



# OpenShift Dedicated 4

## Installing, accessing, and deleting OpenShift Dedicated clusters

Installing, accessing, and deleting OpenShift Dedicated clusters



# OpenShift Dedicated 4 Installing, accessing, and deleting OpenShift Dedicated clusters

---

Installing, accessing, and deleting OpenShift Dedicated clusters

## Legal Notice

Copyright © 2022 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

This document provides information on how to install OpenShift Dedicated clusters. The document also provides details on how to configure identity providers.

---

## Table of Contents

<b>CHAPTER 1. CREATING A CLUSTER ON AWS</b> .....	<b>3</b>
1.1. PREREQUISITES	3
1.2. CREATING A CLUSTER ON AWS WITH CCS	3
1.3. CREATING A CLUSTER ON AWS WITH A RED HAT CLOUD ACCOUNT	7
1.4. ADDITIONAL RESOURCES	10
<b>CHAPTER 2. CREATING A CLUSTER ON GCP</b> .....	<b>11</b>
2.1. PREREQUISITES	11
2.2. CREATING A CLUSTER ON GCP WITH CCS	11
2.3. CREATING A CLUSTER ON GCP WITH A RED HAT CLOUD ACCOUNT	15
2.4. ADDITIONAL RESOURCES	18
<b>CHAPTER 3. CONFIGURING IDENTITY PROVIDERS</b> .....	<b>19</b>
3.1. UNDERSTANDING IDENTITY PROVIDERS	19
3.1.1. Supported identity providers	19
3.1.2. Identity provider parameters	19
3.2. CONFIGURING A GITHUB IDENTITY PROVIDER	20
3.3. CONFIGURING A GITLAB IDENTITY PROVIDER	22
3.4. CONFIGURING A GOOGLE IDENTITY PROVIDER	23
3.5. CONFIGURING A LDAP IDENTITY PROVIDER	24
3.6. CONFIGURING AN OPENID IDENTITY PROVIDER	26
3.7. CONFIGURING AN HTPASSWD IDENTITY PROVIDER	27
3.8. ACCESSING YOUR CLUSTER	29



# CHAPTER 1. CREATING A CLUSTER ON AWS

You can install OpenShift Dedicated on Amazon Web Services (AWS) by using your own AWS account through the Customer Cloud Subscription (CCS) model or by using an AWS infrastructure account that is owned by Red Hat.

## 1.1. PREREQUISITES

- You reviewed the [introduction to OpenShift Dedicated](#) and the documentation on [architecture concepts](#).
- You reviewed the [OpenShift Dedicated cloud deployment options](#).

## 1.2. CREATING A CLUSTER ON AWS WITH CCS

By using the Customer Cloud Subscription (CCS) billing model, you can create an OpenShift Dedicated cluster in an existing Amazon Web Services (AWS) account that you own.

You must meet several prerequisites if you use the CCS model to deploy and manage OpenShift Dedicated into your AWS account.

### Prerequisites

- You have configured your AWS account for use with OpenShift Dedicated.
- You have not deployed any services in your AWS account.
- You have configured the AWS account quotas and limits that are required to support the desired cluster size.
- You have an **osdCcsAdmin** AWS Identity and Access Management (IAM) user with the **AdministratorAccess** policy attached.
- You have set up a service control policy (SCP) in your AWS organization. For more information, see *Minimum required service control policy (SCP)*.
- Consider having **Business Support** or higher from AWS.

### Procedure

1. Log in to [OpenShift Cluster Manager](#) and click **Create cluster**.
2. In the **Cloud** tab, click **Create cluster** in the **Red Hat OpenShift Dedicated** row.
3. Under **Billing model**, configure the subscription type and infrastructure type:
  - a. Select a subscription type. For information about OpenShift Dedicated subscription options, see [Managing OpenShift Dedicated cluster subscriptions](#) in the OpenShift Cluster Manager documentation.



### NOTE

The subscription types that are available to you depend on your OpenShift Dedicated subscriptions and resource quotas. For more information, contact your sales representative or Red Hat support.

- b. Select the **Customer Cloud Subscription** infrastructure type to deploy OpenShift Dedicated in an existing cloud provider account that you own.
  - c. Click **Next**.
4. Select **Run on Amazon Web Services**
5. Click **Prerequisites** to review the prerequisites for installing OpenShift Dedicated on AWS with CCS.
6. Provide your AWS account details:
  - a. Enter your **AWS account ID**.
  - b. Enter your **AWS access key ID** and **AWS secret access key** for your AWS IAM user account.

**NOTE**

Revoking these credentials in AWS results in a loss of access to any cluster created with these credentials.

- c. Optional: You can select **Bypass AWS service control policy (SCP) checks** to disable the SCP checks.

**NOTE**

Some AWS SCPs can cause the installation to fail, even if you have the required permissions. Disabling the SCP checks allows an installation to proceed. The SCP is still enforced even if the checks are bypassed.

7. **Validate** your cloud provider account and then click **Next**.
8. On the **Cluster details** page, provide a name for your cluster and specify the cluster details:
  - a. Add a **Cluster name**.
  - b. Select a cluster version from the **Version** drop-down menu.
  - c. Select a cloud provider region from the **Region** drop-down menu.
  - d. Select a **Single zone** or **Multi-zone** configuration.
  - e. Leave **Enable user workload monitoring** selected to monitor your own projects in isolation from Red Hat Site Reliability Engineer (SRE) platform metrics. This option is enabled by default.
  - f. Optional: Select **Enable additional etcd encryption** if you require etcd key value encryption. With this option, the etcd key values are encrypted, but not the keys. This option is in addition to the control plane storage encryption that encrypts the etcd volumes in OpenShift Dedicated clusters by default.



**NOTE**

By enabling etcd encryption for the key values in etcd, you will incur a performance overhead of approximately 20%. The overhead is a result of introducing this second layer of encryption, in addition to the default control plane storage encryption that encrypts the etcd volumes. Consider enabling etcd encryption only if you specifically require it for your use case.

- g. Optional: Select **Encrypt persistent volumes with customer keys** if you want to provide your own AWS Key Management Service (KMS) key Amazon Resource Name (ARN). These keys are used for encrypting all control plane, infrastructure, and worker node root volumes.
  - h. Click **Next**.
9. On the **Default machine pool** page, select a **Compute node instance type** and a **Compute node count**. The number and types of nodes that are available depend on your OpenShift Dedicated subscription. If you are using multiple availability zones, the compute node count is per zone.

**NOTE**

After your cluster is created, you can change the number of compute nodes in your cluster, but you cannot change the compute node instance type in a machine pool. The number and types of nodes available to you depend on your OpenShift Dedicated subscription.

10. Optional: Expand **Edit node labels** to add labels to your nodes. Click **Add label** to add more node labels and select **Next**.
11. In the **Cluster privacy** section, select **Public** or **Private** to use either public or private API endpoints and application routes for your cluster.

**IMPORTANT**

If you are using private API endpoints, you cannot access your cluster until you update the network settings in your cloud provider account.

12. Optional: To install the cluster in an existing AWS Virtual Private Cloud (VPC):
- a. Select **Install into an existing VPC**
  - b. If you opted to use private API endpoints and are installing into an existing VPC, you can select **Use a PrivateLink** to enable connections to the cluster by Red Hat Site Reliability Engineering (SRE) using only AWS PrivateLink endpoints. This option cannot be changed after a cluster is created.
13. Click **Next**.
14. If you opted to install the cluster in an existing AWS VPC, provide your **Virtual Private Cloud (VPC) subnet settings**.


**NOTE**

You must ensure that your VPC is configured with a public and a private subnet for each availability zone that you want the cluster installed into. If you opted to use PrivateLink, only private subnets are required.

15. In the **CIDR ranges** dialog, configure custom classless inter-domain routing (CIDR) ranges or use the defaults that are provided.

**NOTE**

If you are installing into a VPC, the **Machine CIDR** range must match the VPC subnets.

**IMPORTANT**

CIDR configurations cannot be changed later. Confirm your selections with your network administrator before proceeding.

16. On the **Cluster update strategy** page, configure your update preferences:

- a. Choose a cluster update method:

- Select **Individual updates** if you want to schedule each update individually. This is the default option.
- Select **Recurring updates** to update your cluster on your preferred day and start time, when updates are available.

**NOTE**

You can review the end-of-life dates in the update life cycle documentation for OpenShift Dedicated. For more information, see *OpenShift Dedicated update life cycle*.

- b. Provide administrator approval based on your cluster update method:

- Individual updates: If you select an update version that requires approval, provide an administrator's acknowledgment and click **Approve and continue**.
- Recurring updates: If you selected recurring updates for your cluster, provide an administrator's acknowledgment and click **Approve and continue**. OpenShift Cluster Manager does not start scheduled y-stream updates for minor versions without receiving an administrator's acknowledgment.  
For information about administrator acknowledgment, see [Administrator acknowledgment when upgrading to OpenShift 4.9](#).

- c. If you opted for recurring updates, select a preferred day of the week and upgrade start time in UTC from the drop-down menus.
- d. Optional: You can set a grace period for **Node draining** during cluster upgrades. A **1 hour** grace period is set by default.
- e. Click **Next**.

**NOTE**

In the event of critical security concerns that significantly impact the security or stability of a cluster, Red Hat Site Reliability Engineering (SRE) might schedule automatic updates to the latest z-stream version that is not impacted. The updates are applied within 48 hours after customer notifications are provided. For a description of the critical impact security rating, see [Understanding Red Hat security ratings](#) .

17. Review the summary of your selections and click **Create cluster** to start the cluster installation. The installation takes approximately 30–40 minutes to complete.

**Verification**

- You can monitor the progress of the installation in the **Overview** page for your cluster. You can view the installation logs on the same page. Your cluster is ready when the **Status** in the **Details** section of the page is listed as **Ready**.

## 1.3. CREATING A CLUSTER ON AWS WITH A RED HAT CLOUD ACCOUNT

Through [OpenShift Cluster Manager](#), you can create an OpenShift Dedicated cluster on Amazon Web Services (AWS) using a standard cloud provider account owned by Red Hat.

**Procedure**

1. Log in to [OpenShift Cluster Manager](#) and click **Create cluster**.
2. In the **Cloud** tab, click **Create cluster** in the **Red Hat OpenShift Dedicated** row.
3. Under **Billing model**, configure the subscription type and infrastructure type:
  - a. Select the **Annual** subscription type. Only the **Annual** subscription type is available when you deploy a cluster using a Red Hat cloud account. For information about OpenShift Dedicated subscription options, see [Managing OpenShift Dedicated cluster subscriptions](#) in the OpenShift Cluster Manager documentation.

**NOTE**

You must have the required resource quota for the **Annual** subscription type to be available. For more information, contact your sales representative or Red Hat support.

- b. Select the **Red Hat cloud account** infrastructure type to deploy OpenShift Dedicated in a cloud provider account that is owned by Red Hat.
  - c. Click **Next**.
4. Select **Run on Amazon Web Services** and click **Next**.
5. On the **Cluster details** page, provide a name for your cluster and specify the cluster details:
  - a. Add a **Cluster name**.

- b. Select a cluster version from the **Version** drop-down menu.
- c. Select a cloud provider region from the **Region** drop-down menu.
- d. Select a **Single zone** or **Multi-zone** configuration.
- e. Select a **Persistent storage** capacity for the cluster. For more information, see the *Storage* section in the OpenShift Dedicated service definition.
- f. Specify the number of **Load balancers** that you require for your cluster. For more information, see the *Load balancers* section in the OpenShift Dedicated service definition.
- g. Leave **Enable user workload monitoring** selected to monitor your own projects in isolation from Red Hat Site Reliability Engineer (SRE) platform metrics. This option is enabled by default.
- h. Optional: Select **Enable additional etcd encryption** if you require etcd key value encryption. With this option, the etcd key values are encrypted, but not the keys. This option is in addition to the control plane storage encryption that encrypts the etcd volumes in OpenShift Dedicated clusters by default.

**NOTE**


By enabling etcd encryption for the key values in etcd, you will incur a performance overhead of approximately 20%. The overhead is a result of introducing this second layer of encryption, in addition to the default control plane storage encryption that encrypts the etcd volumes. Consider enabling etcd encryption only if you specifically require it for your use case.

- i. Click **Next**.
6. On the **Default machine pool** page, select a **Compute node instance type** and a **Compute node count**. The number and types of nodes that are available depend on your OpenShift Dedicated subscription. If you are using multiple availability zones, the compute node count is per zone.

**NOTE**

After your cluster is created, you can change the number of compute nodes, but you cannot change the compute node instance type in a machine pool. For clusters that use the CCS model, you can add machine pools after installation that use a different instance type. The number and types of nodes available to you depend on your OpenShift Dedicated subscription.

7. Optional: Expand **Edit node labels** to add labels to your nodes. Click **Add label** to add more node labels and select **Next**.
8. In the **Cluster privacy** dialog, select **Public** or **Private** to use either public or private API endpoints and application routes for your cluster.
9. Click **Next**.
10. In the **CIDR ranges** dialog, configure custom classless inter-domain routing (CIDR) ranges or use the defaults that are provided.



## IMPORTANT

CIDR configurations cannot be changed later. Confirm your selections with your network administrator before proceeding.

If the cluster privacy is set to **Private**, you cannot access your cluster until you configure private connections in your cloud provider.

11. On the **Cluster update strategy** page, configure your update preferences:

a. Choose a cluster update method:

- Select **Individual updates** if you want to schedule each update individually. This is the default option.
- Select **Recurring updates** to update your cluster on your preferred day and start time, when updates are available.



## NOTE

You can review the end-of-life dates in the update life cycle documentation for OpenShift Dedicated. For more information, see *OpenShift Dedicated update life cycle*.

b. Provide administrator approval based on your cluster update method:

- Individual updates: If you select an update version that requires approval, provide an administrator's acknowledgment and click **Approve and continue**.
- Recurring updates: If you selected recurring updates for your cluster, provide an administrator's acknowledgment and click **Approve and continue**. OpenShift Cluster Manager does not start scheduled y-stream updates for minor versions without receiving an administrator's acknowledgment. For information about administrator acknowledgment, see [Administrator acknowledgment when upgrading to OpenShift 4.9](#).

c. If you opted for recurring updates, select a preferred day of the week and upgrade start time in UTC from the drop-down menus.

d. Optional: You can set a grace period for **Node draining** during cluster upgrades. A **1 hour** grace period is set by default.

e. Click **Next**.



## NOTE

In the event of critical security concerns that significantly impact the security or stability of a cluster, Red Hat Site Reliability Engineering (SRE) might schedule automatic updates to the latest z-stream version that is not impacted. The updates are applied within 48 hours after customer notifications are provided. For a description of the critical impact security rating, see [Understanding Red Hat security ratings](#).

12. Review the summary of your selections and click **Create cluster** to start the cluster installation. The installation takes approximately 30-40 minutes to complete.

## Verification

- You can monitor the progress of the installation in the **Overview** page for your cluster. You can view the installation logs on the same page. Your cluster is ready when the **Status** in the **Details** section of the page is listed as **Ready**.

## 1.4. ADDITIONAL RESOURCES

- For details about the AWS service control policies required for CCS deployments, see [Minimum required service control policy \(SCP\)](#).
- For information about persistent storage for OpenShift Dedicated, see the [Storage](#) section in the OpenShift Dedicated service definition.
- For information about load balancers for OpenShift Dedicated, see the [Load balancers](#) section in the OpenShift Dedicated service definition.
- For more information about etcd encryption, see the [etcd encryption service definition](#).
- For information about the end-of-life dates for OpenShift Dedicated versions, see the [OpenShift Dedicated update life cycle](#).

## CHAPTER 2. CREATING A CLUSTER ON GCP

You can install OpenShift Dedicated on Google Cloud Platform (GCP) by using your own GCP account through the Customer Cloud Subscription (CCS) model or by using a GCP infrastructure account that is owned by Red Hat.

### 2.1. PREREQUISITES

- You reviewed the [introduction to OpenShift Dedicated](#) and the documentation on [architecture concepts](#).
- You reviewed the [OpenShift Dedicated cloud deployment options](#).

### 2.2. CREATING A CLUSTER ON GCP WITH CCS

By using the Customer Cloud Subscription (CCS) billing model, you can create an OpenShift Dedicated cluster in an existing Google Cloud Platform (GCP) account that you own.

You must meet several prerequisites if you use the CCS model to deploy and manage OpenShift Dedicated into your GCP account.

#### Prerequisites

- You have configured your GCP account for use with OpenShift Dedicated.
- You have configured the GCP account quotas and limits that are required to support the desired cluster size.
- You have created a GCP project.




#### NOTE

The project name must be 10 characters or less.

- You have enabled the Google Cloud Resource Manager API in your GCP project. For more information about enabling APIs for your project, see [the Google Cloud documentation](#).
- You have an IAM service account in GCP called **osd-ccs-admin** with the following roles attached:
  - DNS Administrator
  - Organization Policy Viewer
  - Owner
  - Project IAM Admin
  - Service Management Administrator
  - Service Usage Admin
  - Storage Admin

- You have created a key for your **osd-ccs-admin** GCP service account and exported it to a file named **osServiceAccount.json**.




#### NOTE

For more information about creating a key for your GCP service account and exporting it to a JSON file, see [Creating service account keys](#) in the Google Cloud documentation.

- Consider having **Production Support** or higher from GCP.
- To prevent potential conflicts, consider having no other resources provisioned in the project prior to installing OpenShift Dedicated.

### Procedure

1. Log in to [OpenShift Cluster Manager](#) and click **Create cluster**.
2. In the **Cloud** tab, click **Create cluster** in the **Red Hat OpenShift Dedicated** row.
3. Under **Billing model**, configure the subscription type and infrastructure type:
  - a. Select a subscription type. For information about OpenShift Dedicated subscription options, see [Managing OpenShift Dedicated cluster subscriptions](#) in the OpenShift Cluster Manager documentation.



#### NOTE

The subscription types that are available to you depend on your OpenShift Dedicated subscriptions and resource quotas. For more information, contact your sales representative or Red Hat support.

- b. Select the **Customer Cloud Subscription** infrastructure type to deploy OpenShift Dedicated in an existing cloud provider account that you own.
  - c. Click **Next**.
4. Select **Run on Google Cloud Platform**
  5. Click **Prerequisites** to review the prerequisites for installing OpenShift Dedicated on GCP with CCS.
  6. Provide your GCP service account private key in JSON format. You can either click **Browse** to locate and attach a JSON file or add the details in the **Service account JSON** field.
  7. **Validate** your cloud provider account and then click **Next**.
  8. On the **Cluster details** page, provide a name for your cluster and specify the cluster details:
    - a. Add a **Cluster name**.
    - b. Select a cluster version from the **Version** drop-down menu.
    - c. Select a cloud provider region from the **Region** drop-down menu.
    - d. Select a **Single zone** or **Multi-zone** configuration.



- e. Leave **Enable user workload monitoring** selected to monitor your own projects in isolation from Red Hat Site Reliability Engineer (SRE) platform metrics. This option is enabled by default.
- f. Optional: Select **Enable additional etcd encryption** if you require etcd key value encryption. With this option, the etcd key values are encrypted, but not the keys. This option is in addition to the control plane storage encryption that encrypts the etcd volumes in OpenShift Dedicated clusters by default.



#### NOTE

By enabling etcd encryption for the key values in etcd, you will incur a performance overhead of approximately 20%. The overhead is a result of introducing this second layer of encryption, in addition to the default control plane storage encryption that encrypts the etcd volumes. Consider enabling etcd encryption only if you specifically require it for your use case.

- g. Optional: Select **Encrypt persistent volumes with customer keys** if you want to provide your own encryption keys through the Google Cloud Key Management Service. These keys are used for encrypting all control plane, infrastructure, and worker node root volumes.
- h. Click **Next**.
9. On the **Default machine pool** page, select a **Compute node instance type** and a **Compute node count**. The number and types of nodes that are available depend on your OpenShift Dedicated subscription. If you are using multiple availability zones, the compute node count is per zone.



#### NOTE

After your cluster is created, you can change the number of compute nodes in your cluster, but you cannot change the compute node instance type in a machine pool. The number and types of nodes available to you depend on your OpenShift Dedicated subscription.

10. Optional: Expand **Edit node labels** to add labels to your nodes. Click **Add label** to add more node labels and select **Next**.
11. In the **Cluster privacy** section, select **Public** or **Private** to use either public or private API endpoints and application routes for your cluster.



#### IMPORTANT


If you are using private API endpoints, you cannot access your cluster until you update the network settings in your cloud provider account.

12. Optional: Select **Install into an existing VPC** to install the cluster in an existing GCP Virtual Private Cloud (VPC).
13. Click **Next**.
14. If you opted to install the cluster in an existing GCP VPC, provide your **Virtual Private Cloud (VPC) subnet settings**.

15. In the **CIDR ranges** dialog, configure custom classless inter-domain routing (CIDR) ranges or use the defaults that are provided.

**NOTE**

If you are installing into a VPC, the **Machine CIDR** range must match the VPC subnets.

**IMPORTANT**

CIDR configurations cannot be changed later. Confirm your selections with your network administrator before proceeding.

16. On the **Cluster update strategy** page, configure your update preferences:

- a. Choose a cluster update method:

- Select **Individual updates** if you want to schedule each update individually. This is the default option.
- Select **Recurring updates** to update your cluster on your preferred day and start time, when updates are available.

**NOTE**

You can review the end-of-life dates in the update life cycle documentation for OpenShift Dedicated. For more information, see *OpenShift Dedicated update life cycle*.

- b. Provide administrator approval based on your cluster update method:

- Individual updates: If you select an update version that requires approval, provide an administrator's acknowledgment and click **Approve and continue**.
- Recurring updates: If you selected recurring updates for your cluster, provide an administrator's acknowledgment and click **Approve and continue**. OpenShift Cluster Manager does not start scheduled y-stream updates for minor versions without receiving an administrator's acknowledgment. For information about administrator acknowledgment, see [Administrator acknowledgment when upgrading to OpenShift 4.9](#).

- c. If you opted for recurring updates, select a preferred day of the week and upgrade start time in UTC from the drop-down menus.
- d. Optional: You can set a grace period for **Node draining** during cluster upgrades. A **1 hour** grace period is set by default.
- e. Click **Next**.

**NOTE**

In the event of critical security concerns that significantly impact the security or stability of a cluster, Red Hat Site Reliability Engineering (SRE) might schedule automatic updates to the latest z-stream version that is not impacted. The updates are applied within 48 hours after customer notifications are provided. For a description of the critical impact security rating, see [Understanding Red Hat security ratings](#).

17. Review the summary of your selections and click **Create cluster** to start the cluster installation. The installation takes approximately 30–40 minutes to complete.

**Verification**

- You can monitor the progress of the installation in the **Overview** page for your cluster. You can view the installation logs on the same page. Your cluster is ready when the **Status** in the **Details** section of the page is listed as **Ready**.

## 2.3. CREATING A CLUSTER ON GCP WITH A RED HAT CLOUD ACCOUNT

Through [OpenShift Cluster Manager](#), you can create an OpenShift Dedicated cluster on Google Cloud Platform (GCP) using a standard cloud provider account owned by Red Hat.

**Procedure**

1. Log in to [OpenShift Cluster Manager](#) and click **Create cluster**.
2. In the **Cloud** tab, click **Create cluster** in the **Red Hat OpenShift Dedicated** row.
3. Under **Billing model**, configure the subscription type and infrastructure type:
  - a. Select the **Annual** subscription type. Only the **Annual** subscription type is available when you deploy a cluster using a Red Hat cloud account. For information about OpenShift Dedicated subscription options, see [Managing OpenShift Dedicated cluster subscriptions](#) in the OpenShift Cluster Manager documentation.

**NOTE**

You must have the required resource quota for the **Annual** subscription type to be available. For more information, contact your sales representative or Red Hat support.

- b. Select the **Red Hat cloud account** infrastructure type to deploy OpenShift Dedicated in a cloud provider account that is owned by Red Hat.
  - c. Click **Next**.
4. Select **Run on Google Cloud Platform** and click **Next**.
5. On the **Cluster details** page, provide a name for your cluster and specify the cluster details:
  - a. Add a **Cluster name**.

- b. Select a cluster version from the **Version** drop-down menu.
- c. Select a cloud provider region from the **Region** drop-down menu.
- d. Select a **Single zone** or **Multi-zone** configuration.
- e. Select a **Persistent storage** capacity for the cluster. For more information, see the *Storage* section in the OpenShift Dedicated service definition.
- f. Specify the number of **Load balancers** that you require for your cluster. For more information, see the *Load balancers* section in the OpenShift Dedicated service definition.
- g. Leave **Enable user workload monitoring** selected to monitor your own projects in isolation from Red Hat Site Reliability Engineer (SRE) platform metrics. This option is enabled by default.
- h. Optional: Select **Enable additional etcd encryption** if you require etcd key value encryption. With this option, the etcd key values are encrypted, but not the keys. This option is in addition to the control plane storage encryption that encrypts the etcd volumes in OpenShift Dedicated clusters by default.



#### NOTE

By enabling etcd encryption for the key values in etcd, you will incur a performance overhead of approximately 20%. The overhead is a result of introducing this second layer of encryption, in addition to the default control plane storage encryption that encrypts the etcd volumes. Consider enabling etcd encryption only if you specifically require it for your use case.

- i. Click **Next**.
6. On the **Default machine pool** page, select a **Compute node instance type** and a **Compute node count**. The number and types of nodes that are available depend on your OpenShift Dedicated subscription. If you are using multiple availability zones, the compute node count is per zone.



#### NOTE

After your cluster is created, you can change the number of compute nodes, but you cannot change the compute node instance type in a machine pool. For clusters that use the CCS model, you can add machine pools after installation that use a different instance type. The number and types of nodes available to you depend on your OpenShift Dedicated subscription.

7. Optional: Expand **Edit node labels** to add labels to your nodes. Click **Add label** to add more node labels and select **Next**.
8. In the **Cluster privacy** dialog, select **Public** or **Private** to use either public or private API endpoints and application routes for your cluster.
9. Click **Next**.
10. In the **CIDR ranges** dialog, configure custom classless inter-domain routing (CIDR) ranges or use the defaults that are provided.



## IMPORTANT

CIDR configurations cannot be changed later. Confirm your selections with your network administrator before proceeding.

If the cluster privacy is set to **Private**, you cannot access your cluster until you configure private connections in your cloud provider.

11. On the **Cluster update strategy** page, configure your update preferences:

a. Choose a cluster update method:

- Select **Individual updates** if you want to schedule each update individually. This is the default option.
- Select **Recurring updates** to update your cluster on your preferred day and start time, when updates are available.



## NOTE

You can review the end-of-life dates in the update life cycle documentation for OpenShift Dedicated. For more information, see *OpenShift Dedicated update life cycle*.

b. Provide administrator approval based on your cluster update method:

- Individual updates: If you select an update version that requires approval, provide an administrator's acknowledgment and click **Approve and continue**.
- Recurring updates: If you selected recurring updates for your cluster, provide an administrator's acknowledgment and click **Approve and continue**. OpenShift Cluster Manager does not start scheduled y-stream updates for minor versions without receiving an administrator's acknowledgment. For information about administrator acknowledgment, see [Administrator acknowledgment when upgrading to OpenShift 4.9](#).

c. If you opted for recurring updates, select a preferred day of the week and upgrade start time in UTC from the drop-down menus.

d. Optional: You can set a grace period for **Node draining** during cluster upgrades. A **1 hour** grace period is set by default.

e. Click **Next**.



## NOTE

In the event of critical security concerns that significantly impact the security or stability of a cluster, Red Hat Site Reliability Engineering (SRE) might schedule automatic updates to the latest z-stream version that is not impacted. The updates are applied within 48 hours after customer notifications are provided. For a description of the critical impact security rating, see [Understanding Red Hat security ratings](#).

12. Review the summary of your selections and click **Create cluster** to start the cluster installation. The installation takes approximately 30-40 minutes to complete.

## Verification

- You can monitor the progress of the installation in the **Overview** page for your cluster. You can view the installation logs on the same page. Your cluster is ready when the **Status** in the **Details** section of the page is listed as **Ready**.

## 2.4. ADDITIONAL RESOURCES

- For information about persistent storage for OpenShift Dedicated, see the [Storage](#) section in the OpenShift Dedicated service definition.
- For information about load balancers for OpenShift Dedicated, see the [Load balancers](#) section in the OpenShift Dedicated service definition.
- For more information about etcd encryption, see the [etcd encryption service definition](#) .
- For information about the end-of-life dates for OpenShift Dedicated versions, see the [OpenShift Dedicated update life cycle](#) .

## CHAPTER 3. CONFIGURING IDENTITY PROVIDERS

After your OpenShift Dedicated cluster is created, you must configure identity providers to determine how users log in to access the cluster.

### 3.1. UNDERSTANDING IDENTITY PROVIDERS

OpenShift Dedicated includes a built-in OAuth server. Developers and administrators obtain OAuth access tokens to authenticate themselves to the API. As an administrator, you can configure OAuth to specify an identity provider after you install your cluster. Configuring identity providers allows users to log in and access the cluster.

#### 3.1.1. Supported identity providers

You can configure the following types of identity providers:

Identity provider	Description
GitHub or GitHub Enterprise	Configure a GitHub identity provider to validate usernames and passwords against GitHub or GitHub Enterprise's OAuth authentication server.
GitLab	Configure a GitLab identity provider to use <a href="https://gitlab.com">GitLab.com</a> or any other GitLab instance as an identity provider.
Google	Configure a Google identity provider using <a href="#">Google's OpenID Connect integration</a> .
LDAP	Configure an LDAP identity provider to validate usernames and passwords against an LDAPv3 server, using simple bind authentication.
OpenID Connect	Configure an OpenID Connect (OIDC) identity provider to integrate with an OIDC identity provider using an <a href="#">Authorization Code Flow</a> .
HTPasswd	<p>Configure an HTPasswd identity provider for a single, static administration user. You can log in to the cluster as the user to troubleshoot issues.</p> <div data-bbox="422 1512 531 1742" data-label="Image"> </div> <p><b>IMPORTANT</b></p> <p>The HTPasswd identity provider option is included only to enable the creation of a single, static administration user. HTPasswd is not supported as a general-use identity provider for OpenShift Dedicated. For the steps to configure the single user, see <i>Configuring an HTPasswd identity provider</i>.</p>

#### 3.1.2. Identity provider parameters

The following parameters are common to all identity providers:

Parameter	Description
<b>name</b>	The provider name is prefixed to provider user names to form an identity name.
<b>mappingMethod</b>	<p>Defines how new identities are mapped to users when they log in. Enter one of the following values:</p> <p><b>claim</b> The default value. Provisions a user with the identity's preferred user name. Fails if a user with that user name is already mapped to another identity.</p> <p><b>lookup</b> Looks up an existing identity, user identity mapping, and user, but does not automatically provision users or identities. This allows cluster administrators to set up identities and users manually, or using an external process. Using this method requires you to manually provision users.</p> <p><b>generate</b> Provisions a user with the identity's preferred user name. If a user with the preferred user name is already mapped to an existing identity, a unique user name is generated. For example, <b>myuser2</b>. This method should not be used in combination with external processes that require exact matches between OpenShift Dedicated user names and identity provider user names, such as LDAP group sync.</p> <p><b>add</b> Provisions a user with the identity's preferred user name. If a user with that user name already exists, the identity is mapped to the existing user, adding to any existing identity mappings for the user. Required when multiple identity providers are configured that identify the same set of users and map to the same user names.</p>

**NOTE**

When adding or changing identity providers, you can map identities from the new provider to existing users by setting the **mappingMethod** parameter to **add**.

## 3.2. CONFIGURING A GITHUB IDENTITY PROVIDER

Configure a GitHub identity provider to validate user names and passwords against GitHub or GitHub Enterprise's OAuth authentication server and access your OpenShift Dedicated cluster. OAuth facilitates a token exchange flow between OpenShift Dedicated and GitHub or GitHub Enterprise.

**WARNING**

Configuring GitHub authentication allows users to log in to OpenShift Dedicated with their GitHub credentials. To prevent anyone with any GitHub user ID from logging in to your OpenShift Dedicated cluster, you must restrict access to only those in specific GitHub organizations or teams.

### Prerequisites



- The OAuth application must be created directly within the GitHub [organization settings](#) by the GitHub organization administrator.
- [GitHub organizations or teams](#) are set up in your GitHub account.

### Procedure

1. From [OpenShift Cluster Manager](#), navigate to the **Clusters** page and select the cluster that you need to configure identity providers for.
2. Click the **Access control** tab.
3. Click **Add identity provider**.



#### NOTE

You can also click the **Add OAuth configuration** link in the warning message displayed after cluster creation to configure your identity providers.

4. Select **GitHub** from the drop-down menu.
5. Enter a unique name for the identity provider. This name cannot be changed later.
  - An **OAuth callback URL** is automatically generated in the provided field. You will use this to register the GitHub application.

```
https://oauth-openshift.apps.<cluster_name>.  
<cluster_domain>/oauth2callback/<idp_provider_name>
```

For example:

```
https://oauth-openshift.apps.openshift-cluster.example.com/oauth2callback/github
```

6. [Register an application on GitHub](#).
7. Return to OpenShift Dedicated and select a mapping method from the drop-down menu. **Claim** is recommended in most cases.
8. Enter the **Client ID** and **Client secret** provided by GitHub.
9. Enter a **hostname**. A hostname must be entered when using a hosted instance of GitHub Enterprise.
10. Optional: You can use a certificate authority (CA) file to validate server certificates for the configured GitHub Enterprise URL. Click **Browse** to locate and attach a **CA file** to the identity provider.
11. Select **Use organizations** or **Use teams** to restrict access to a particular GitHub organization or a GitHub team.
12. Enter the name of the organization or team you would like to restrict access to. Click **Add more** to specify multiple organizations or teams that users can be a member of.
13. Click **Confirm**.

## Verification

- The configured identity provider is now visible on the **Access control** tab of the **Clusters** page.

## 3.3. CONFIGURING A GITLAB IDENTITY PROVIDER

Configure a GitLab identity provider to use [GitLab.com](https://gitlab.com) or any other GitLab instance as an identity provider.

### Prerequisites

- If you use GitLab version 7.7.0 to 11.0, you connect using the [OAuth integration](#). If you use GitLab version 11.1 or later, you can use [OpenID Connect \(OIDC\)](#) to connect instead of OAuth.

### Procedure

1. From [OpenShift Cluster Manager](#), navigate to the **Clusters** page and select the cluster that you need to configure identity providers for.
2. Click the **Access control** tab.
3. Click **Add identity provider**.



#### NOTE

You can also click the **Add OAuth configuration** link in the warning message displayed after cluster creation to configure your identity providers.

4. Select **GitLab** from the drop-down menu.
5. Enter a unique name for the identity provider. This name cannot be changed later.
  - An **OAuth callback URL** is automatically generated in the provided field. You will provide this URL to GitLab.

```
https://oauth-openshift.apps.<cluster_name>.  
<cluster_domain>/oauth2callback/<idp_provider_name>
```

For example:

```
https://oauth-openshift.apps.openshift-cluster.example.com/oauth2callback/gitlab
```

6. [Add a new application in GitLab](#) .
7. Return to OpenShift Dedicated and select a mapping method from the drop-down menu. **Claim** is recommended in most cases.
8. Enter the **Client ID** and **Client secret** provided by GitLab.
9. Enter the **URL** of your GitLab provider.
10. Optional: You can use a certificate authority (CA) file to validate server certificates for the configured GitLab URL. Click **Browse** to locate and attach a **CA file** to the identity provider.

11. Click **Confirm**.

### Verification

- The configured identity provider is now visible on the **Access control** tab of the **Clusters** page.

## 3.4. CONFIGURING A GOOGLE IDENTITY PROVIDER

Configure a Google identity provider to allow users to authenticate with their Google credentials.



### WARNING

Using Google as an identity provider allows any Google user to authenticate to your server. You can limit authentication to members of a specific hosted domain with the **hostedDomain** configuration attribute.

### Procedure

1. From [OpenShift Cluster Manager](#), navigate to the **Clusters** page and select the cluster that you need to configure identity providers for.
2. Click the **Access control** tab.
3. Click **Add identity provider**.



### NOTE

You can also click the **Add OAuth configuration** link in the warning message displayed after cluster creation to configure your identity providers.

4. Select **Google** from the drop-down menu.
5. Enter a unique name for the identity provider. This name cannot be changed later.
  - An **OAuth callback URL** is automatically generated in the provided field. You will provide this URL to Google.

```
https://oauth-openshift.apps.<cluster_name>.  
<cluster_domain>/oauth2callback/<idp_provider_name>
```

For example:

```
https://oauth-openshift.apps.openshift-cluster.example.com/oauth2callback/github
```

6. Configure a Google identity provider using [Google's OpenID Connect integration](#).
7. Return to OpenShift Dedicated and select a mapping method from the drop-down menu. **Claim** is recommended in most cases.

8. Enter the **Client ID** of a registered Google project and the **Client secret** issued by Google.
9. Enter a hosted domain to restrict users to a Google Apps domain.
10. Click **Confirm**.

### Verification

- The configured identity provider is now visible on the **Access control** tab of the **Clusters** page.

## 3.5. CONFIGURING A LDAP IDENTITY PROVIDER

Configure the LDAP identity provider to validate user names and passwords against an LDAPv3 server, using simple bind authentication.

### Prerequisites

- When configuring a LDAP identity provider, you will need to enter a configured **LDAP URL**. The configured URL is an RFC 2255 URL, which specifies the LDAP host and search parameters to use. The syntax of the URL is:

```
ldap://host:port/basedn?attribute?scope?filter
```

URL component	Description
<b>ldap</b>	For regular LDAP, use the string <b>ldap</b> . For secure LDAP (LDAPS), use <b>ldaps</b> instead.
<b>host:port</b>	The name and port of the LDAP server. Defaults to <b>localhost:389</b> for ldap and <b>localhost:636</b> for LDAPS.
<b>basedn</b>	The DN of the branch of the directory where all searches should start from. At the very least, this must be the top of your directory tree, but it could also specify a subtree in the directory.
<b>attribute</b>	The attribute to search for. Although RFC 2255 allows a comma-separated list of attributes, only the first attribute will be used, no matter how many are provided. If no attributes are provided, the default is to use <b>uid</b> . It is recommended to choose an attribute that will be unique across all entries in the subtree you will be using.
<b>scope</b>	The scope of the search. Can be either <b>one</b> or <b>sub</b> . If the scope is not provided, the default is to use a scope of <b>sub</b> .
<b>filter</b>	A valid LDAP search filter. If not provided, defaults to <b>(objectClass=*)</b>

When doing searches, the attribute, filter, and provided user name are combined to create a search filter that looks like:

```
(<filter>(<attribute>=<username>))
```

**IMPORTANT**

If the LDAP directory requires authentication to search, specify a **bindDN** and **bindPassword** to use to perform the entry search.

**Procedure**

1. From [OpenShift Cluster Manager](#), navigate to the **Clusters** page and select the cluster that you need to configure identity providers for.
2. Click the **Access control** tab.
3. Click **Add identity provider**.

**NOTE**

You can also click the **Add OAuth configuration** link in the warning message displayed after cluster creation to configure your identity providers.

4. Select **LDAP** from the drop-down menu.
5. Enter a unique name for the identity provider. This name cannot be changed later.
6. Select a mapping method from the drop-down menu. **Claim** is recommended in most cases.
7. Enter a **LDAP URL** to specify the LDAP search parameters to use.
8. Optional: Enter a **Bind DN** and **Bind password**.
9. Enter the attributes that will map LDAP attributes to identities.
  - Enter an **ID** attribute whose value should be used as the user ID. Click **Add more** to add multiple ID attributes.
  - Optional: Enter a **Preferred username** attribute whose value should be used as the display name. Click **Add more** to add multiple preferred username attributes.
  - Optional: Enter an **Email** attribute whose value should be used as the email address. Click **Add more** to add multiple email attributes.
10. Optional: Click **Show advanced Options** to add a certificate authority (CA) file to your LDAP identity provider to validate server certificates for the configured URL. Click **Browse** to locate and attach a **CA file** to the identity provider.
11. Optional: Under the advanced options, you can choose to make the LDAP provider **Insecure**. If you select this option, a CA file cannot be used.

**IMPORTANT**

If you are using an insecure LDAP connection (ldap:// or port 389), then you must check the **Insecure** option in the configuration wizard.


12. Click **Confirm**.

**Verification**

- The configured identity provider is now visible on the **Access control** tab of the **Clusters** page.

### 3.6. CONFIGURING AN OPENID IDENTITY PROVIDER

Configure an OpenID identity provider to integrate with an OpenID Connect identity provider using an [Authorization Code Flow](#).



#### IMPORTANT

The Authentication Operator in OpenShift Dedicated requires that the configured OpenID Connect identity provider implements the [OpenID Connect Discovery](#) specification.

Claims are read from the JWT **id\_token** returned from the OpenID identity provider and, if specified, from the JSON returned by the Issuer URL.

At least one claim must be configured to use as the user's identity.

You can also indicate which claims to use as the user's preferred user name, display name, and email address. If multiple claims are specified, the first one with a non-empty value is used. The standard claims are:

Claim	Description
<b>preferred_username</b>	The preferred user name when provisioning a user. A shorthand name that the user wants to be referred to as, such as <b>janedoe</b> . Typically a value that corresponding to the user's login or username in the authentication system, such as username or email.
<b>email</b>	Email address.
<b>name</b>	Display name.

See the [OpenID claims documentation](#) for more information.

#### Prerequisites

- Before you configure OpenID Connect, check the installation prerequisites for any Red Hat product or service you want to use with your OpenShift Dedicated cluster.

#### Procedure

1. From [OpenShift Cluster Manager](#), navigate to the **Clusters** page and select the cluster that you need to configure identity providers for.
2. Click the **Access control** tab.
3. Click **Add identity provider**.

**NOTE**

You can also click the **Add OAuth configuration** link in the warning message displayed after cluster creation to configure your identity providers.

4. Select **OpenID** from the drop-down menu.
5. Enter a unique name for the identity provider. This name cannot be changed later.
  - An **OAuth callback URL** is automatically generated in the provided field.
 

```
https://oauth-openshift.apps.<cluster_name>.<cluster_domain>/oauth2callback/<idp_provider_name>
```

For example:


```
https://oauth-openshift.apps.openshift-cluster.example.com/oauth2callback/openid
```
6. [Create an authorization request using an Authorization Code Flow](#) .
7. Return to OpenShift Dedicated and select a mapping method from the drop-down menu. **Claim** is recommended in most cases.
8. Enter a **Client ID** and **Client secret** provided from OpenID.
9. Enter an **Issuer URL**. This is the URL that the OpenID provider asserts as the Issuer Identifier. It must use the https scheme with no URL query parameters or fragments.
10. Enter an **Email** attribute whose value should be used as the email address. Click **Add more** to add multiple email attributes.
11. Enter a **Name** attribute whose value should be used as the preferred username. Click **Add more** to add multiple preferred usernames.
12. Enter a **Preferred username** attribute whose value should be used as the display name. Click **Add more** to add multiple display names.
13. Optional: Click **Show advanced Options** to add a certificate authority (CA) file to your OpenID identity provider.
14. Optional: Under the advanced options, you can add **Additional scopes**. By default, the **OpenID** scope is requested.
15. Click **Confirm**.

**Verification**

- The configured identity provider is now visible on the **Access control** tab of the **Clusters** page.

**3.7. CONFIGURING AN HTPASSWD IDENTITY PROVIDER**

Configure an HTPasswd identity provider to create a single, static user with cluster administration privileges. You can log in to your cluster as the user to troubleshoot issues.




## IMPORTANT

The HTPasswd identity provider option is included only to enable the creation of a single, static administration user. HTPasswd is not supported as a general-use identity provider for OpenShift Dedicated.

### Procedure

1. From [OpenShift Cluster Manager](#), navigate to the **Clusters** page and select your cluster.
2. Select **Access control** → **Identity providers**.
3. Click **Add identity provider**.
4. Select **HTPasswd** from the **Identity Provider** drop-down menu.
5. Add a unique name in the **Name** field for the identity provider.
6. Use the suggested username and password for the static user, or create your own.



## NOTE

The credentials defined in this step are not visible after you select **Add** in the following step. If you lose the credentials, you must recreate the identity provider and define the credentials again.

7. Select **Add** to create the HTPasswd identity provider and the single, static user.
8. Grant the static user permission to manage the cluster:
  - a. Under **Access control** → **Cluster Roles and Access**, select **Add user**.
  - b. Enter the **User ID** of the static user that you created in the preceding step.
  - c. Select a **Group**.
    - If you are installing OpenShift Dedicated using the Customer Cloud Subscription (CCS) infrastructure type, choose either the **dedicated-admins** or **cluster-admins** group. Users in the **dedicated-admins** group have standard administrative privileges for OpenShift Dedicated. Users in the **cluster-admins** group have full administrative access to the cluster.
    - If you are installing OpenShift Dedicated using the Red Hat cloud account infrastructure type, the **dedicated-admins** group is automatically selected.
  - d. Select **Add user** to grant the administration privileges to the user.

### Verification

- The configured HTPasswd identity provider is visible on the **Access control** → **Identity providers** page.



**NOTE**

After creating the identity provider, synchronization usually completes within two minutes. You can log in to the cluster as the user after the HTPasswd identity provider becomes available.

- The single, administrative user is visible on the **Access control** → **Cluster Roles and Access** page. The administration group membership of the user is also displayed.

## 3.8. ACCESSING YOUR CLUSTER

After you have configured your identity providers, users can access the cluster from Red Hat OpenShift Cluster Manager.

### Prerequisites

- You logged in to [OpenShift Cluster Manager](#).
- You created an OpenShift Dedicated cluster.
- You configured an identity provider for your cluster.
- You added your user account to the configured identity provider.

### Procedure

1. From [OpenShift Cluster Manager](#), click on the cluster you want to access.
2. Click **Open Console**.
3. Click on your identity provider and provide your credentials to log into the cluster.
4. Click **Open console** to open the web console for your cluster.
5. Click on your identity provider and provide your credentials to log in to the cluster. Complete any authorization requests that are presented by your provider.