. The CodeReady Workspaces server checks if the launched process is running.

Some installers activate special agents, such as the workspace, terminal, and execution agents. If a workspace agent fails to start, the workspace is treated as if it has been started but the IDE is not usable. If the execution agent fails, the **commands** widget is unavailable.

### 14.15.3. What happens when enabling and disabling installers

You can enable or disable installers per machine by using the **User Dashboard** or by updating the workspace machine configuration. When an installer is enabled, the bootstrapper executes an installer script after the workspace has started.

The following shows installers that are enabled and disabled:

Installers Servers Env Variables Volumes Config SSH

Workspace Installers

STATE	INSTALLER	VERSION	DESCRIPTION
<input type="checkbox"/>	C# language server	1.0.1	C# Intellisense
<input checked="" type="checkbox"/>	Exec	1.0.0	Agent for command execution
<input type="checkbox"/>	File sync	1.0.0	Unison File Synchronizer
<input type="checkbox"/>	Git credentials	1.0.0	Agent fetches SSH keys, Git username and email from CHE user preferences, and injects to console Git
<input type="checkbox"/>	JSON language server	1.0.1	JSON Intellisense
<input type="checkbox"/>	PHP language server	2.0.1	PHP Intellisense
<input type="checkbox"/>	Python language server	1.0.3	Python Intellisense
<input type="checkbox"/>	Simple Test language server	1.0.0	Test LS
<input checked="" type="checkbox"/>	SSH	1.0.0	SSH server, key-pair generation
<input checked="" type="checkbox"/>	Terminal	1.0.0	Embedded web terminal
<input type="checkbox"/>	TypeScript language server	1.0.1	TypeScript Intellisense
<input type="checkbox"/>	Workspace API	1.0.1	Workspace API support
<input type="checkbox"/>	Workspace API	1.0.0	Workspace API support
<input type="checkbox"/>	Yaml language server	1.0.0	Yaml Intellisense

Make a wish Docs Community Diagnostics

## 14.15.4. Troubleshooting installer failures

### 14.15.4.1. Permission denied failure

Installers run as if a user in the container has **sudoers** privileges. If the user does not have the privileges, the installer fails with **permission denied** issues.

This problem can occur with OpenShift when a pod is run by a user with no permissions to use **sudo** or to access or modify resources on the file system.

In most cases, this problem can be solved by rebuilding the base image so that it already has all of the dependencies for particular agents that an installer activates.

### 14.15.4.2. Permission to files and directories failures

Another possible issue can be with permissions to files and directories. For example, an installer may need to write to the [user home directory](#).

## 14.15.5. Installer registry and REST API

CodeReady Workspaces installers are stored in the Installer Registry. They can be viewed and edited through the Installer Registry REST API:

Path	HTTP Method	Description
<b>/installer</b>	GET	Get installers
<b>/installer/{id}/version</b>	GET	Get versions of installers by given id
<b>/installer/orders</b>	GET	Get installers, ordered by their dependencies
<b>/installer/</b>	POST	Add installer to the registry
<b>/installer/{key}</b>	PUT	Update installer in the registry
<b>/installer/{key}</b>	DELETE	Remove installer from the registry

## 14.16. VOLUMES

### 14.16.1. Default volumes for workspace containers

By default, workspace containers start with a default volume and have a minimum of one PVC that is located in the **/projects** directory.

Workspace projects are physically located in the **/projects** directory. When a workspace stops, the machines are destroyed, but the volumes persist.

### 14.16.2. Adding volumes

In order for your data to persist for a local Maven repository, the **node\_modules/** directory, Ruby gems, or the **authorized\_keys** file for SSH connections, your workspace will need additional volumes. Each machine can add as many volumes as the underlying infrastructure can support. OpenShift may impose a limit on the number of volumes.

You can add volumes either by using the **User Dashboard** or by updating the machine configuration. The following is an example of the configuration file:

```
"volumes": {
  "myvolume": {
    "path": "/absolute/path/in/workspace"
  }
}
```

To avoid failures when updating the workspace configuration using REST APIs:

- Use an absolute path.
- The name and path cannot contain special characters, including dashes (-) or underscores (\_).

**NOTE**

To allow machines to share the same volume, create a volume for each machine with an identical name.

### 14.16.3. Configuring workspaces

To configure workspaces on the OpenShift and Kubernetes infrastructure as ephemeral, set the **persistVolumes** attribute to **false** in the workspace configuration.

The following is an example of the configuration file:

```
"attributes": {  
  "persistVolumes": "false"  
}
```

In this case, regardless of the PVC strategy, all volumes would be created as **emptyDir** for the given workspace. When a workspace pod is removed for any reason, the data in the **emptyDir** volume is deleted forever.

## 14.17. ENVIRONMENT VARIABLES

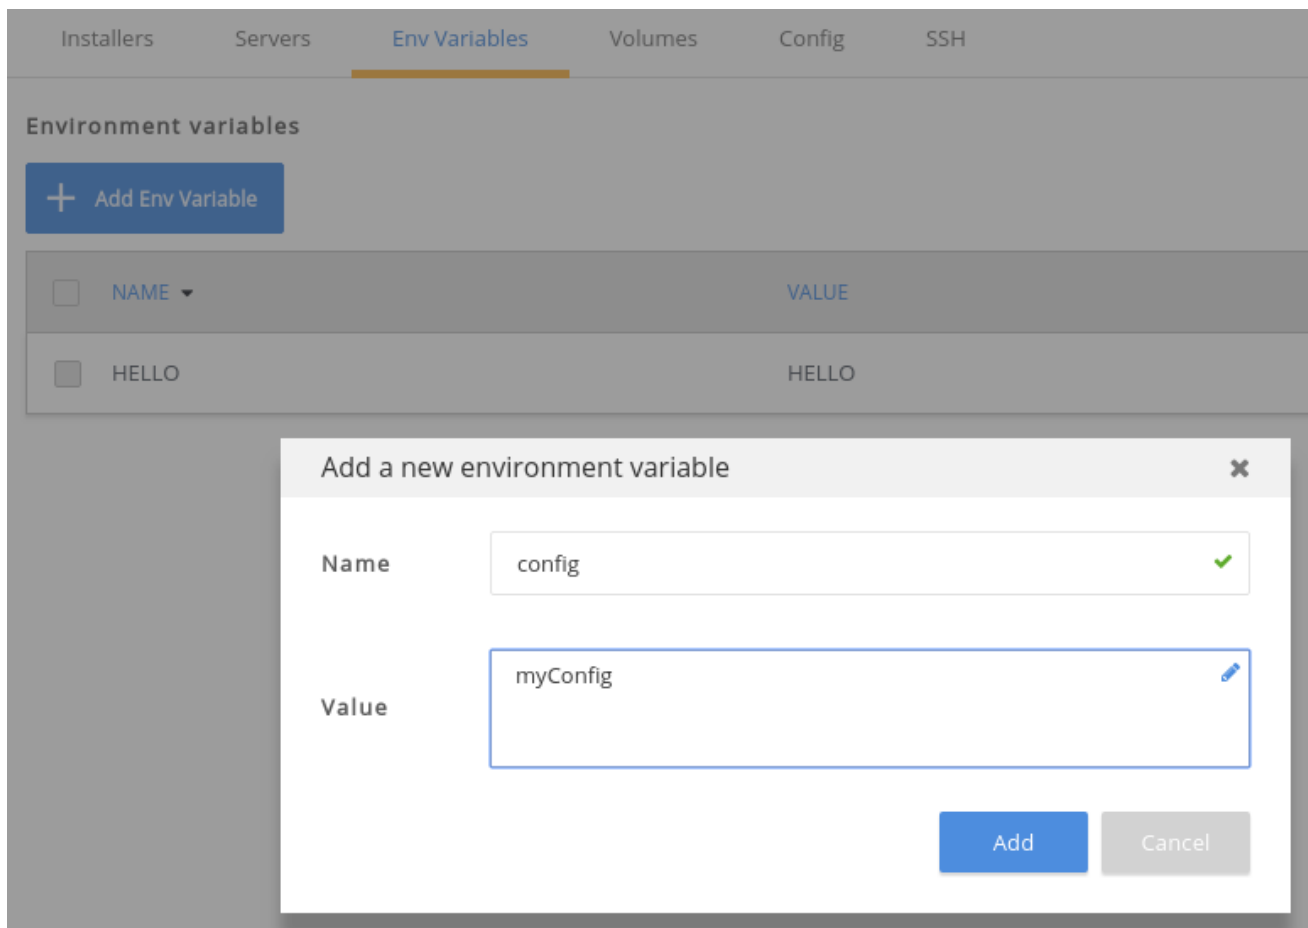
Environment variables are defined per machine. Depending on the infrastructure, they are added either to the container or the Kubernetes pod definition. You can add, edit, and remove environment variables either in the **User Dashboard** or directly in the workspace machine configuration.

The following is an example of an environment variable:

```
"env": {  
  "key": "value"  
}
```

You can use environment variables in applications running in a workspace, in [commands](#), and in the terminal. The CodeReady Workspaces server also adds some environment variables that a user does not control, although they are available to use. For example, they can be used as an API endpoint or workspace ID.

The following shows how to add a new environment variable:



## 14.18. PROJECTS

### 14.18.1. Creating projects in workspaces

Projects are always associated with a workspace and saved in a workspace configuration.

The following is an example of the project YAML file:

```
"projects": [
  {
    "description": "A basic example using Spring servlets. The app returns values entered into a
submit form.",
    "source": {
      "location": "https://github.com/che-samples/web-java-spring.git",
      "type": "git",
      "parameters": {}
    },
    "links": [],
    "mixins": [],
    "problems": [],
    "name": "web-java-spring",
    "type": "maven",
    "path": "/web-java-spring",
    "attributes": {}
  }
]
```




Once a project is saved into a workspace configuration, the IDE checks if the project exists on a file system. Use the **source.location** URL to import projects that do yet exist on the file system. This happens during the IDE initialization stage.

You can add the following projects:

- Git projects
- remotely hosted archives
- GitHub projects
- example projects provided by CodeReady Workspaces

Project import tools can be found on the **User Dashboard** when you are creating a new workspace or editing an existing workspace in the IDE. Project import tools can also be found in the **Workspace** menu.

The following shows example projects:



### 14.18.2. Defining project types

Plug-in developers can define their own project types. Since project types trigger certain behaviors within the IDE, the construction of the projects is important to understand.

- A project type is defined as one primary type and zero or more mixin types.
  - A primary project type is one where the project is editable, buildable, and runnable.
  - A mixin project type defines additional restrictions and behaviors of the project, but it cannot be a primary project type by itself.

The collection of primary and mixin types for a single project defines the aggregate set of attributes that will be stored as meta data within the project.

- Project types describe different aspects of a project, such as:
  - the types of source files inside
  - the structure of the explorer tree
  - the way in which a command is executed

- associated workflows
- which plug-ins must be installed
- A project defines a set of attributes. The attributes of a project can be mandatory or optional. Attributes that are optional can be dynamically set at runtime or during configuration.
- Sub-projects may have different project types than their parents. Modules may physically exist within the tree structure of the parent. For example, subdirectories exist within the tree structure of the parent. Also, modules may physically exist outside the tree structure of the parent, such as when the parent is a soft link to the module project.

### 14.18.3. Creating a sub-project

A sub-project is a portion of a project that can have sets of commands run against it where the sub-directory is treated as the root working directory. Sub-projects make it possible to organize a single repository into multiple, independently buildable, and runnable units.

To create a module, right-click on a directory in the IDE explorer tree and select **Convert to Project**. You can then execute commands directly against this sub-project.

### 14.18.4. Navigating the project tree

You can step into or out of the project tree. When you step into a directory, that directory is set as the project tree root and the explorer refreshes the view. All commands are then executed against this directory root.

## 14.19. TROUBLESHOOTING FAILURES IN STARTING THE WORKSPACE

Failures to start a workspace may be caused by the following factors:

- Incorrect environment recipe
- Restrictive network settings

### 14.19.1. Incorrect environment recipes

When a workspace is starting, an environment [recipe](#) is sent to Docker or to the OpenShift API. The CodeReady Workspaces server then listens to events provided by the given infrastructure. The CodeReady Workspaces server expects a running Docker container or an OpenShift pod. The server fails to start an environment and consequently the starting of the workspace fails if the infrastructure is unable to create and start a container or a pod from the provided recipe.

A recipe can be incorrect due to the following reasons:

- The Docker build fails with the provided **Dockerfile**. This can be because of a broken **Dockerfile** or because of CodeReady Workspaces. If the Docker build in CodeReady Workspaces does not support context, consider editing the Docker recipe locally to ensure that it is a valid **Dockerfile**. Add or copy resources into an image locally on your machine, push the image to a registry, such as DockerHub, and use the resulting images in the recipe.
- CodeReady Workspaces does not support certain Docker Compose syntax. Ensure that the **Composefile** is supported by CodeReady Workspaces.

- Installing packages in your **Dockerfile** instructions can take time. This may be influenced by network settings.

#### 14.19.1.1. Viewing logs from a failed workspace start

No **installer** logs are shown when a workspace fails to start because its container or pod are not launched. In most cases, only logs from infrastructure and image pull and build are shown. Analyse these logs to find the problem. The CodeReady Workspaces server also produces logs that are helpful in debugging the problem.

#### 14.19.2. Restrictive network settings

The CodeReady Workspaces server and agents, which run in a workspace container or pod, and the user's browser communicate with each other. Firewall, filtered ports, and other network restrictions may cause trouble when starting a workspace.

A workspace is considered to be in a **RUNNING** state after the CodeReady Workspaces server verifies that the workspace agent is up. The workspace agent also tries to reach the CodeReady Workspaces server. All this happens in separate containers or pods, and the user's browser is not yet involved. The **workspace started by user \$userName** message in the CodeReady Workspaces server logs indicates the following:

- The workspace container or pod is up.
- The workspace agent has successfully started.
- The CodeReady Workspaces server can reach it.

#### 14.19.2.1. Troubleshooting network setting when workspace agent cannot be reached

An error message saying that the IDE cannot be initialized indicates that the client (browser) cannot reach the workspace agent. This is caused by the CodeReady Workspaces server using an internal IP address to reach the workspace agent, while you are accessing the workspace from a machine that is located on a different network. To confirm this, open the browser developer console and check failed requests. The failed requests are to **project** and **project-type** API.

To access a workspace from a different network than the one on which the CodeReady Workspaces server is running, enable access to the ephemeral port range on the CodeReady Workspaces server network.

#### 14.19.3. Failure in bootstrapping

When a workspace starts, the CodeReady Workspaces server creates and starts a container or a pod or a set of containers and pods as per the environment recipe. After the container or pod is running, a bootstrapping process begins - the bootstrapper binary is downloaded and launched. If the server logs show bootstrapping failures, or you do not see any output in the **Machine** tab of the **Workspaces** view, the reason is that bootstrapper is not downloaded. The following are possible the reasons for the bootstrapper download failure:

- Network conditions (for example, firewall restrictions).
- Incorrect bootstrapper binary URL that the CodeReady Workspaces server uses (often reproduced when deploying to OpenShift and missing necessary environment variables).

To work around the problem, download the bootstrapper binary manually. On OpenShift, access the pod on the command line (shell or the terminal in the web console) and run the following commands:

```
$ cd /tmp/bootstrapper
$ ls -la 1
$ curl ${CHE_URL}/agent-binaries/linux_amd64/bootstrapper/bootstrapper
```

**1** to check for the existence of the bootstrapper binary

To prevent the **curl** command from failing, unblock port **80** on your network. On OpenShift with **https** routes, unblock port **443**.

## 14.20. WORKSPACE DATA MODEL

The following table lists the data types and their description.

Data Types	Description
environments: Map<String, getEnvironments>	Workspace environment variables. A workspace can have multiple environment variables.
defaultEnv: STRING	A workspace must have a default environment.
projects: []	Projects associated with a workspace.
commands: []	Commands associated with a workspace.
name: STRING	Workspace name that has to be unique in a namespace.
links: []	-

### 14.20.1. Environment recipes

<b>environments</b>	[-] Object, 1 property	
<b>default</b>	[-] Object, 2 properties	
<b>machines</b>	[-] Object, 1 property	
<b>dev-machine</b>	[-] Object, 5 properties	
<b>attributes</b>	[-] Object, 1 property	
	<b>memoryLimitBytes</b>	2147483648
<b>servers</b>	[-] Object data structure, 2 properties	
	<b>[key]</b>	<b>attributes</b>
	<b>tomcat8</b>	[-] Object, 1 property
		<b>port</b>
		<b>protocol</b>
	<b>key</b>	<b>value</b>
	<b>tomcat8-debug</b>	{ Empty Object }
	<b>port</b>	8080
	<b>protocol</b>	http
	<b>tomcat8-debug</b>	{ Empty Object }
	<b>port</b>	8000
	<b>protocol</b>	http
<b>volumes</b>	[-] Object data structure, 2 properties	
	<b>[key]</b>	<b>path</b>
	<b>myvolume</b>	/myfolder
	<b>myvolume1</b>	/myfolder1
<b>installers</b>	[-] Array, 3 items	
	<b>0</b>	org.eclipse.che.exec
	<b>1</b>	org.eclipse.che.terminal
	<b>2</b>	org.eclipse.che.ws-agent
<b>env</b>	[-] Object, 2 properties	
	<b>myenv1</b>	myvalue1
	<b>myenv</b>	myvalue
<b>recipe</b>	[-] Object, 2 properties	
	<b>location</b>	eclipse/ubuntu_jdk8
	<b>type</b>	dockerimage

For [recipe](#) types of **dockerfile**, **compose**, or **openshift**, content, not location, is specified.

```
"recipe": {
  "type": "compose",
  "content": "services:\n db:\n image: eclipse/mysql\n environment:\n
  MYSQL_ROOT_PASSWORD: password\n MYSQL_DATABASE: petclinic\n MYSQL_USER:
  petclinic\n MYSQL_PASSWORD: password\n mem_limit: 1073741824\n dev-machine:\n image:
  eclipse/ubuntu_jdk8\n mem_limit: 2147483648\n depends_on:\n - db",
  "contentType": "application/x-yaml"
}
```

## 14.20.2. Projects

<b>projects</b>	[-] Array, 1 item	
	<b>0</b>	[-] Object, 9 properties
	<b>links</b>	[ Empty Array ]
	<b>name</b>	console-java-simple
	<b>attributes</b>	{ Empty Object }
	<b>type</b>	maven
	<b>source</b>	[-] Object, 3 properties
		<b>location</b> https://github.com/che-samples/console-java-simple.git
		<b>type</b> git
		<b>parameters</b> [-] Object, 1 property
		<b>branch</b> che6
	<b>path</b>	/console-java-simple
<b>description</b>	A hello world Java application.	
<b>problems</b>	[ Empty Array ]	
<b>mixins</b>	[ Empty Array ]	

The project object structure has the **source.location** and **source.type** parameters. There are two importer types: **git** and **zip**. New location types can be provided by custom plugins, such as **svn**.

Incorrectly configured projects or projects missing sources are marked with error codes and messages explaining the error. In the example above, the project does not have errors and mixins.

A mixin adds additional behaviors to the project, the IDE panels, and menus. Mixins are reusable across any project type. To define the mixins to add to a project, specify an array of strings, with each string containing the identifier for the mixin.

Mixin ID	Description
<b>git</b>	Initiates the project with a Git repository. Adds Git-menu functionality to the IDE. To add a mixin to the project, create a new project and then initialize a Git repository.
<b>pullrequest</b>	Enables pull-request workflow where a server handles the local and remote branching, forking, and pull-request issuance. Pull requests generated from within the server have another Factory placed into the comments of pull requests that a PR reviewer can consume. Adds contribution panel to the IDE. Set this mixin to use attribute values for <b>project.attributes.local_branch</b> and <b>project.attributes.contribute_to_branch</b> .

The **pullrequest** mixin requires additional configuration from the **attributes** object of the project.

The **project** object can include **source.parameters**, which is a map that can contain additional parameters. Example: related to project importer.

Parameter name	Description
<b>skipFirstLevel</b>	Used for projects with type <b>zip</b> . When value is 'true', the first directory inside ZIP will be omitted.

### 14.20.3. Commands

Commands can be both tied to a workspace and an individual project. In the example below, a command is saved to workspace configuration.

<b>commands</b>	[ <b>-</b> ] Array, 1 item
<b>0</b>	[ <b>-</b> ] Object, 4 properties
<b>commandLine</b>	mvn clean install -f \${current.project.path}
<b>name</b>	build
<b>attributes</b>	[ <b>-</b> ] Object, 2 properties
<b>goal</b>	Build
<b>previewUrl</b>	[zero-length string]
<b>type</b>	mvn

The following image shows ways to save commands in the project configuration.

<b>projects</b>	[ <b>-</b> ] Array, 1 item
<b>0</b>	[ <b>-</b> ] Object, 8 properties
<b>links</b>	[ Empty Array ]
<b>name</b>	spring
<b>attributes</b>	[ <b>-</b> ] Object, 2 properties
<b>contribute_to_branch</b>	[ <b>-</b> ] Array, 1 item
<b>0</b>	master
<b>commands</b>	[ <b>-</b> ] Array, 5 items
<b>0</b>	{"commandLine": "mvn -f /projects/spring clean install \ncp /projects/s...
<b>1</b>	{"commandLine": "\$TOMCAT_HOME/bin/catalina.sh run 2>&1", "name": "spring...
<b>2</b>	{"commandLine": "\$TOMCAT_HOME/bin/catalina.sh stop", "name": "spring: st...
<b>3</b>	{"commandLine": "mvn -f /projects/spring clean install \ncp /projects/s...
<b>4</b>	{"commandLine": "mvn -f /projects/spring clean install \ncp /projects/s...
<b>type</b>	maven
<b>source</b>	[ <b>-</b> ] Object, 3 properties
<b>location</b>	https://github.com/che-samples/web-java-spring.git
<b>type</b>	git
<b>parameters</b>	{ Empty Object }
<b>path</b>	/spring
<b>problems</b>	[ Empty Array ]
<b>mixins</b>	[ <b>-</b> ] Array, 1 item
<b>0</b>	pullrequest

### 14.20.4. Runtime

A runtime object is created when a workspace is in a running state. Runtime returns server URLs, internal or external, depending on the server configuration. Interested clients, like the **User Dashboard** and the IDE, use these URLs.

<b>runtime</b>	[ - ] Object, 5 properties	
<b>machines</b>	[ - ] Object, 1 property	
<b>dev-machine</b>	[ - ] Object, 2 properties	
<b>attributes</b>	{ Empty Object }	
<b>servers</b>	[ - ] Object data structure, 10 properties	
<b>[key]</b>	<b>attributes</b>	<b>url</b>
<b>22tcp</b>	{ Empty Object }	tcp://dev-machine:22
<b>codeserver</b>	{ Empty Object }	http://172.19.20.180:32836
<b>exec-agent/http</b>	[ - ] Object, 3 properties	
<b>url</b>	http://172.19.20.180:32839/process	
<b>attributes</b>	{ Empty Object }	
<b>status</b>	RUNNING	
<b>exec-agent/ws</b>	{ Empty Object }	ws://172.19.20.180:32839/connect
<b>terminal</b>	[ - ] Object, 3 properties	
<b>url</b>	ws://172.19.20.180:32840/pty	
<b>attributes</b>	{ Empty Object }	
<b>status</b>	RUNNING	
<b>tomcat8</b>	{ Empty Object }	http://172.19.20.180:32837
<b>tomcat8-debug</b>	{ Empty Object }	http://172.19.20.180:32838
<b>wsagent-debug</b>	{ Empty Object }	tcp://172.19.20.180:32841
<b>wsagent/http</b>	[ - ] Object, 3 properties	
<b>url</b>	http://172.19.20.180:32842/api	
<b>attributes</b>	{ Empty Object }	
<b>status</b>	RUNNING	
<b>wsagent/ws</b>	{ Empty Object }	ws://172.19.20.180:32842/wsagent
<b>warnings</b>	[ Empty Array ]	
<b>activeEnv</b>	default	
<b>machineToken</b>	[ zero-length string ]	
<b>links</b>	[ Empty Array ]	

## 14.21. GETTING STARTED WITH FACTORIES

A factory is a template containing configuration to automate the generation of a new workspace using a factory identifier added to the IDE URL. Factories can be used to create replicas of existing workspaces or to automate the provisioning of statically or dynamically defined workspaces.

### 14.21.1. Trying a factory

Clone a public workspace on **che.openshift.io** by clicking [try a factory](#).

### 14.21.2. Using factories

Factories can be invoked from a factory URL built in multiple ways. You can replace the **localhost:8080** domain with the hostname of any CodeReady Workspaces installation.



Using factories on **che.openshift.io** requires the user to be authenticated. Users who are not authenticated see a login screen after they click on the factory URL. Users without an account can create one using the same dialog.

### 14.21.3. Invoking factories using their unique hashcodes

Format	<code>/f?id={hashcode}</code> <code>/factory?id={hashcode}</code>
Sample	<a href="https://localhost:8080/f?id=factorymtyoroly0qt8tq2j">https://localhost:8080/f?id=factorymtyoroly0qt8tq2j</a>

### 14.21.4. Invoking a named factory

Format	<code>/f?user={username}&amp;name={factoryname}</code> <code>/factory?user={username}&amp;name={factoryname}</code>
Sample	<a href="https://localhost:8080/f?user=che&amp;name=starwars">https://localhost:8080/f?user=che&amp;name=starwars</a> <a href="https://localhost:8080/factory?user=che&amp;name=starwars">https://localhost:8080/factory?user=che&amp;name=starwars</a>

### 14.21.5. Invoking a factory for a specific git repository

Format	<code>/f?url={git URL}</code>
Sample	<a href="http://localhost:8080/f?url=https://github.com/eclipse/che">http://localhost:8080/f?url=https://github.com/eclipse/che</a> <a href="http://localhost:8080/f?url=https://github.com/eclipse/che/tree/language-server">http://localhost:8080/f?url=https://github.com/eclipse/che/tree/language-server</a> <a href="http://localhost:8080/f?url=https://gitlab.com/benoitf/simple-project">http://localhost:8080/f?url=https://gitlab.com/benoitf/simple-project</a>

Once a factory is executed, it either loads an existing workspace or generates a new one, depending on the factory configuration. The name of the workspace is determined by the factory configuration, and its name becomes a part of the URL used to access the factory. The format is: **{hostname}/{username}/{workspace}**.

### 14.21.6. Next steps

You have just created your first developer workspace using factories. Read on to learn more about:

- How to [create factories](#)
- Customizing factories with the [factory JSON reference](#)

### 14.21.7. Creating Factories

#### 14.21.7.1. Creating a factory in the dashboard

You can create a factory based on an existing workspace. You can also create factories based on a template or by pasting in a **.factory.json** file and then generating a factory URL using the CodeReady Workspaces CLI or API. To learn more about the JSON structure and options, see [Factory JSON reference](#).

A factory created from the dashboard is persisted on CodeReady Workspaces and retained when upgrading to a newer version.

To create a factory on the dashboard:

1. In the IDE, click **Dashboard > Factories > Create Factory**.

Sample factory: <https://che.openshift.io/f?id=factorymtyoro1y0qt8tq2j>.

### 14.21.7.2. Creating a factory in the IDE

Creating a factory in the IDE in a running workspace generates a factory to replicate that workspace including the runtime and project settings.

A factory created from the dashboard is persisted on CodeReady Workspaces and retained when upgrading to a newer version.

To create a factory in the IDE:

1. In the IDE, click **Workspace > Create Factory**.

Sample factory: <https://che.openshift.io/f?id=factorymtyoro1y0qt8tq2j>.

### 14.21.7.3. Creating a factory based on a repository

URL factories work with GitHub and GitLab repositories. By using URL factories, the project referenced by the URL is automatically imported.

To create a factory based on a repository:

1. Specify the repository URL. Ensure that you store the configuration in the repository.

Sample factories:

- <http://che.openshift.io/f?url=https://github.com/eclipse/che>
- <http://che.openshift.io/f?url=https://github.com/eclipse/che/tree/language-server>
- <http://che.openshift.io/f?url=https://gitlab.com/benoitf/simple-project>

The factory URL can include a branch or a subdirectory. Following are examples of optional parameters:

- **?url=https://github.com/eclipse/che** CodeReady Workspaces is imported with the **master** branch.
- **?url=https://github.com/eclipse/che/tree/5.0.0** CodeReady Workspaces is imported by using the **5.0.0** branch.
- **?url=https://github.com/eclipse/che/tree/5.0.0/dashboard** subdirectory **dashboard/** is imported by using the **5.0.0** branch.

#### 14.21.7.3.1. Customizing URL factories

The following are two ways to customize the runtime and configuration:

##### Customizing only the runtime

Providing a **.factory.json** file inside the repository signals to CodeReady Workspaces URL factory to configure the project and runtime according to this configuration file. When a **.factory.json** file is stored inside the repository, any **Dockerfile** content is ignored because the workspace runtime configuration is defined inside the JSON file.

### Customizing the Dockerfile

(This only works on Docker infrastructure. On recent CodeReady Workspaces versions, support of this feature may be reduced or dropped.) Providing a **.factory.dockerfile** inside the repository signals to the URL factory to use this **Dockerfile** for the workspace agent runtime. By default, imported projects are set to a **blank** project type. You can also set the project type in the **.factory.json** file or in the workspace definition that the factory inherits from.

#### 14.21.7.4. Configuring factory policies

Policies are a way to send instructions to the automation engine about the number of workspaces to create and their meta data such as lifespan and resource allocation.

##### 14.21.7.4.1. Setting factory limitations

###### Referer

CodeReady Workspaces checks the hostname of the acceptor and only allows the factory to execute if there is a match.

###### Since and Until

Defines the time window in which the factory can be activated. For example, instructors who want to create an exercise that can only be accessed for two hours should set these properties.

##### 14.21.7.4.2. Setting factory multiplicity

Multiplicity defines the number of workspaces that can be created from the factory.

###### Multiple workspaces (perClick)

Every click of the factory URL generates a different workspace, each with its own identifier, name, and resources.

###### Single workspace (perUser)

Exactly one workspace is generated for each unique user that clicks on the factory URL. Existing workspaces are reopened.

To learn how to configure factory policies, see the [JSON reference](#).

#### 14.21.7.5. Customizing the IDE

You can instruct the factory to invoke a series of IDE actions based on events in the lifecycle of the workspace.

##### 14.21.7.6. Lifecycle Events

The lifecycle of the workspace is defined by the following events:

- **onAppLoaded**: Triggered when the IDE is loaded.
- **onProjectsLoaded**: Triggered when the workspace and all projects have been activated.
- **onAppClosed**: Triggered when the IDE is closed.

Each event type has a set of actions that can be triggered. There is no ordering of actions executed when you provide a list; CodeReady Workspaces asynchronously invokes multiple actions if appropriate.

#### 14.21.7.7. Factory actions

The following is a list of all possible actions that can be configured with your factory.

##### Run Command

Specify the name of the command to invoke after the IDE is loaded.

*Associated Event:* **onProjectsLoaded**

##### Open File

Open project files in the editor. Optionally, define the line to be highlighted.

*Associated Event:* **onProjectsLoaded**

##### Open a Welcome Page

Customize content of a welcome panel displayed when the workspace is loaded.

*Associated Event:* **onAppLoaded**

##### Warn on Uncommitted Changes

Opens a warning pop-up window when the user closes the browser tab with a project that has uncommitted changes.

*Associated Event:* **onAppClosed**

To learn how to configure factory actions, see the [Factory JSON reference](#).

#### 14.21.7.8. Finding and replacing variables

Factories make it possible to replace variables or placeholders in the source code – used to avoid exposing sensitive information (passwords, URLs, account names, API keys) – with real values. To find and replace a value, you can use the **run** command during an **onProjectsLoaded** event. You can use **sed**, **awk**, or other tools available in your workspace environment.

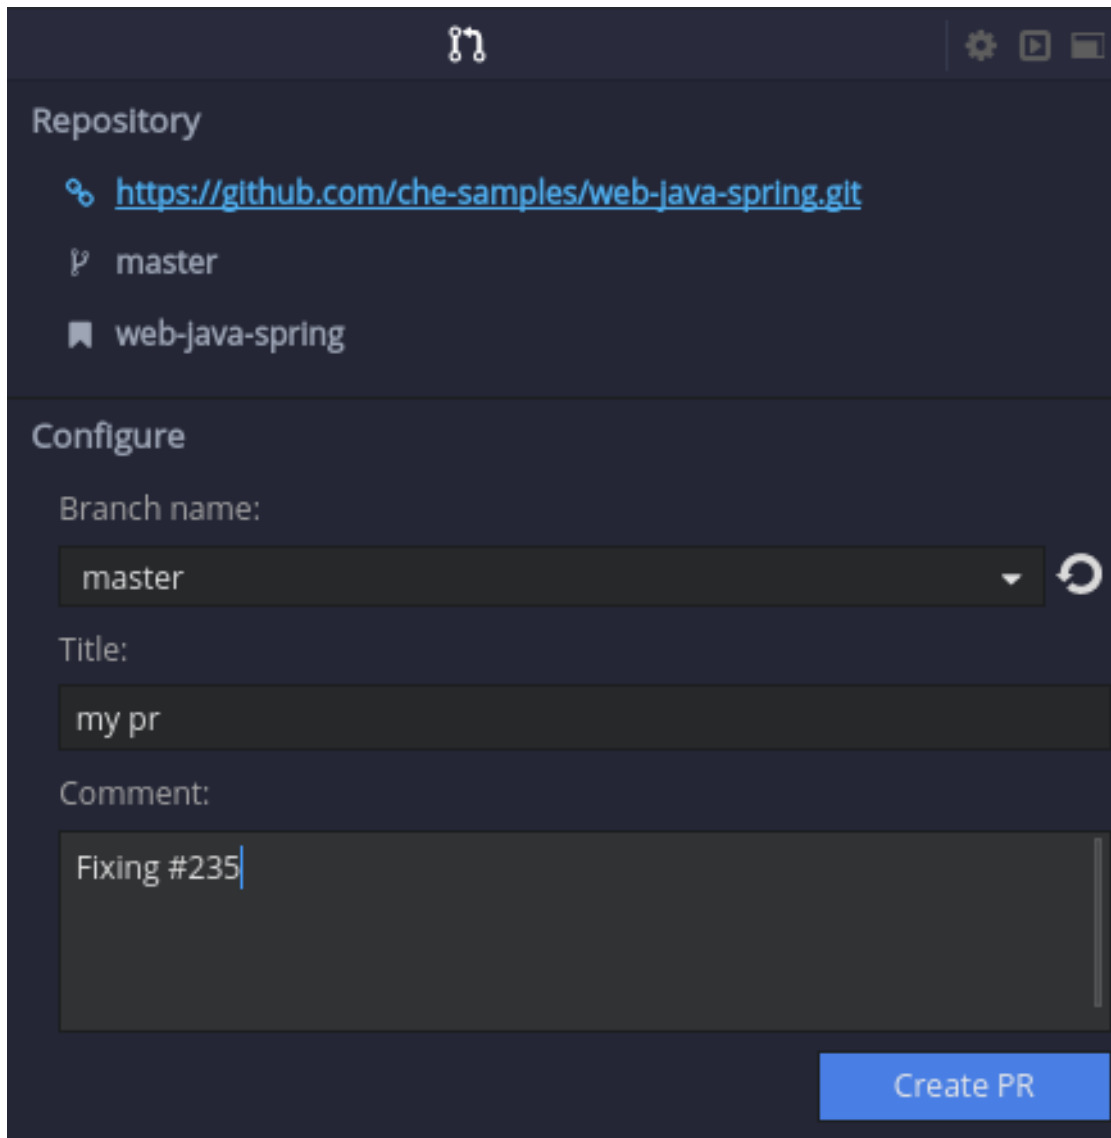
For a sample of how to configure finding and replacing a value, see the [Factory JSON reference](#) section. Alternatively, you can add IDE actions in the **Factory** tab, on the user **Dashboard**.

Use [regular expressions](#) in **sed**, both in find-replace and file-file type patterns.

#### 14.21.7.9. Pull request workflow

Factories can be configured with a dedicated pull request workflow. The PR workflow handles local and remote branching, forking, and issuing the pull request. Pull requests generated from within CodeReady Workspaces have another factory placed into the comments of the pull requests that a PR reviewer can use to quickly start the workspace.

When enabled, the pull request workflow adds a contribution panel to the IDE.



Repository

<https://github.com/che-samples/web-java-spring.git>

master

web-java-spring

Configure

Branch name:

master

Title:

my pr

Comment:

Fixing #235

Create PR

#### 14.21.7.10. Repository badging

If you have projects in GitHub or GitLab, you can help your contributors to get started by providing them ready-to-code developer workspaces. Create a factory and add the following badge on your repositories **readme.md**:

```
[[Developer Workspace]](https://che.openshift.io/factory/resources/factory-contribute.svg)](your-factory-url)
```

#### 14.21.7.11. Next steps

- Read about customizing factories with the [Factory JSON reference](#).

#### 14.21.8. Factories JSON Reference

A factory configuration is a JSON snippet either stored within CodeReady Workspaces or as a **.factory.json** file. You can create factories within the IDE using the CodeReady Workspaces URL syntax, within the dashboard, or on the command line with the API and CLI.

```
factory : {
  "v"      : 4.0,
```

1

```

"workspace" : {},
"policies" : {},
"ide"      : {},
"creator"  : {},
}

```

- 1 Version of the configuration format.
- 2 Identical to **workspace: {}** object for CodeReady Workspaces.
- 3 (Optional) Restrictions that limit behaviors.
- 4 (Optional) Trigger IDE UI actions tied to workspace events.
- 5 (Optional) Identifying information of author.

The **factory.workspace** is identical to the **workspace: {}** object for CodeReady Workspaces and contains the structure of the workspace. To learn more about the workspace JSON object, see [Workspace Data Model](#).

You can export workspaces and then reuse the workspace definition within a factory. workspaces are composed of the following:

- 0..n projects
- 0..n environments that contain machines to run the code
- 0..n commands to execute against the code and machines
- a type

The **factory.policies**, **factory.ide**, and **factory.creator** objects are unique to factories. They provide meta information to the automation engine that alter the presentation of the factory URL or the behavior of the provisioning.

#### 14.21.8.1. Mixins

A mixin adds additional behaviors to a project as a set of new project type attributes. Mixins are reusable across any project type. To define the mixins to add to a project, specify an array of strings, with each string containing the identifier for the mixin. For example, **"mixins" : [ "pullrequest" ]**.

Mixin ID	Description
<b>pullrequest</b>	Enables pull request workflow where CodeReady Workspaces handles local and remote branching, forking, and pull request issuance. Pull requests generated from within CodeReady Workspaces have another factory placed into the comments of pull requests that a PR reviewer can consume. Adds contribution panel to the IDE. If this mixin is set, it uses attribute values for <b>project.attributes.local_branch</b> and <b>project.attributes.contribute_to_branch</b>

- The **pullrequest** mixin requires additional configuration from the **attributes** object of the project. If present, `{{ site.product_mini_name }}` will use the project attributes as defined in the factory. If not provided, `{{ site.product_mini_name }}` will set defaults for the attributes.
- Learn more about other link:TODO: link to project API docs[mixins]

### 14.21.8.2. Pull Request mixin attributes

Project attributes alter the behavior of the IDE or workspace.

Different CodeReady Workspaces plug-ins can add their own attributes to affect the behavior of the IDE or workspace. Attribute configuration is always optional and if not provided within a factory definition, the IDE or workspace sets it.

Attribute	Description
<b>local_branch</b>	Used in conjunction with the <b>pullrequest</b> mixin. If provided, the local branch for the project is set with this value. If not provided, the local branch is set with the value of <b>project.source.parameters.branch</b> (the name of the branch from the remote). If both <b>local_branch</b> and <b>project.source.parameters.branch</b> are not provided, the local branch is set to the name of the checked out branch.
<b>contribute_to_branch</b>	Name of the branch that a pull request will be contributed to. The value of <b>project.source.parameters.branch</b> is default. It is the name of the branch that this project was cloned from.

Following is a snippet that demonstrates full configuration of the **contribution** mixin.

```
factory.workspace.project : {
  "mixins" : [ "pullrequest" ],

  "attributes" : {
    "local_branch" : [ "timing" ],
    "contribute_to_branch" : [ "master" ]
  },

  "source" : {
    "type" : "git",
    "location" : "https://github.com/codenvy/che.git",
    "parameters" : {
      "keepVcs" : "true"
    }
  }
}
```

### 14.21.8.3. Policies

Following is an example of a factory policy.

```
factory.policies : {
  "referrer" : STRING,      1
  "since"    : EPOCHTIME,  2
  "until"    : EPOCHTIME,  3
  "create"   : [perClick | perUser] 4
}
```

- 1 Works only for clients from a referrer.
- 2 Factory works only after this date.
- 3 Factory works only before this date.
- 4 Create one workspace per click, user, or account.

#### 14.21.8.4. Limitations

You can use **since : EPOCHTIME**, **until : EPOCHTIME**, and **referrer** as a way to prevent the factory from executing under certain conditions. **since** and **until** represent a valid time window that allows the factory to activate. The **referrer** checks the hostname of the acceptor and only allows the factory to execute if there is a match.

#### 14.21.8.5. Multiplicity

Using **create : perClick** causes every click of the factory URL to generate a new workspace, each with its own identifier, name, and resources. Using **create : perUser** causes only one workspace to be generated for each unique user that clicks on the factory URL. If the workspace has previously been generated, the existing workspace is reopened.

#### 14.21.8.6. Customizing the IDE

```
factory.ide.{event} : {
  "actions" : [{}],      1
}                          2

factory.ide.{event}.actions : [{
  "id"      : String,    3
  "properties" : {}      4
}]
```

- 1 event = **onAppLoaded**, **onProjectsLoaded**, **onAppClosed**.
- 2 List of IDE actions to be executed when the event is triggered.
- 3 Action for the IDE to perform when the event is triggered.
- 4 Properties to customize action behavior.

You can instruct the factory to invoke a series of IDE actions based on events in the lifecycle of the workspace.



**onAppLoaded**

Triggered when the IDE is loaded.

**onProjectsLoaded**

Triggered when the workspace and all projects have been activated or imported.

**onAppClosed**

Triggered when the IDE is closed.

Following is an example that associates a variety of actions with all of the events.

```

"ide" : {
  "onProjectsLoaded" : {
    "actions" : [{
      "id" : "openFile",
      "properties" : {
        "file" : "/my-project/pom.xml"
      }
    }],
    {
      "id" : "runCommand",
      "properties" : {
        "name" : "MCI"
      }
    }
  ]},
  "onAppLoaded": {
    "actions": [
      {
        "properties":{
          "greetingTitle": "Getting Started",
          "greetingContentUrl": "http://example.com/README.html"
        },
        "id": "openWelcomePage"
      }
    ]
  },
  "onAppClosed" : {
    "actions" : [{
      "id" : "warnOnClose"
    }]
  }
}

```

- 1 Actions triggered when a project is opened.
- 2 Opens a file in the editor. Can add multiple.
- 3 The file to be opened (include project name).
- 4 Launch command after the IDE opens.
- 5 Command name.
- 6 Title of a **Welcome** tab.

- 7 HTML file to be loaded into a tab.
- 8 Actions to be triggered when the IDE is closed.
- 9 Show warning when closing a browser tab.

Each event type has a set of actions that can be triggered. There is no ordering of actions executed when you provide a list; `{{ site.product_mini_name }}` will asynchronously invoke multiple actions if appropriate. Some actions can be configured in how they perform and will have an associated **properties : {}** object.

#### onProjectsLoaded Event

Action	Properties?	Description
<b>runCommand</b>	Yes	Specify the name of the command to invoke after the IDE is loaded. Specify the commands in the <b>factory.workspace.commands : []</b> array.
<b>openFile</b>	Yes	Open project files as a tab in the editor.

#### onAppLoaded Event

Action	Properties?	Description
<b>openWelcomePage</b>	Yes	Customize the content of the welcome panel when the workspace is loaded. Note that browsers block http resources that are loaded into https pages.

#### onAppClosed Event

Action	Properties?	Description
<b>warnOnClose</b>	No	Opens a warning pop-up window when the user closes the browser tab with a project that has uncommitted changes. Requires <b>project.parameters.keepVcs</b> to be <b>true</b> .

#### 14.21.8.7. Action: Open File

This action opens a file as a tab in the editor. You can provide this action multiple times to have multiple

files open. The file property is a relative reference to a file in the project source tree. The **file** parameter is the relative path within the workspace to the file that should be opened by the editor. The **line** parameter is optional and can be used to move the editor cursor to a specific line when the file is opened. Projects are located in the **/projects/** directory of a workspace.

```
{
  "id": "openFile",
  "properties": {
    "file": "/my-project/pom.xml",
    "line": "50"
  }
}
```

#### 14.21.8.8. Action: Find and Replace

In projects created from a factory, CodeReady Workspaces can find and replace values in the source code after it is imported into the project tree. This lets you parameterize your source code. Find and replace can be run as a **Run Command** during **onProjectsLoaded** event. You can use **sed**, **awk**, or any other tools that are available in your workspace environment.

To define a command for your workspace in **factory.workspace.workspaceConfig.commands**:

```
{
  "commandLine": "sed -i 's/*/userld984hfy6/g' /projects/console-java-simple/README.md",
  "name": "replace",
  "attributes": {
    "goal": "Common",
    "previewUrl": ""
  },
  "type": "custom"
}
```

In the preceding example, a named command **replace** is created. The command replaces each occurrence of **\*** with the string **userld984hfy6** in the **README.md** file of the project.

Then register this command to the execution list linked to the **onProjectsLoaded** event. In this example, the **replace** command is executed after the project is imported into a workspace.

```
"ide": {
  "onProjectsLoaded": {
    "actions": [
      {
        "properties": {
          "name": "replace"
        },
        "id": "runCommand"
      }
    ]
  }
}
```

Use [regular expressions](#) in **sed**, both in find-replace and file-file type patterns.

#### 14.21.8.9. Creator

This object has meta information that you can embed within the factory. These attributes do not affect the automation behavior or the behavior of the generated workspace.

```
factory.creator : {  
  "name"      : STRING,      1  
  "email"     : STRING,      2  
  "created"   : EPOCHTIME,   3  
  "userId"    : STRING       4  
}
```

- 1 The name of the author of this configuration file.
- 2 The author's email address.
- 3 This value is set by the system.
- 4 This value is set by the system.