



OpenShift Dedicated 4

Authentication and authorization

Securing OpenShift Dedicated.

OpenShift Dedicated 4 Authentication and authorization

Securing OpenShift Dedicated.

Legal Notice

Copyright © Red Hat.

Except as otherwise noted below, the text of and illustrations in this documentation are licensed by Red Hat under the Creative Commons Attribution–Share Alike 3.0 Unported license . 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, the Red Hat logo, JBoss, Hibernate, and RHCE are trademarks or registered trademarks of Red Hat, LLC. or its subsidiaries in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

XFS is a trademark or registered trademark of Hewlett Packard Enterprise Development LP or its subsidiaries in the United States and other countries.

The OpenStack[®] Word Mark and OpenStack logo are trademarks or registered trademarks of the Linux Foundation, used under license.

All other trademarks are the property of their respective owners.

Abstract

This document provides information about securing OpenShift Dedicated clusters.

Table of Contents

CHAPTER 1. OVERVIEW OF AUTHENTICATION AND AUTHORIZATION	5
1.1. GLOSSARY OF COMMON TERMS FOR OPENSIFT DEDICATED AUTHENTICATION AND AUTHORIZATION	5
1.2. ABOUT AUTHENTICATION IN OPENSIFT DEDICATED	6
1.3. ABOUT AUTHORIZATION IN OPENSIFT DEDICATED	6
CHAPTER 2. UNDERSTANDING AUTHENTICATION	8
2.1. USERS	8
2.2. GROUPS	8
2.3. API AUTHENTICATION	9
2.3.1. OpenShift Dedicated OAuth server	9
2.3.1.1. OAuth token requests	9
CHAPTER 3. MANAGING USER-OWNED OAUTH ACCESS TOKENS	11
3.1. LISTING USER-OWNED OAUTH ACCESS TOKENS	11
3.2. VIEWING THE DETAILS OF A USER-OWNED OAUTH ACCESS TOKEN	11
3.3. DELETING USER-OWNED OAUTH ACCESS TOKENS	12
3.4. ADDING UNAUTHENTICATED GROUPS TO CLUSTER ROLES	13
CHAPTER 4. IDENTITY PROVIDERS OVERVIEW	15
4.1. UNDERSTANDING IDENTITY PROVIDERS	15
4.1.1. Supported identity providers	15
4.1.2. Identity provider parameters	15
4.2. CONFIGURING A GITHUB IDENTITY PROVIDER	16
4.3. CONFIGURING A GITLAB IDENTITY PROVIDER	17
4.4. CONFIGURING A GOOGLE IDENTITY PROVIDER	19
4.5. CONFIGURING A LDAP IDENTITY PROVIDER	20
4.6. CONFIGURING AN OPENID IDENTITY PROVIDER	21
4.7. CONFIGURING AN HTPASSWD IDENTITY PROVIDER	23
4.7.1. Configuring an htpasswd identity provider	24
4.7.2. Configuring an htpasswd identity provider with the CLI	24
4.8. ACCESSING YOUR CLUSTER	25
4.9. ADDITIONAL RESOURCES	25
CHAPTER 5. REVOKING PRIVILEGES AND ACCESS TO AN OPENSIFT DEDICATED CLUSTER	26
5.1. REVOKING ADMINISTRATOR PRIVILEGES FROM A USER	26
5.2. REVOKING USER ACCESS TO A CLUSTER	26
CHAPTER 6. MANAGING ADMINISTRATION ROLES AND USERS	28
6.1. UNDERSTANDING ADMINISTRATION ROLES	28
6.1.1. The cluster-admin role	28
6.1.2. The dedicated-admin role	28
6.2. MANAGING OPENSIFT DEDICATED ADMINISTRATORS	28
CHAPTER 7. USING RBAC TO DEFINE AND APPLY PERMISSIONS	30
7.1. RBAC OVERVIEW	30
7.1.1. Default cluster roles	31
7.1.2. Evaluating authorization	32
7.1.2.1. Cluster role aggregation	33
7.2. PROJECTS AND NAMESPACES	33
7.3. DEFAULT PROJECTS	34
7.4. VIEWING CLUSTER ROLES AND BINDINGS	35
7.5. VIEWING LOCAL ROLES AND BINDINGS	41

7.6. ADDING ROLES TO USERS	43
7.7. CREATING A LOCAL ROLE	45
7.8. LOCAL ROLE BINDING COMMANDS	46
7.9. CLUSTER ROLE BINDING COMMANDS	47
7.10. GRANTING ADMINISTRATOR PRIVILEGES TO A USER	47
7.11. CLUSTER ROLE BINDINGS FOR UNAUTHENTICATED GROUPS	48
CHAPTER 8. UNDERSTANDING AND CREATING SERVICE ACCOUNTS	49
8.1. SERVICE ACCOUNTS OVERVIEW	49
8.1.1. Automatically generated image pull secrets	49
8.2. CREATING SERVICE ACCOUNTS	50
8.3. GRANTING ROLES TO SERVICE ACCOUNTS	51
CHAPTER 9. USING SERVICE ACCOUNTS IN APPLICATIONS	54
9.1. SERVICE ACCOUNTS OVERVIEW	54
9.2. DEFAULT SERVICE ACCOUNTS	54
9.2.1. Default cluster service accounts	54
9.2.2. Default project service accounts and roles	54
9.2.3. Automatically generated image pull secrets	55
9.3. CREATING SERVICE ACCOUNTS	56
CHAPTER 10. USING A SERVICE ACCOUNT AS AN OAUTH CLIENT	58
10.1. SERVICE ACCOUNTS AS OAUTH CLIENTS	58
10.1.1. Redirect URIs for service accounts as OAuth clients	58
CHAPTER 11. SCOPING TOKENS	61
11.1. ABOUT SCOPING TOKENS	61
11.1.1. User scopes	61
11.1.2. Role scope	61
11.2. ADDING UNAUTHENTICATED GROUPS TO CLUSTER ROLES	61
CHAPTER 12. USING BOUND SERVICE ACCOUNT TOKENS	63
12.1. ABOUT BOUND SERVICE ACCOUNT TOKENS	63
12.2. CONFIGURING BOUND SERVICE ACCOUNT TOKENS USING VOLUME PROJECTION	63
12.3. CREATING BOUND SERVICE ACCOUNT TOKENS OUTSIDE THE POD	64
CHAPTER 13. MANAGING SECURITY CONTEXT CONSTRAINTS	66
13.1. ABOUT SECURITY CONTEXT CONSTRAINTS	66
13.1.1. Default security context constraints	67
13.1.2. Security context constraints settings	69
13.1.3. Security context constraints strategies	70
13.1.4. Controlling volumes for CCS clusters	71
13.1.5. Admission control	73
13.1.6. Security context constraints prioritization	74
13.2. ABOUT PRE-ALLOCATED SECURITY CONTEXT CONSTRAINTS VALUES	74
13.3. EXAMPLE SECURITY CONTEXT CONSTRAINTS	75
13.4. CREATING SECURITY CONTEXT CONSTRAINTS FOR CUSTOMER CLOUD SUBSCRIPTION CLUSTERS	78
13.5. CONFIGURING A WORKLOAD TO REQUIRE A SPECIFIC SCC	80
13.6. ROLE-BASED ACCESS TO SECURITY CONTEXT CONSTRAINTS	81
13.7. REFERENCE OF SECURITY CONTEXT CONSTRAINTS COMMANDS	82
13.7.1. Listing security context constraints	82
13.7.2. Examining security context constraints	83
13.8. ADDITIONAL RESOURCES	84

CHAPTER 14. UNDERSTANDING AND MANAGING POD SECURITY ADMISSION	85
14.1. ABOUT POD SECURITY ADMISSION	85
14.1.1. Pod security admission modes	85
14.1.2. Pod security admission profiles	85
14.1.3. Privileged namespaces	86
14.1.4. Pod security admission and security context constraints	86
14.2. ABOUT POD SECURITY ADMISSION SYNCHRONIZATION	86
14.2.1. Pod security admission synchronization namespace exclusions	87
14.3. CONTROLLING POD SECURITY ADMISSION SYNCHRONIZATION	87
14.4. CONFIGURING POD SECURITY ADMISSION FOR A NAMESPACE	88
14.5. ABOUT POD SECURITY ADMISSION ALERTS	88
14.6. ADDITIONAL RESOURCES	88
CHAPTER 15. SYNCING LDAP GROUPS	90
15.1. ABOUT CONFIGURING LDAP SYNC	90
15.1.1. About the RFC 2307 configuration file	92
15.1.2. About the Active Directory configuration file	93
15.1.3. About the augmented Active Directory configuration file	94
15.2. RUNNING LDAP SYNC	95
15.2.1. Syncing the LDAP server with OpenShift Dedicated	95
15.2.2. Syncing OpenShift Dedicated groups with the LDAP server	95
15.2.3. Syncing subgroups from the LDAP server with OpenShift Dedicated	96
15.3. RUNNING A GROUP PRUNING JOB	97
15.4. LDAP GROUP SYNC EXAMPLES	97
15.4.1. Syncing groups using the RFC 2307 schema	97
15.4.2. Syncing groups using the RFC2307 schema with user-defined name mappings	99
15.4.3. Syncing groups using RFC 2307 with user-defined error tolerances	100
15.4.4. Syncing groups using the Active Directory schema	103
15.4.5. Syncing groups using the augmented Active Directory schema	105
15.4.5.1. LDAP nested membership sync example	106
15.5. LDAP SYNC CONFIGURATION SPECIFICATION	110
15.5.1. v1.LDAPSyncConfig	110
15.5.2. v1.StringSource	112
15.5.3. v1.LDAPQuery	113
15.5.4. v1.RFC2307Config	114
15.5.5. v1.ActiveDirectoryConfig	115
15.5.6. v1.AugmentedActiveDirectoryConfig	116

CHAPTER 1. OVERVIEW OF AUTHENTICATION AND AUTHORIZATION

1.1. GLOSSARY OF COMMON TERMS FOR OPENSIFT DEDICATED AUTHENTICATION AND AUTHORIZATION

This glossary defines common terms that are used in OpenShift Dedicated authentication and authorization.

authentication

An authentication determines access to an OpenShift Dedicated cluster and ensures only authenticated users access the OpenShift Dedicated cluster.

authorization

Authorization determines whether the identified user has permissions to perform the requested action.

bearer token

Bearer token is used to authenticate to API with the header **Authorization: Bearer <token>**.

config map

A config map provides a way to inject configuration data into the pods. You can reference the data stored in a config map in a volume of type **ConfigMap**. Applications running in a pod can use this data.

containers

Lightweight and executable images that consist of software and all its dependencies. Because containers virtualize the operating system, you can run containers in a data center, public or private cloud, or your local host.

Custom Resource (CR)

A CR is an extension of the Kubernetes API.

group

A group is a set of users. A group is useful for granting permissions to multiple users one time.

HTPasswd

HTPasswd updates the files that store usernames and password for authentication of HTTP users.

Keystone

Keystone is an Red Hat OpenStack Platform (RHOSP) project that provides identity, token, catalog, and policy services.

Lightweight directory access protocol (LDAP)

LDAP is a protocol that queries user information.

namespace

A namespace isolates specific system resources that are visible to all processes. Inside a namespace, only processes that are members of that namespace can see those resources.

node

A node is a worker machine in the OpenShift Dedicated cluster. A node is either a virtual machine (VM) or a physical machine.

OAuth client

OAuth client is used to get a bearer token.

OAuth server

The OpenShift Dedicated control plane includes a built-in OAuth server that determines the user's identity from the configured identity provider and creates an access token.

OpenID Connect

The OpenID Connect is a protocol to authenticate the users to use single sign-on (SSO) to access sites that use OpenID Providers.

pod

A pod is the smallest logical unit in Kubernetes. A pod is comprised of one or more containers to run in a worker node.

regular users

Users that are created automatically in the cluster upon first login or via the API.

request header

A request header is an HTTP header that is used to provide information about HTTP request context, so that the server can track the response of the request.

role-based access control (RBAC)

A key security control to ensure that cluster users and workloads have access to only the resources required to execute their roles.

service accounts

Service accounts are used by the cluster components or applications.

system users

Users that are created automatically when the cluster is installed.

users

Users is an entity that can make requests to API.

1.2. ABOUT AUTHENTICATION IN OPENSIFT DEDICATED

To control access to an OpenShift Dedicated cluster, an administrator with the **dedicated-admin** role can configure [user authentication](#) and ensure only approved users access the cluster.

To interact with an OpenShift Dedicated cluster, users must first authenticate to the OpenShift Dedicated API in some way. You can authenticate by providing an [OAuth access token or an X.509 client certificate](#) in your requests to the OpenShift Dedicated API.



NOTE

If you do not present a valid access token or certificate, your request is unauthenticated and you receive an HTTP 401 error.

An administrator can configure authentication by configuring an identity provider. You can define any [supported identity provider in OpenShift Dedicated](#) and add it to your cluster.

1.3. ABOUT AUTHORIZATION IN OPENSIFT DEDICATED

Authorization involves determining whether the identified user has permissions to perform the requested action.

Administrators can define permissions and assign them to users using the [RBAC objects, such as rules, roles, and bindings](#). To understand how authorization works in OpenShift Dedicated, see [Evaluating authorization](#).

You can also control access to an OpenShift Dedicated cluster through [projects and namespaces](#).

Along with controlling user access to a cluster, you can also control the actions a pod can perform and the resources it can access using [security context constraints \(SCCs\)](#).

You can manage authorization for OpenShift Dedicated through the following tasks:

- Viewing [local](#) and [cluster](#) roles and bindings.
- Creating a [local role](#) and assigning it to a user or group.
- Assigning a cluster role to a user or group: OpenShift Dedicated includes a set of [default cluster roles](#). You can [add them to a user or group](#).
- Granting administrator privileges to users: You can [grant **dedicated-admin** privileges to users](#).
- Creating service accounts: [Service accounts](#) provide a flexible way to control API access without sharing a regular user's credentials. A user can [create and use a service account in applications](#) and also as [an OAuth client](#).
- [Scoping tokens](#): A scoped token is a token that identifies as a specific user who can perform only specific operations. You can create scoped tokens to delegate some of your permissions to another user or a service account.
- Syncing LDAP groups: You can manage user groups in one place by [syncing the groups stored in an LDAP server](#) with the OpenShift Dedicated user groups.

CHAPTER 2. UNDERSTANDING AUTHENTICATION

For users to interact with OpenShift Dedicated, they must first authenticate to the cluster. The authentication layer identifies the user associated with requests to the OpenShift Dedicated API. The authorization layer then uses information about the requesting user to determine if the request is allowed.

2.1. USERS

A *user* in OpenShift Dedicated is an entity that can make requests to the OpenShift Dedicated API. An OpenShift Dedicated **User** object represents an actor which can be granted permissions in the system by adding roles to them or to their groups. Typically, this represents the account of a developer or administrator that is interacting with OpenShift Dedicated.

Several types of users can exist:

User type	Description
Regular users	This is the way most interactive OpenShift Dedicated users are represented. Regular users are created automatically in the system upon first login or can be created via the API. Regular users are represented with the User object. Examples: joe alice
System users	Many of these are created automatically when the infrastructure is defined, mainly for the purpose of enabling the infrastructure to interact with the API securely. They include a cluster administrator (with access to everything), a per-node user, users for use by routers and registries, and various others. Finally, there is an anonymous system user that is used by default for unauthenticated requests. Examples: system:admin system:openshift-registry system:node:node1.example.com
Service accounts	These are special system users associated with projects; some are created automatically when the project is first created, while project administrators can create more for the purpose of defining access to the contents of each project. Service accounts are represented with the ServiceAccount object. Examples: system:serviceaccount:default:deployer system:serviceaccount:foo:builder

Each user must authenticate in some way to access OpenShift Dedicated. API requests with no authentication or invalid authentication are authenticated as requests by the **anonymous** system user. After authentication, policy determines what the user is authorized to do.

2.2. GROUPS

A user can be assigned to one or more *groups*, each of which represent a certain set of users. Groups are useful when managing authorization policies to grant permissions to multiple users at once, for example allowing access to objects within a project, versus granting them to users individually.

In addition to explicitly defined groups, there are also system groups, or *virtual groups*, that are automatically provisioned by the cluster.

The following default virtual groups are most important:

Virtual group	Description
system:authenticated	Automatically associated with all authenticated users.
system:authenticated:oauth	Automatically associated with all users authenticated with an OAuth access token.
system:unauthenticated	Automatically associated with all unauthenticated users.

2.3. API AUTHENTICATION

Requests to the OpenShift Dedicated API are authenticated using the following methods:

OAuth access tokens

- Obtained from the OpenShift Dedicated OAuth server using the `<namespace_route>/oauth/authorize` and `<namespace_route>/oauth/token` endpoints.
- Sent as an **Authorization: Bearer...** header.
- Sent as a websocket subprotocol header in the form **base64url.bearer.authorization.k8s.io.<base64url-encoded-token>** for websocket requests.

X.509 client certificates

- Requires an HTTPS connection to the API server.
- Verified by the API server against a trusted certificate authority bundle.
- The API server creates and distributes certificates to controllers to authenticate themselves.

Any request with an invalid access token or an invalid certificate is rejected by the authentication layer with a **401** error.

If no access token or certificate is presented, the authentication layer assigns the **system:anonymous** virtual user and the **system:unauthenticated** virtual group to the request. This allows the authorization layer to determine which requests, if any, an anonymous user is allowed to make.

2.3.1. OpenShift Dedicated OAuth server

The OpenShift Dedicated Control Plane includes a built-in OAuth server. Users obtain OAuth access tokens to authenticate themselves to the API.

When a person requests a new OAuth token, the OAuth server uses the configured identity provider to determine the identity of the person making the request.

It then determines what user that identity maps to, creates an access token for that user, and returns the token for use.

2.3.1.1. OAuth token requests

Every request for an OAuth token must specify the OAuth client that will receive and use the token. The following OAuth clients are automatically created when starting the OpenShift Dedicated API:

OAuth client	Usage
openshift-browser-client	Requests tokens at <namespace_route>/oauth/token/request with a user-agent that can handle interactive logins. ^[1]
openshift-challenging-client	Requests tokens with a user-agent that can handle WWW-Authenticate challenges.

1. **<namespace_route>** refers to the namespace route. This is found by running the following command:

```
$ oc get route oauth-openshift -n openshift-authentication -o json | jq .spec.host
```

All requests for OAuth tokens involve a request to **<namespace_route>/oauth/authorize**. Most authentication integrations place an authenticating proxy in front of this endpoint, or configure OpenShift Dedicated to validate credentials against a backing identity provider. Requests to **<namespace_route>/oauth/authorize** can come from user-agents that cannot display interactive login pages, such as the CLI. Therefore, OpenShift Dedicated supports authenticating using a **WWW-Authenticate** challenge in addition to interactive login flows.

If an authenticating proxy is placed in front of the **<namespace_route>/oauth/authorize** endpoint, it sends unauthenticated, non-browser user-agents **WWW-Authenticate** challenges rather than displaying an interactive login page or redirecting to an interactive login flow.



NOTE

To prevent cross-site request forgery (CSRF) attacks against browser clients, only send Basic authentication challenges with if a **X-CSRF-Token** header is on the request. Clients that expect to receive Basic **WWW-Authenticate** challenges must set this header to a non-empty value.

If the authenticating proxy cannot support **WWW-Authenticate** challenges, or if OpenShift Dedicated is configured to use an identity provider that does not support WWW-Authenticate challenges, you must use a browser to manually obtain a token from **<namespace_route>/oauth/token/request**.

CHAPTER 3. MANAGING USER-OWNED OAUTH ACCESS TOKENS

Users can review their own OAuth access tokens and delete any that are no longer needed.

3.1. LISTING USER-OWNED OAUTH ACCESS TOKENS

You can list your user-owned OAuth access tokens. Token names are not sensitive and cannot be used to log in.

Procedure

- List all user-owned OAuth access tokens:

```
$ oc get useroauthaccesstokens
```

Example output

```
NAME      CLIENT NAME          CREATED          EXPIRES
REDIRECT URI          SCOPES
<token1> openshift-challenging-client 2021-01-11T19:25:35Z 2021-01-12 19:25:35
+0000 UTC https://oauth-openshift.apps.example.com/oauth/token/implicit user:full
<token2> openshift-browser-client   2021-01-11T19:27:06Z 2021-01-12 19:27:06 +0000
UTC https://oauth-openshift.apps.example.com/oauth/token/display user:full
<token3> console                2021-01-11T19:26:29Z 2021-01-12 19:26:29 +0000 UTC
https://console-openshift-console.apps.example.com/auth/callback user:full
```

- List user-owned OAuth access tokens for a particular OAuth client:

```
$ oc get useroauthaccesstokens --field-selector=clientName="console"
```

Example output

```
NAME      CLIENT NAME          CREATED          EXPIRES
REDIRECT URI          SCOPES
<token3> console                2021-01-11T19:26:29Z 2021-01-12 19:26:29 +0000 UTC
https://console-openshift-console.apps.example.com/auth/callback user:full
```

3.2. VIEWING THE DETAILS OF A USER-OWNED OAUTH ACCESS TOKEN

You can view the details of a user-owned OAuth access token.

Procedure

- Describe the details of a user-owned OAuth access token:

```
$ oc describe useroauthaccesstokens <token_name>
```

Example output

```
-
```

```
Name: <token_name> 1
Namespace:
Labels: <none>
Annotations: <none>
API Version: oauth.openshift.io/v1
Authorize Token: sha256~Ksckkug-9Fg_RWn_AUysPolg-_HqmFI9zUL_CgD8wr8
Client Name: openshift-browser-client 2
Expires In: 86400 3
Inactivity Timeout Seconds: 317 4
Kind: UserOAuthAccessToken
Metadata:
  Creation Timestamp: 2021-01-11T19:27:06Z
  Managed Fields:
    API Version: oauth.openshift.io/v1
    Fields Type: FieldsV1
    fieldsV1:
      f:authorizeToken:
      f:clientName:
      f:expiresIn:
      f:redirectURI:
      f:scopes:
      f:userName:
      f:userUID:
    Manager: oauth-server
    Operation: Update
    Time: 2021-01-11T19:27:06Z
  Resource Version: 30535
  Self Link: /apis/oauth.openshift.io/v1/useroauthaccesstokens/<token_name>
  UID: f9d00b67-ab65-489b-8080-e427fa3c6181
  Redirect URI: https://oauth-openshift.apps.example.com/oauth/token/display
  Scopes:
    user:full 5
  User Name: <user_name> 6
  User UID: 82356ab0-95f9-4fb3-9bc0-10f1d6a6a345
  Events: <none>
```

- 1 The token name, which is the sha256 hash of the token. Token names are not sensitive and cannot be used to log in.
- 2 The client name, which describes where the token originated from.
- 3 The value in seconds from the creation time before this token expires.
- 4 If there is a token inactivity timeout set for the OAuth server, this is the value in seconds from the creation time before this token can no longer be used.
- 5 The scopes for this token.
- 6 The user name associated with this token.

3.3. DELETING USER-OWNED OAUTH ACCESS TOKENS

The **oc logout** command only invalidates the OAuth token for the active session. You can use the following procedure to delete any user-owned OAuth tokens that are no longer needed.

Deleting an OAuth access token logs out the user from all sessions that use the token.

Procedure

- Delete the user-owned OAuth access token:

```
$ oc delete useroauthaccesstokens <token_name>
```

Example output


```
useroauthaccesstoken.oauth.openshift.io "<token_name>" deleted
```

3.4. ADDING UNAUTHENTICATED GROUPS TO CLUSTER ROLES

As a cluster administrator, you can add unauthenticated users to the following cluster roles in OpenShift Dedicated by creating a cluster role binding. Unauthenticated users do not have access to non-public cluster roles. This should only be done in specific use cases when necessary.

You can add unauthenticated users to the following cluster roles:

- **system:scope-impersonation**
- **system:webhook**
- **system:oauth-token-deleter**
- **self-access-reviewer**



IMPORTANT

Always verify compliance with your organization's security standards when modifying unauthenticated access.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- You have installed the OpenShift CLI (**oc**).

Procedure

1. Create a YAML file named **add-<cluster_role>-unauth.yaml** and add the following content:

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  annotations:
    rbac.authorization.kubernetes.io/autoupdate: "true"
  name: <cluster_role>access-unauthenticated
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: <cluster_role>
```

```
subjects:  
- apiGroup: rbac.authorization.k8s.io  
  kind: Group  
  name: system:unauthenticated
```

2. Apply the configuration by running the following command:

```
$ oc apply -f add-<cluster_role>.yaml
```

CHAPTER 4. IDENTITY PROVIDERS OVERVIEW

After you create your OpenShift Dedicated cluster, configure identity providers so users can log in and access the cluster.

4.1. 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.

4.1.1. 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 GitLab.com or any other GitLab instance as an identity provider.
Google	Configure a Google identity provider using Google's OpenID Connect integration .
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 Authorization Code Flow .
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 1740" data-label="Image"> </div> <p>IMPORTANT</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>

4.1.2. Identity provider parameters

The following parameters are common to all identity providers:

Parameter	Description
name	The provider name is prefixed to provider user names to form an identity name.
mappingMethod	<p>Defines how new identities are mapped to users when they log in. Enter one of the following values:</p> <p>claim 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>lookup 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>add 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**.

4.2. 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 **Cluster List** 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 **Cluster List** page.

4.3. 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 **Cluster List** 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 **Cluster List** page.

4.4. 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 **Cluster List** 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/google
```

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 **Cluster List** page.

4.5. 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
ldap	For regular LDAP, use the string ldap . For secure LDAP (LDAPS), use ldaps instead.
host:port	The name and port of the LDAP server. Defaults to localhost:389 for ldap and localhost:636 for LDAPS.
basedn	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.
attribute	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 uid . It is recommended to choose an attribute that will be unique across all entries in the subtree you will be using.
scope	The scope of the search. Can be either one or sub . If the scope is not provided, the default is to use a scope of sub .
filter	A valid LDAP search filter. If not provided, defaults to (objectClass=*)

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 **Cluster List** 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 **Cluster List** page.

4.6. 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
preferred_username	The preferred user name when provisioning a user. A shorthand name that the user wants to be referred to as, such as janedoe . Typically a value that corresponding to the user's login or username in the authentication system, such as username or email.
email	Email address.
name	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 **Cluster List** 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. Register a new OpenID Connect client in the OpenID identity provider by following the steps to [create an authorization request](#).
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 **Cluster List** page.

4.7. 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 problems. You can use the web user interface (UI) or your command-line interface (CLI) to create an htpasswd identity provider.



IMPORTANT

The htpasswd identity provider option is included only to create a single, static administration user. htpasswd is not supported as a general-use identity provider for OpenShift Dedicated.

4.7.1. Configuring an htpasswd identity provider

You can create an htpasswd identity provider with the OpenShift Cluster Manager web user interface (UI).

Procedure

1. Select your cluster from the the **Cluster List** page on [OpenShift Cluster Manager](#).
2. Select **Access control** → **Identity providers**.
3. Click **Add identity provider**.
4. Select **HTPasswd** from the **Identity Provider** list.
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

You cannot retrieve the credentials defined in this step 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. Select **Access control** → **Cluster Roles and Access** > **Add user**.
 - b. Enter the User ID of the static user that you created in the preceding step.
 - c. 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.

4.7.2. Configuring an htpasswd identity provider with the CLI

You can create an htpasswd identity provider (IDP) with the OCM CLI (**ocm**) tool.

Prerequisites

- You have installed and configured the latest version of the OCM CLI (**ocm**).

Procedure

- Run the following command to create an htpasswd IDP by passing the usernames and passwords in the command-line interface:

```
$ ocm create idp --type htpasswd --cluster <cluster_name> --name <idp_name> --username <user_name> --password '<password>'
```



NOTE

You must include the password within quotes like **'password'** to prevent your password from disrupting the Bash commands.

Passwords must include uppercase letters, lowercase letters, and numbers or symbols, specifically, ASCII-standard characters only. The password must be at least 14 characters.

4.8. 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](#), select the cluster you want to access.
2. Click **Open console** to open the web console for your cluster.
3. Select your identity provider and enter your credentials to log in to the cluster. Complete any authorization requests from your provider.

4.9. ADDITIONAL RESOURCES

- [Apache Password Formats](#)

CHAPTER 5. REVOKING PRIVILEGES AND ACCESS TO AN OPENSIFT DEDICATED CLUSTER

As cluster owner, you can revoke admin privileges and user access to a OpenShift Dedicated cluster.

5.1. REVOKING ADMINISTRATOR PRIVILEGES FROM A USER


After you have granted **dedicated-admin** privileges to a user, you can revoke those privileges when they are no longer needed.

Prerequisites

- You logged in to [OpenShift Cluster Manager](#).
- You created an OpenShift Dedicated cluster.
- You have configured a GitHub identity provider for your cluster and added an identity provider user.
- You granted **dedicated-admin** privileges to a user.

Procedure

1. Navigate to [OpenShift Cluster Manager](#) and select your cluster.
2. Click the **Access control** tab.

3. In the **Cluster Roles and Access** tab, select  next to a user and click **Delete**.

Verification

- After revoking the privileges, the user is no longer listed as part of the **dedicated-admins** group under **Access control** → **Cluster Roles and Access** on the OpenShift Cluster Manager page for your cluster.

5.2. REVOKING USER ACCESS TO A CLUSTER

You can revoke cluster access from an identity provider user by removing them from your configured identity provider.

You can configure different types of identity providers for your OpenShift Dedicated cluster. The following example procedure revokes cluster access for a member of a GitHub organization or team that is configured for identity provision to the cluster.

Prerequisites

- You have an OpenShift Dedicated cluster.
- You have a GitHub user account.
- You have configured a GitHub identity provider for your cluster and added an identity provider user.

Procedure

1. Navigate to github.com and log in to your GitHub account.
2. Remove the user from your GitHub organization or team:
 - If your identity provider configuration uses a GitHub organization, follow the steps in [Removing a member from your organization](#) in the GitHub documentation.
 - If your identity provider configuration uses a team within a GitHub organization, follow the steps in [Removing organization members from a team](#) in the GitHub documentation.

Verification

- After removing the user from your identity provider, the user cannot authenticate into the cluster.

CHAPTER 6. MANAGING ADMINISTRATION ROLES AND USERS

Keep your OpenShift Dedicated cluster secure and up-to-date by managing administration roles and users effectively.

6.1. UNDERSTANDING ADMINISTRATION ROLES

As an administrator of an OpenShift Dedicated cluster, you have access to the **cluster-admin** and **dedicated-admin** roles. These roles have different permissions and access levels that allow you to manage and configure your cluster effectively. Understanding the differences between these roles can help you assign the appropriate permissions to users and manage your cluster more efficiently.

6.1.1. The cluster-admin role

As an administrator of an OpenShift Dedicated cluster with Customer Cloud Subscriptions (CCS), you have access to the **cluster-admin** role. The user who created the cluster can add the **cluster-admin** user role to an account to have the maximum administrator privileges. These privileges are not automatically assigned to your user account when you create the cluster. While logged in to an account with the cluster-admin role, users have mostly unrestricted access to control and configure the cluster. There are some configurations that are blocked with webhooks to prevent destabilizing the cluster, or because they are managed in [OpenShift Cluster Manager](#) and any in-cluster changes would be overwritten. Usage of the cluster-admin role is subject to the restrictions listed in your Appendix 4 agreement with Red Hat. As a best practice, limit the number of **cluster-admin** users to as few as possible.

6.1.2. The dedicated-admin role

As an administrator of an OpenShift Dedicated cluster, your account has additional permissions and access to all user-created projects in your organization's cluster. While logged in to an account with the **dedicated-admin** role, the developer CLI commands (under the **oc** command) allow you increased visibility and management capabilities over objects across projects, while the administrator CLI commands (under the **oc adm** command) allow you to complete additional operations.



NOTE

While your account does have these increased permissions, the actual cluster maintenance and host configuration is still performed by the Red Hat Site Reliability Engineering (SRE) team.


6.2. MANAGING OPENSIFT DEDICATED ADMINISTRATORS

Administrator roles are managed using a **cluster-admin** or **dedicated-admin** group on the cluster. Existing members of this group can edit membership through [OpenShift Cluster Manager](#).


Procedure

1. Navigate to the **Cluster Details** page and select the **Access Control** tab.
2. Select the **Cluster Roles and Access** tab and click **Add user**.
3. Enter the user name and select your group.

4. Click **Add user**.

**NOTE**

Adding a user to the **cluster-admin** group can take several minutes to complete.

5. Optional: To remove a OpenShift Dedicated administrator, click the Options menu  to the right of the user and group combination and click **Delete**.

CHAPTER 7. USING RBAC TO DEFINE AND APPLY PERMISSIONS

7.1. RBAC OVERVIEW

Role-based access control (RBAC) objects determine whether a user is allowed to perform a given action within a project.

Administrators with the **dedicated-admin** role can use the cluster roles and bindings to control who has various access levels to the OpenShift Dedicated platform itself and all projects.

Developers can use local roles and bindings to control who has access to their projects. Note that authorization is a separate step from authentication, which is more about determining the identity of who is taking the action.

Authorization is managed using:

Authorization object	Description
Rules	Sets of permitted verbs on a set of objects. For example, whether a user or service account can create pods.
Roles	Collections of rules. You can associate, or bind, users and groups to multiple roles.
Bindings	Associations between users and/or groups with a role.

There are two levels of RBAC roles and bindings that control authorization:

RBAC level	Description
Cluster RBAC	Roles and bindings that are applicable across all projects. <i>Cluster roles</i> exist cluster-wide, and <i>cluster role bindings</i> can reference only cluster roles.
Local RBAC	Roles and bindings that are scoped to a given project. While <i>local roles</i> exist only in a single project, local role bindings can reference <i>both</i> cluster and local roles.

A cluster role binding is a binding that exists at the cluster level. A role binding exists at the project level. The cluster role *view* must be bound to a user using a local role binding for that user to view the project. Create local roles only if a cluster role does not provide the set of permissions needed for a particular situation.

This two-level hierarchy allows reuse across multiple projects through the cluster roles while allowing customization inside of individual projects through local roles.

During evaluation, both the cluster role bindings and the local role bindings are used. For example:

1. Cluster-wide "allow" rules are checked.
2. Locally-bound "allow" rules are checked.

- Deny by default.

7.1.1. Default cluster roles

OpenShift Dedicated includes a set of default cluster roles that you can bind to users and groups cluster-wide or locally.



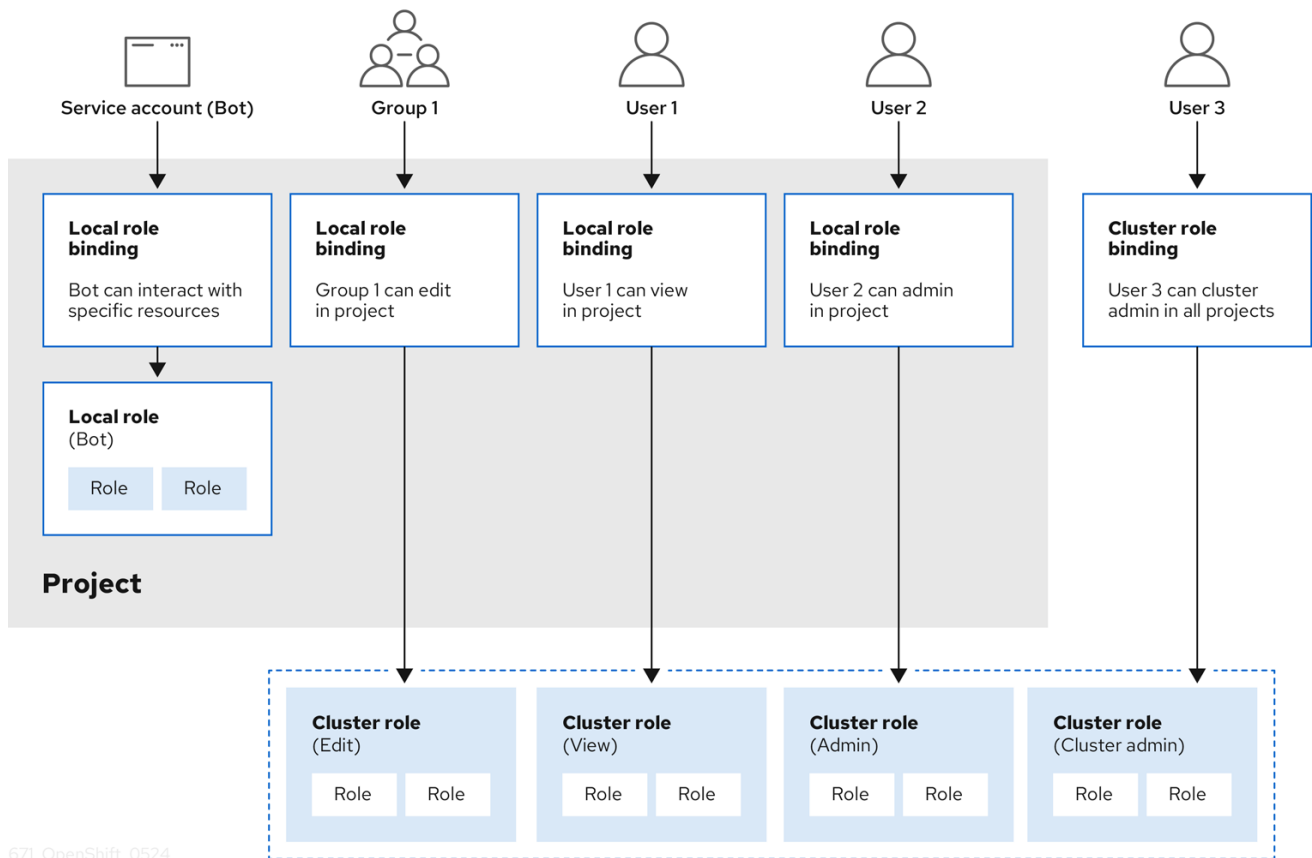
IMPORTANT

It is not recommended to manually modify the default cluster roles. Modifications to these system roles can prevent a cluster from functioning properly.

Default cluster role	Description
admin	A project manager. If used in a local binding, an admin has rights to view any resource in the project and modify any resource in the project except for quota.
basic-user	A user that can get basic information about projects and users.
cluster-admin	A super-user that can perform any action in any project. When bound to a user with a local binding, they have full control over quota and every action on every resource in the project.
cluster-status	A user that can get basic cluster status information.
cluster-reader	A user that can get or view most of the objects but cannot modify them.
edit	A user that can modify most objects in a project but does not have the power to view or modify roles or bindings.
self-provisioner	A user that can create their own projects.
view	A user who cannot make any modifications, but can see most objects in a project. They cannot view or modify roles or bindings.

Be mindful of the difference between local and cluster bindings. For example, if you bind the **cluster-admin** role to a user by using a local role binding, it might appear that this user has the privileges of a cluster administrator. This is not the case. Binding the **cluster-admin** to a user in a project grants super administrator privileges for only that project to the user. That user has the permissions of the cluster role **admin**, plus a few additional permissions like the ability to edit rate limits, for that project. This binding can be confusing via the web console UI, which does not list cluster role bindings that are bound to true cluster administrators. However, it does list local role bindings that you can use to locally bind **cluster-admin**.

The relationships between cluster roles, local roles, cluster role bindings, local role bindings, users, groups and service accounts are illustrated below.



WARNING

The **get pods/exec**, **get pods/***, and **get *** rules grant execution privileges when they are applied to a role. Apply the principle of least privilege and assign only the minimal RBAC rights required for users and agents. For more information, see [RBAC rules allow execution privileges](#).

7.1.2. Evaluating authorization

OpenShift Dedicated evaluates authorization by using:

Identity

The user name and list of groups that the user belongs to.

Action

The action you perform. In most cases, this consists of:

- **Project:** The project you access. A project is a Kubernetes namespace with additional annotations that allows a community of users to organize and manage their content in isolation from other communities.
- **Verb :** The action itself: **get**, **list**, **create**, **update**, **delete**, **deletecollection**, or **watch**.
- **Resource name:** The API endpoint that you access.

Bindings

The full list of bindings, the associations between users or groups with a role.

OpenShift Dedicated evaluates authorization by using the following steps:

1. The identity and the project-scoped action is used to find all bindings that apply to the user or their groups.
2. Bindings are used to locate all the roles that apply.
3. Roles are used to find all the rules that apply.
4. The action is checked against each rule to find a match.
5. If no matching rule is found, the action is then denied by default.

TIP

Remember that users and groups can be associated with, or bound to, multiple roles at the same time.

Project administrators can use the CLI to view local roles and bindings, including a matrix of the verbs and resources each are associated with.



IMPORTANT

The cluster role bound to the project administrator is limited in a project through a local binding. It is not bound cluster-wide like the cluster roles granted to the **cluster-admin** or **system:admin**.

Cluster roles are roles defined at the cluster level but can be bound either at the cluster level or at the project level.

7.1.2.1. Cluster role aggregation

The default admin, edit, view, and cluster-reader cluster roles support [cluster role aggregation](#), where the cluster rules for each role are dynamically updated as new rules are created. This feature is relevant only if you extend the Kubernetes API by creating custom resources.

7.2. PROJECTS AND NAMESPACES

A Kubernetes *namespace* provides a mechanism to scope resources in a cluster. The [Kubernetes documentation](#) has more information on namespaces.

Namespaces provide a unique scope for:

- Named resources to avoid basic naming collisions.
- Delegated management authority to trusted users.
- The ability to limit community resource consumption.

Most objects in the system are scoped by namespace, but some are excepted and have no namespace, including nodes and users.

A *project* is a Kubernetes namespace with additional annotations and is the central vehicle by which

access to resources for regular users is managed. A project allows a community of users to organize and manage their content in isolation from other communities. Users must be given access to projects by administrators, or if allowed to create projects, automatically have access to their own projects.

Projects can have a separate **name**, **displayName**, and **description**.

- The mandatory **name** is a unique identifier for the project and is most visible when using the CLI tools or API. The maximum name length is 63 characters.
- The optional **displayName** is how the project is displayed in the web console (defaults to **name**).
- The optional **description** can be a more detailed description of the project and is also visible in the web console.

Each project scopes its own set of:

Object	Description
Objects	Pods, services, replication controllers, etc.
Policies	Rules for which users can or cannot perform actions on objects.
Constraints	Quotas for each kind of object that can be limited.
Service accounts	Service accounts act automatically with designated access to objects in the project.

Administrators with the **dedicated-admin** role can create projects and delegate administrative rights for the project to any member of the user community. Administrators with the **dedicated-admin** role can also allow developers to create their own projects.

Developers and administrators can interact with projects by using the CLI or the web console.

7.3. DEFAULT PROJECTS

OpenShift Dedicated comes with a number of default projects, and projects starting with **openshift-** are the most essential to users. These projects host master components that run as pods and other infrastructure components. The pods created in these namespaces that have a [critical pod annotation](#) are considered critical, and they have guaranteed admission by kubelet. Pods created for master components in these namespaces are already marked as critical.



IMPORTANT

Do not run workloads in or share access to default projects. Default projects are reserved for running core cluster components.

The following default projects are considered highly privileged: **default**, **kube-public**, **kube-system**, **openshift**, **openshift-infra**, **openshift-node**, and other system-created projects that have the **openshift.io/run-level** label set to **0** or **1**. Functionality that relies on admission plugins, such as pod security admission, security context constraints, cluster resource quotas, and image reference resolution, does not work in highly privileged projects.

7.4. VIEWING CLUSTER ROLES AND BINDINGS

You can use the **oc** CLI to view cluster roles and bindings by using the **oc describe** command.

Prerequisites

- Install the **oc** CLI.
- Obtain permission to view the cluster roles and bindings.

Procedure

1. To view the cluster roles and their associated rule sets:

```
$ oc describe clusterrole.rbac
```

Example output

```
Name:      admin
Labels:    kubernetes.io/bootstrapping=rbac-defaults
Annotations: rbac.authorization.kubernetes.io/autoupdate: true
PolicyRule:
  Resources          Non-Resource URLs  Resource Names  Verbs
-----
.packages.apps.redhat.com          []              []              [* create update
patch delete get list watch]
imagestreams                    []              []              [create delete
deletecollection get list patch update watch create get list watch]
imagestreams.image.openshift.io    []              []              [create delete
deletecollection get list patch update watch create get list watch]
secrets                          []              []              [create delete deletecollection
get list patch update watch get list watch create delete deletecollection patch update]
buildconfigs/webhooks              []              []              [create delete
deletecollection get list patch update watch get list watch]
buildconfigs                       []              []              [create delete
deletecollection get list patch update watch get list watch]
buildlogs                           []              []              [create delete deletecollection
get list patch update watch get list watch]
deploymentconfigs/scale              []              []              [create delete
deletecollection get list patch update watch get list watch]
deploymentconfigs                    []              []              [create delete
deletecollection get list patch update watch get list watch]
```

imagestreamimages	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
imagestreammappings	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
imagestreamtags	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
processedtemplates	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
routes	[]	[]	[create delete deletecollection
get list patch update watch get list watch]			
templateconfigs	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
templateinstances	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
templates	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
deploymentconfigs.apps.openshift.io/scale	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
deploymentconfigs.apps.openshift.io	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
buildconfigs.build.openshift.io/webhooks	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
buildconfigs.build.openshift.io	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
buildlogs.build.openshift.io	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
imagestreamimages.image.openshift.io	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
imagestreammappings.image.openshift.io	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
imagestreamtags.image.openshift.io	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
routes.route.openshift.io	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
processedtemplates.template.openshift.io	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
templateconfigs.template.openshift.io	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
templateinstances.template.openshift.io	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
templates.template.openshift.io	[]	[]	[create delete
deletecollection get list patch update watch get list watch]			
serviceaccounts	[]	[]	[create delete
deletecollection get list patch update watch impersonate create delete deletecollection patch			
update get list watch]			
imagestreams/secrets	[]	[]	[create delete
deletecollection get list patch update watch]			
rolebindings	[]	[]	[create delete
deletecollection get list patch update watch]			
roles	[]	[]	[create delete deletecollection
get list patch update watch]			
rolebindings.authorization.openshift.io	[]	[]	[create delete
deletecollection get list patch update watch]			
roles.authorization.openshift.io	[]	[]	[create delete
deletecollection get list patch update watch]			
imagestreams.image.openshift.io/secrets	[]	[]	[create delete

```

deletecollection get list patch update watch]
  rolebindings.rbac.authorization.k8s.io          []          []          [create delete]
deletecollection get list patch update watch]
  roles.rbac.authorization.k8s.io                 []          []          [create delete]
deletecollection get list patch update watch]
  networkpolicies.extensions                     []          []          [create delete]
deletecollection patch update create delete deletecollection get list patch update watch get
list watch]
  networkpolicies.networking.k8s.io              []          []          [create delete]
deletecollection patch update create delete deletecollection get list patch update watch get
list watch]
  configmaps                                     []          []          [create delete]
deletecollection patch update get list watch]
  endpoints                                     []          []          [create delete]
deletecollection patch update get list watch]
  persistentvolumeclaims                       []          []          [create delete]
deletecollection patch update get list watch]
  pods                                           []          []          [create delete deletecollection
patch update get list watch]
  replicationcontrollers/scale                   []          []          [create delete]
deletecollection patch update get list watch]
  replicationcontrollers                       []          []          [create delete]
deletecollection patch update get list watch]
  services                                     []          []          [create delete deletecollection
patch update get list watch]
  daemonsets.apps                              []          []          [create delete]
deletecollection patch update get list watch]
  deployments.apps/scale                       []          []          [create delete]
deletecollection patch update get list watch]
  deployments.apps                             []          []          [create delete]
deletecollection patch update get list watch]
  replicaset.apps/scale                       []          []          [create delete]
deletecollection patch update get list watch]
  replicaset.apps                              []          []          [create delete]
deletecollection patch update get list watch]
  statefulsets.apps/scale                     []          []          [create delete]
deletecollection patch update get list watch]
  statefulsets.apps                           []          []          [create delete]
deletecollection patch update get list watch]
  horizontalpodautoscalers.autoscaling         []          []          [create delete]
deletecollection patch update get list watch]
  cronjobs.batch                               []          []          [create delete]
deletecollection patch update get list watch]
  jobs.batch                                   []          []          [create delete]
deletecollection patch update get list watch]
  daemonsets.extensions                       []          []          [create delete]
deletecollection patch update get list watch]
  deployments.extensions/scale                 []          []          [create delete]
deletecollection patch update get list watch]
  deployments.extensions                       []          []          [create delete]
deletecollection patch update get list watch]
  ingresses.extensions                         []          []          [create delete]
deletecollection patch update get list watch]
  replicaset.extensions/scale                 []          []          [create delete]
deletecollection patch update get list watch]
  replicaset.extensions                       []          []          [create delete]

```

deletecollection patch update get list watch]				
replicationcontrollers.extensions.scale	[]	[]		[create delete]
deletecollection patch update get list watch]				
poddisruptionbudgets.policy	[]	[]		[create delete]
deletecollection patch update get list watch]				
deployments.apps/rollback	[]	[]		[create delete]
deletecollection patch update]				
deployments.extensions/rollback	[]	[]		[create delete]
deletecollection patch update]				
catalogsources.operators.coreos.com	[]	[]		[create update]
patch delete get list watch]				
clusterserviceversions.operators.coreos.com	[]	[]		[create update]
patch delete get list watch]				
installplans.operators.coreos.com	[]	[]		[create update]
patch delete get list watch]				
packagemanifests.operators.coreos.com	[]	[]		[create update]
patch delete get list watch]				
subscriptions.operators.coreos.com	[]	[]		[create update]
patch delete get list watch]				
buildconfigs/instantiate	[]	[]		[create]
buildconfigs/instantiatebinary	[]	[]		[create]
builds/clone	[]	[]		[create]
deploymentconfigrollbacks	[]	[]		[create]
deploymentconfigs/instantiate	[]	[]		[create]
deploymentconfigs/rollback	[]	[]		[create]
imagestreamimports	[]	[]		[create]
localresourceaccessreviews	[]	[]		[create]
localsubjectaccessreviews	[]	[]		[create]
podsecuritypolicyreviews	[]	[]		[create]
podsecuritypolicyselfsubjectreviews	[]	[]		[create]
podsecuritypolicysubjectreviews	[]	[]		[create]
resourceaccessreviews	[]	[]		[create]
routes/custom-host	[]	[]		[create]
subjectaccessreviews	[]	[]		[create]
subjectrulesreviews	[]	[]		[create]
deploymentconfigrollbacks.apps.openshift.io	[]	[]		[create]
deploymentconfigs.apps.openshift.io/instantiate	[]	[]		[create]
deploymentconfigs.apps.openshift.io/rollback	[]	[]		[create]
localsubjectaccessreviews.authorization.k8s.io	[]	[]		[create]
localresourceaccessreviews.authorization.openshift.io	[]	[]		[create]
localsubjectaccessreviews.authorization.openshift.io	[]	[]		[create]
resourceaccessreviews.authorization.openshift.io	[]	[]		[create]
subjectaccessreviews.authorization.openshift.io	[]	[]		[create]
subjectrulesreviews.authorization.openshift.io	[]	[]		[create]
buildconfigs.build.openshift.io/instantiate	[]	[]		[create]
buildconfigs.build.openshift.io/instantiatebinary	[]	[]		[create]
builds.build.openshift.io/clone	[]	[]		[create]
imagestreamimports.image.openshift.io	[]	[]		[create]
routes.route.openshift.io/custom-host	[]	[]		[create]
podsecuritypolicyreviews.security.openshift.io	[]	[]		[create]
podsecuritypolicyselfsubjectreviews.security.openshift.io	[]	[]		[create]
podsecuritypolicysubjectreviews.security.openshift.io	[]	[]		[create]
jenkins.build.openshift.io	[]	[]		[edit view view admin]
edit view]				
builds	[]	[]		[get create delete]
deletecollection get list patch update watch get list watch]				

builds.build.openshift.io	[]	[]	[get create delete
deletecollection get list patch update watch get list watch]			
projects	[]	[]	[get delete get delete get patch
update]			
projects.project.openshift.io	[]	[]	[get delete get delete
get patch update]			
namespaces	[]	[]	[get get list watch]
Pods/attach	[]	[]	[get list watch create delete
deletecollection patch update]			
Pods/exec	[]	[]	[get list watch create delete
deletecollection patch update]			
Pods/portforward	[]	[]	[get list watch create
delete deletecollection patch update]			
Pods/proxy	[]	[]	[get list watch create delete
deletecollection patch update]			
services/proxy	[]	[]	[get list watch create delete
deletecollection patch update]			
routes/status	[]	[]	[get list watch update]
routes.route.openshift.io/status	[]	[]	[get list watch update]
appliedclusterresourcequotas	[]	[]	[get list watch]
bindings	[]	[]	[get list watch]
builds/log	[]	[]	[get list watch]
deploymentconfigs/log	[]	[]	[get list watch]
deploymentconfigs/status	[]	[]	[get list watch]
events	[]	[]	[get list watch]
imagestreams/status	[]	[]	[get list watch]
limitranges	[]	[]	[get list watch]
namespaces/status	[]	[]	[get list watch]
Pods/log	[]	[]	[get list watch]
Pods/status	[]	[]	[get list watch]
replicationcontrollers/status	[]	[]	[get list watch]
resourcequotas/status	[]	[]	[get list watch]
resourcequotas	[]	[]	[get list watch]
resourcequotausages	[]	[]	[get list watch]
rolebindingrestrictions	[]	[]	[get list watch]
deploymentconfigs.apps.openshift.io/log		[]	[get list watch]
deploymentconfigs.apps.openshift.io/status		[]	[get list watch]
controllerrevisions.apps	[]	[]	[get list watch]
rolebindingrestrictions.authorization.openshift.io		[]	[get list watch]
builds.build.openshift.io/log	[]	[]	[get list watch]
imagestreams.image.openshift.io/status		[]	[get list watch]
appliedclusterresourcequotas.quota.openshift.io		[]	[get list watch]
imagestreams/layers	[]	[]	[get update get]
imagestreams.image.openshift.io/layers		[]	[get update get]
builds/details	[]	[]	[update]
builds.build.openshift.io/details		[]	[update]

Name: basic-user

Labels: <none>

Annotations: openshift.io/description: A user that can get basic information about projects.
rbac.authorization.kubernetes.io/autoupdate: true

PolicyRule:

Resources	Non-Resource URLs	Resource Names	Verbs
-----	-----	-----	-----
selfsubjectrulesreviews	[]	[]	[create]

```

selfsubjectaccessreviews.authorization.k8s.io    []          []          [create]
selfsubjectrulesreviews.authorization.openshift.io []        []          [create]
clusterroles.rbac.authorization.k8s.io          []          []          [get list watch]
clusterroles                                    []          []          [get list]
clusterroles.authorization.openshift.io         []          []          [get list]
storageclasses.storage.k8s.io                  []          []          [get list]
users                                           []          [~]         [get]
users.user.openshift.io                        []          [~]         [get]
projects                                       []          []          [list watch]
projects.project.openshift.io                  []          []          [list watch]
projectrequests                                []          []          [list]
projectrequests.project.openshift.io           []          []          [list]

```

```

Name:      cluster-admin
Labels:    kubernetes.io/bootstrapping=rbac-defaults
Annotations: rbac.authorization.kubernetes.io/autoupdate: true
PolicyRule:

```

Resources Non-Resource URLs Resource Names Verbs

```

-----
*: * [] [] [*]
  [*] [] [*]
...

```

- To view the current set of cluster role bindings, which shows the users and groups that are bound to various roles:

```
$ oc describe clusterrolebinding.rbac
```

Example output

```

Name:      alertmanager-main
Labels:    <none>
Annotations: <none>
Role:
  Kind: ClusterRole
  Name: alertmanager-main
Subjects:
  Kind      Name      Namespace
  ----      -
  ServiceAccount alertmanager-main openshift-monitoring

Name:      basic-users
Labels:    <none>
Annotations: rbac.authorization.kubernetes.io/autoupdate: true
Role:
  Kind: ClusterRole
  Name: basic-user
Subjects:
  Kind Name      Namespace
  ---- ----      -
  Group system:authenticated

```

```

Name:      cloud-credential-operator-rolebinding
Labels:    <none>
Annotations: <none>
Role:
  Kind: ClusterRole
  Name: cloud-credential-operator-role
Subjects:
  Kind      Name      Namespace
  ----      -
  ServiceAccount default openshift-cloud-credential-operator

Name:      cluster-admin
Labels:    kubernetes.io/bootstrapping=rbac-defaults
Annotations: rbac.authorization.kubernetes.io/autoupdate: true
Role:
  Kind: ClusterRole
  Name: cluster-admin
Subjects:
  Kind Name      Namespace
  ---- -
  Group system:masters

Name:      cluster-admins
Labels:    <none>
Annotations: rbac.authorization.kubernetes.io/autoupdate: true
Role:
  Kind: ClusterRole
  Name: cluster-admin
Subjects:
  Kind Name      Namespace
  ---- -
  Group system:cluster-admins
  User system:admin

Name:      cluster-api-manager-rolebinding
Labels:    <none>
Annotations: <none>
Role:
  Kind: ClusterRole
  Name: cluster-api-manager-role
Subjects:
  Kind      Name      Namespace
  ----      -
  ServiceAccount default openshift-machine-api

...

```

7.5. VIEWING LOCAL ROLES AND BINDINGS

You can use the **oc** CLI to view local roles and bindings by using the **oc describe** command.

Prerequisites

- Install the **oc** CLI.
- Obtain permission to view the local roles and bindings:
 - Users with the **admin** default cluster role bound locally can view and manage roles and bindings in that project.

Procedure

1. To view the current set of local role bindings, which show the users and groups that are bound to various roles for the current project:

```
$ oc describe rolebinding.rbac
```

2. To view the local role bindings for a different project, add the **-n** flag to the command:

```
$ oc describe rolebinding.rbac -n joe-project
```

Example output

```
Name:      admin
Labels:    <none>
Annotations: <none>
```

Role:

```
Kind: ClusterRole
```

```
Name: admin
```

Subjects:

```
Kind Name      Namespace
---- ----      -
```

```
-----
```

```
User kube:admin
```

```
Name:      system:deployers
```

```
Labels:    <none>
```

```
Annotations: openshift.io/description:
```

```
Allows deploymentconfigs in this namespace to rollout pods in
this namespace. It is auto-managed by a controller; remove
subjects to disa...
```

Role:

```
Kind: ClusterRole
```

```
Name: system:deployer
```

Subjects:

```
Kind      Name      Namespace
----      -
```

```
-----
```

```
ServiceAccount deployer joe-project
```

```
Name:      system:image-builders
```

```
Labels:    <none>
```

```
Annotations: openshift.io/description:
```

```
Allows builds in this namespace to push images to this
namespace. It is auto-managed by a controller; remove subjects
to disable.
```

Role:

```
Kind: ClusterRole
```

```
Name: system:image-builder
Subjects:
Kind      Name      Namespace
----      -
ServiceAccount builder  joe-project
```

```
Name:      system:image-pullers
Labels:    <none>
Annotations: openshift.io/description:
           Allows all pods in this namespace to pull images from this
           namespace. It is auto-managed by a controller; remove subjects
           to disable.
```

```
Role:
Kind: ClusterRole
Name: system:image-puller
Subjects:
Kind Name      Namespace
---- ----      -
Group system:serviceaccounts:joe-project
```

7.6. ADDING ROLES TO USERS

You can use the **oc adm** administrator CLI to manage the roles and bindings.

Binding, or adding, a role to users or groups gives the user or group the access that is granted by the role. You can add and remove roles to and from users and groups using **oc adm policy** commands.

You can bind any of the default cluster roles to local users or groups in your project.

Procedure

1. Add a role to a user in a specific project:

```
$ oc adm policy add-role-to-user <role> <user> -n <project>
```

For example, you can add the **admin** role to the **alice** user in **joe** project by running:

```
$ oc adm policy add-role-to-user admin alice -n joe
```

TIP

You can alternatively apply the following YAML to add the role to the user:

```
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: admin-0
  namespace: joe
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: admin
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: User
  name: alice
```

2. View the local role bindings and verify the addition in the output:

```
$ oc describe rolebinding.rbac -n <project>
```

For example, to view the local role bindings for the **joe** project:

```
$ oc describe rolebinding.rbac -n joe
```

Example output

```
Name:      admin
Labels:    <none>
Annotations: <none>
Role:
  Kind: ClusterRole
  Name: admin
Subjects:
  Kind Name      Namespace
  ---- ----      -
  User kube:admin
```

```
Name:      admin-0
Labels:    <none>
Annotations: <none>
Role:
  Kind: ClusterRole
  Name: admin
Subjects:
  Kind Name      Namespace
  ---- ----      -
  User alice 1
```

```
Name:      system:deployers
Labels:    <none>
```

```

Annotations: openshift.io/description:
    Allows deploymentconfigs in this namespace to rollout pods in
    this namespace. It is auto-managed by a controller; remove
    subjects to disa...

Role:
  Kind: ClusterRole
  Name: system:deployer
Subjects:
  Kind      Name      Namespace
  ----      -
  ServiceAccount deployer joe

Name:      system:image-builders
Labels:    <none>
Annotations: openshift.io/description:
    Allows builds in this namespace to push images to this
    namespace. It is auto-managed by a controller; remove subjects
    to disable.

Role:
  Kind: ClusterRole
  Name: system:image-builder
Subjects:
  Kind      Name      Namespace
  ----      -
  ServiceAccount builder joe

Name:      system:image-pullers
Labels:    <none>
Annotations: openshift.io/description:
    Allows all pods in this namespace to pull images from this
    namespace. It is auto-managed by a controller; remove subjects
    to disable.

Role:
  Kind: ClusterRole
  Name: system:image-puller
Subjects:
  Kind Name      Namespace
  ---- ----      -
  Group system:serviceaccounts:joe

```

- 1 The **alice** user has been added to the **admins RoleBinding**.

7.7. CREATING A LOCAL ROLE

You can create a local role for a project and then bind it to a user.

Procedure

1. To create a local role for a project, run the following command:

```
$ oc create role <name> --verb=<verb> --resource=<resource> -n <project>
```

In this command, specify:

- **<name>**, the local role's name
- **<verb>**, a comma-separated list of the verbs to apply to the role
- **<resource>**, the resources that the role applies to
- **<project>**, the project name

For example, to create a local role that allows a user to view pods in the **blue** project, run the following command:

```
$ oc create role podview --verb=get --resource=pod -n blue
```

2. To bind the new role to a user, run the following command:

```
$ oc adm policy add-role-to-user podview user2 --role-namespace=blue -n blue
```

7.8. LOCAL ROLE BINDING COMMANDS

When you manage a user or group's associated roles for local role bindings using the following operations, a project may be specified with the **-n** flag. If it is not specified, then the current project is used.

You can use the following commands for local RBAC management.

Table 7.1. Local role binding operations

Command	Description
\$ oc adm policy who-can <verb> <resource>	Indicates which users can perform an action on a resource.
\$ oc adm policy add-role-to-user <role> <username>	Binds a specified role to specified users in the current project.
\$ oc adm policy remove-role-from-user <role> <username>	Removes a given role from specified users in the current project.
\$ oc adm policy remove-user <username>	Removes specified users and all of their roles in the current project.
\$ oc adm policy add-role-to-group <role> <groupname>	Binds a given role to specified groups in the current project.
\$ oc adm policy remove-role-from-group <role> <groupname>	Removes a given role from specified groups in the current project.
\$ oc adm policy remove-group <groupname>	Removes specified groups and all of their roles in the current project.

7.9. CLUSTER ROLE BINDING COMMANDS

You can also manage cluster role bindings using the following operations. The **-n** flag is not used for these operations because cluster role bindings use non-namespaced resources.

Table 7.2. Cluster role binding operations

Command	Description
\$ oc adm policy add-cluster-role-to-user <role> <username>	Binds a given role to specified users for all projects in the cluster.
\$ oc adm policy remove-cluster-role-from-user <role> <username>	Removes a given role from specified users for all projects in the cluster.
\$ oc adm policy add-cluster-role-to-group <role> <groupname>	Binds a given role to specified groups for all projects in the cluster.
\$ oc adm policy remove-cluster-role-from-group <role> <groupname>	Removes a given role from specified groups for all projects in the cluster.

7.10. GRANTING ADMINISTRATOR PRIVILEGES TO A USER

After you have configured an identity provider for your cluster and added a user to the identity provider, you can grant **dedicated-admin** cluster privileges to the user.

Prerequisites

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


Procedure

1. Navigate to [OpenShift Cluster Manager](#) and select your cluster.
2. Click the **Access control** tab.
3. In the **Cluster Roles and Access** tab, click **Add user**.
4. Enter the user ID of an identity provider user.
5. Click **Add user** to grant **dedicated-admin** cluster privileges to the user.

Verification

- After granting the privileges, the user is listed as part of the **dedicated-admins** group under **Access control** → **Cluster Roles and Access** on the OpenShift Cluster Manager page for your cluster.

7.11. CLUSTER ROLE BINDINGS FOR UNAUTHENTICATED GROUPS



NOTE

Before OpenShift Dedicated 4.17, unauthenticated groups were allowed access to some cluster roles. Clusters updated from versions before OpenShift Dedicated 4.17 retain this access for unauthenticated groups.

For security reasons OpenShift Dedicated 4 does not allow unauthenticated groups to have default access to cluster roles.

There are use cases where it might be necessary to add **system:unauthenticated** to a cluster role.

Cluster administrators can add unauthenticated users to the following cluster roles:

- **system:scope-impersonation**
- **system:webhook**
- **system:oauth-token-deleter**
- **self-access-reviewer**



IMPORTANT

Always verify compliance with your organization's security standards when modifying unauthenticated access.

CHAPTER 8. UNDERSTANDING AND CREATING SERVICE ACCOUNTS

8.1. SERVICE ACCOUNTS OVERVIEW

A service account is an OpenShift Dedicated account that allows a component to directly access the API. Service accounts are API objects that exist within each project. Service accounts provide a flexible way to control API access without sharing a regular user's credentials.

When you use the OpenShift Dedicated CLI or web console, your API token authenticates you to the API. You can associate a component with a service account so that they can access the API without using a regular user's credentials.

Each service account's user name is derived from its project and name:

```
system:serviceaccount:<project>:<name>
```

Every service account is also a member of two groups:

Group	Description
system:serviceaccounts	Includes all service accounts in the system.
system:serviceaccounts:<project>	Includes all service accounts in the specified project.

8.1.1. Automatically generated image pull secrets

By default, OpenShift Dedicated creates an image pull secret for each service account.



NOTE

Prior to OpenShift Dedicated 4.16, a long-lived service account API token secret was also generated for each service account that was created. Starting with OpenShift Dedicated 4.16, this service account API token secret is no longer created.

After upgrading to 4, any existing long-lived service account API token secrets are not deleted and will continue to function. For information about detecting long-lived API tokens that are in use in your cluster or deleting them if they are not needed, see the Red Hat Knowledgebase article [Long-lived service account API tokens in OpenShift Container Platform](#).

This image pull secret is necessary to integrate the OpenShift image registry into the cluster's user authentication and authorization system.

However, if you do not enable the **ImageRegistry** capability or if you disable the integrated OpenShift image registry in the Cluster Image Registry Operator's configuration, an image pull secret is not generated for each service account.

When the integrated OpenShift image registry is disabled on a cluster that previously had it enabled, the previously generated image pull secrets are deleted automatically.

8.2. CREATING SERVICE ACCOUNTS

You can create a service account in a project and grant it permissions by binding it to a role.

Procedure

- Optional: To view the service accounts in the current project:

```
$ oc get sa
```

Example output

```
NAME      SECRETS  AGE
builder   1        2d
default   1        2d
deployer  1        2d
```

- To create a new service account in the current project:

```
$ oc create sa <service_account_name> 1
```

- 1** To create a service account in a different project, specify **-n <project_name>**.

Example output

```
serviceaccount "robot" created
```

TIP

You can alternatively apply the following YAML to create the service account:

```
apiVersion: v1
kind: ServiceAccount
metadata:
  name: <service_account_name>
  namespace: <current_project>
```

- Optional: View the secrets for the service account:

```
$ oc describe sa robot
```

Example output

```
Name:          robot
Namespace:     project1
Labels:        <none>
Annotations:   openshift.io/internal-registry-pull-secret-ref: robot-dockercfg-qzbhb
Image pull secrets: robot-dockercfg-qzbhb
Mountable secrets: robot-dockercfg-qzbhb
Tokens:        <none>
Events:        <none>
```

■

8.3. GRANTING ROLES TO SERVICE ACCOUNTS

You can grant roles to service accounts in the same way that you grant roles to a regular user account.

Procedure

1. You can modify the service accounts for the current project. For example, to add the **view** role to the **robot** service account in the **top-secret** project:

```
$ oc policy add-role-to-user view system:serviceaccount:top-secret:robot
```

TIP

You can alternatively apply the following YAML to add the role:

```
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: view
  namespace: top-secret
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: view
subjects:
- kind: ServiceAccount
  name: robot
  namespace: top-secret
```

2. You can also grant access to a specific service account in a project. For example, from the project to which the service account belongs, use the **-z** flag and specify the **<service_account_name>**

```
$ oc policy add-role-to-user <role_name> -z <service_account_name>
```



IMPORTANT

If you want to grant access to a specific service account in a project, use the **-z** flag. Using this flag helps prevent typos and ensures that access is granted to only the specified service account.

TIP

You can alternatively apply the following YAML to add the role:

```

apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: <rolebinding_name>
  namespace: <current_project_name>
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: <role_name>
subjects:
- kind: ServiceAccount
  name: <service_account_name>
  namespace: <current_project_name>

```

- To modify a different namespace, you can use the **-n** option to indicate the project namespace it applies to, as shown in the following examples.
 - For example, to allow all service accounts in all projects to view resources in the **my-project** project:

```
$ oc policy add-role-to-group view system:serviceaccounts -n my-project
```

TIP

You can alternatively apply the following YAML to add the role:

```

apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: view
  namespace: my-project
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: view
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: Group
  name: system:serviceaccounts

```

- To allow all service accounts in the **managers** project to edit resources in the **my-project** project:

```
$ oc policy add-role-to-group edit system:serviceaccounts:managers -n my-project
```

TIP

You can alternatively apply the following YAML to add the role:

```
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: edit
  namespace: my-project
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: edit
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: Group
  name: system:serviceaccounts:managers
```

CHAPTER 9. USING SERVICE ACCOUNTS IN APPLICATIONS

9.1. SERVICE ACCOUNTS OVERVIEW

A service account is an OpenShift Dedicated account that allows a component to directly access the API. Service accounts are API objects that exist within each project. Service accounts provide a flexible way to control API access without sharing a regular user's credentials.

When you use the OpenShift Dedicated CLI or web console, your API token authenticates you to the API. You can associate a component with a service account so that they can access the API without using a regular user's credentials.

Each service account's user name is derived from its project and name:

```
system:serviceaccount:<project>:<name>
```

Every service account is also a member of two groups:

Group	Description
system:serviceaccounts	Includes all service accounts in the system.
system:serviceaccounts:<project>	Includes all service accounts in the specified project.

9.2. DEFAULT SERVICE ACCOUNTS

Your OpenShift Dedicated cluster contains default service accounts for cluster management and generates more service accounts for each project.



9.2.1. Default cluster service accounts

Several infrastructure controllers run using service account credentials. The following service accounts are created in the OpenShift Dedicated infrastructure project (**openshift-infra**) at server start, and given the following roles cluster-wide:

Service account	Description
replication-controller	Assigned the system:replication-controller role
deployment-controller	Assigned the system:deployment-controller role
build-controller	Assigned the system:build-controller role. Additionally, the build-controller service account is included in the privileged security context constraint to create privileged build pods.

9.2.2. Default project service accounts and roles

Three service accounts are automatically created in each project:

Service account	Usage
builder	<p>Used by build pods. It is given the system:image-builder role, which allows pushing images to any imagestream in the project using the internal Docker registry.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>The builder service account is not created if the Build cluster capability is not enabled.</p> </div> </div>
deployer	<p>Used by deployment pods and given the system:deployer role, which allows viewing and modifying replication controllers and pods in the project.</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>The deployer service account is not created if the DeploymentConfig cluster capability is not enabled.</p> </div> </div>
default	Used to run all other pods unless they specify a different service account.

All service accounts in a project are given the **system:image-puller** role, which allows pulling images from any image stream in the project using the internal container image registry.

9.2.3. Automatically generated image pull secrets

By default, OpenShift Dedicated creates an image pull secret for each service account.



NOTE

Prior to OpenShift Dedicated 4.16, a long-lived service account API token secret was also generated for each service account that was created. Starting with OpenShift Dedicated 4.16, this service account API token secret is no longer created.

After upgrading to 4, any existing long-lived service account API token secrets are not deleted and will continue to function. For information about detecting long-lived API tokens that are in use in your cluster or deleting them if they are not needed, see the Red Hat Knowledgebase article [Long-lived service account API tokens in OpenShift Container Platform](#).

This image pull secret is necessary to integrate the OpenShift image registry into the cluster's user authentication and authorization system.

However, if you do not enable the **ImageRegistry** capability or if you disable the integrated OpenShift image registry in the Cluster Image Registry Operator's configuration, an image pull secret is not generated for each service account.

When the integrated OpenShift image registry is disabled on a cluster that previously had it enabled, the previously generated image pull secrets are deleted automatically.

9.3. CREATING SERVICE ACCOUNTS

You can create a service account in a project and grant it permissions by binding it to a role.

Procedure

1. Optional: To view the service accounts in the current project:

```
$ oc get sa
```

Example output

```
NAME      SECRETS  AGE
builder   1        2d
default   1        2d
deployer  1        2d
```

2. To create a new service account in the current project:

```
$ oc create sa <service_account_name> 1
```

- 1** To create a service account in a different project, specify **-n <project_name>**.

Example output

```
serviceaccount "robot" created
```

TIP

You can alternatively apply the following YAML to create the service account:

```
apiVersion: v1
kind: ServiceAccount
metadata:
  name: <service_account_name>
  namespace: <current_project>
```

3. Optional: View the secrets for the service account:

```
$ oc describe sa robot
```

Example output

```
Name:          robot
Namespace:     project1
Labels:        <none>
Annotations:   openshift.io/internal-registry-pull-secret-ref: robot-dockercfg-qzbhb
Image pull secrets: robot-dockercfg-qzbhb
Mountable secrets: robot-dockercfg-qzbhb
Tokens:        <none>
Events:        <none>
```

■

CHAPTER 10. USING A SERVICE ACCOUNT AS AN OAUTH CLIENT

10.1. SERVICE ACCOUNTS AS OAUTH CLIENTS

You can use a service account as a constrained form of OAuth client. Service accounts can request only a subset of scopes that allow access to some basic user information and role-based power inside of the service account's own namespace:

- **user:info**
- **user:check-access**
- **role:<any_role>:<service_account_namespace>**
- **role:<any_role>:<service_account_namespace>:!**

When using a service account as an OAuth client:

- **client_id** is **system:serviceaccount:<service_account_namespace>:<service_account_name>**.
- **client_secret** can be any of the API tokens for that service account. For example:

```
$ oc sa get-token <service_account_name>
```

- To get **WWW-Authenticate** challenges, set an **serviceaccounts.openshift.io/oauth-want-challenges** annotation on the service account to **true**.
- **redirect_uri** must match an annotation on the service account.

10.1.1. Redirect URIs for service accounts as OAuth clients

Annotation keys must have the prefix **serviceaccounts.openshift.io/oauth-redirecturi.** or **serviceaccounts.openshift.io/oauth-redirectreference.** such as:

```
serviceaccounts.openshift.io/oauth-redirecturi.<name>
```

In its simplest form, the annotation can be used to directly specify valid redirect URIs. For example:

```
"serviceaccounts.openshift.io/oauth-redirecturi.first": "https://example.com"
"serviceaccounts.openshift.io/oauth-redirecturi.second": "https://other.com"
```

The **first** and **second** postfixes in the above example are used to separate the two valid redirect URIs.

In more complex configurations, static redirect URIs may not be enough. For example, perhaps you want all Ingresses for a route to be considered valid. This is where dynamic redirect URIs via the **serviceaccounts.openshift.io/oauth-redirectreference.** prefix come into play.

For example:

```
"serviceaccounts.openshift.io/oauth-redirectreference.first": "
{"kind":"OAuthRedirectReference","apiVersion":"v1","reference":
{"kind":"Route","name":"jenkins\"}}
```

Since the value for this annotation contains serialized JSON data, it is easier to see in an expanded format:

```
{
  "kind": "OAuthRedirectReference",
  "apiVersion": "v1",
  "reference": {
    "kind": "Route",
    "name": "jenkins"
  }
}
```

Now you can see that an **OAuthRedirectReference** allows us to reference the route named **jenkins**. Thus, all Ingresses for that route will now be considered valid. The full specification for an **OAuthRedirectReference** is:

```
{
  "kind": "OAuthRedirectReference",
  "apiVersion": "v1",
  "reference": {
    "kind": ..., 1
    "name": ..., 2
    "group": ... 3
  }
}
```

- 1** **kind** refers to the type of the object being referenced. Currently, only **route** is supported.
- 2** **name** refers to the name of the object. The object must be in the same namespace as the service account.
- 3** **group** refers to the group of the object. Leave this blank, as the group for a route is the empty string.

Both annotation prefixes can be combined to override the data provided by the reference object. For example:

```
"serviceaccounts.openshift.io/oauth-redirecturi.first": "custompath"
"serviceaccounts.openshift.io/oauth-redirectreference.first": "
{"kind":"OAuthRedirectReference","apiVersion":"v1","reference":
{"kind":"Route","name":"jenkins\"}}
```

The **first** postfix is used to tie the annotations together. Assuming that the **jenkins** route had an Ingress of **https://example.com**, now **https://example.com/custompath** is considered valid, but **https://example.com** is not. The format for partially supplying override data is as follows:

Type	Syntax
Scheme	"https://"
Hostname	"//website.com"
Port	"//:8000"
Path	"examplepath"

**NOTE**

Specifying a hostname override will replace the hostname data from the referenced object, which is not likely to be desired behavior.

Any combination of the above syntax can be combined using the following format:

<scheme>://<hostname><:port>/<path>

The same object can be referenced more than once for more flexibility:

```
"serviceaccounts.openshift.io/oauth-redirecturi.first": "custompath"
"serviceaccounts.openshift.io/oauth-redirectreference.first": "
{"kind":"OAuthRedirectReference","apiVersion":"v1","reference":
{"kind":"Route","name":"jenkins"}}"
"serviceaccounts.openshift.io/oauth-redirecturi.second": "//:8000"
"serviceaccounts.openshift.io/oauth-redirectreference.second": "
{"kind":"OAuthRedirectReference","apiVersion":"v1","reference":
{"kind":"Route","name":"jenkins"}}"
```

Assuming that the route named **jenkins** has an Ingress of **https://example.com**, then both **https://example.com:8000** and **https://example.com/custompath** are considered valid.

Static and dynamic annotations can be used at the same time to achieve the desired behavior:

```
"serviceaccounts.openshift.io/oauth-redirectreference.first": "
{"kind":"OAuthRedirectReference","apiVersion":"v1","reference":
{"kind":"Route","name":"jenkins"}}"
"serviceaccounts.openshift.io/oauth-redirecturi.second": "https://other.com"
```

CHAPTER 11. SCOPING TOKENS

11.1. ABOUT SCOPING TOKENS

You can create scoped tokens to delegate some of your permissions to another user or service account. For example, a project administrator might want to delegate the power to create pods.

A scoped token is a token that identifies as a given user but is limited to certain actions by its scope. Only a user with the **dedicated-admin** role can create scoped tokens.

Scopes are evaluated by converting the set of scopes for a token into a set of **PolicyRules**. Then, the request is matched against those rules. The request attributes must match at least one of the scope rules to be passed to the "normal" authorizer for further authorization checks.

11.1.1. User scopes

User scopes are focused on getting information about a given user. They are intent-based, so the rules are automatically created for you:

- **user:full** - Allows full read/write access to the API with all of the user's permissions.
- **user:info** - Allows read-only access to information about the user, such as name and groups.
- **user:check-access** - Allows access to **self-localsubjectaccessreviews** and **self-subjectaccessreviews**. These are the variables where you pass an empty user and groups in your request object.
- **user:list-projects** - Allows read-only access to list the projects the user has access to.

11.1.2. Role scope

The role scope allows you to have the same level of access as a given role filtered by namespace.

- **role:<cluster-role name>:<namespace or * for all>** - Limits the scope to the rules specified by the cluster-role, but only in the specified namespace .



NOTE

Caveat: This prevents escalating access. Even if the role allows access to resources like secrets, rolebindings, and roles, this scope will deny access to those resources. This helps prevent unexpected escalations. Many people do not think of a role like **edit** as being an escalating role, but with access to a secret it is.


- **role:<cluster-role name>:<namespace or * for all>:!** - This is similar to the example above, except that including the bang causes this scope to allow escalating access.

11.2. ADDING UNAUTHENTICATED GROUPS TO CLUSTER ROLES

As a cluster administrator, you can add unauthenticated users to the following cluster roles in OpenShift Dedicated by creating a cluster role binding. Unauthenticated users do not have access to non-public cluster roles. This should only be done in specific use cases when necessary.

You can add unauthenticated users to the following cluster roles:

- **system:scope-impersonation**
- **system:webhook**
- **system:oauth-token-deleter**
- **self-access-reviewer**



IMPORTANT

Always verify compliance with your organization's security standards when modifying unauthenticated access.

Prerequisites

- You have access to the cluster as a user with the **cluster-admin** role.
- You have installed the OpenShift CLI (**oc**).

Procedure

1. Create a YAML file named **add-<cluster_role>-unauth.yaml** and add the following content:

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  annotations:
    rbac.authorization.kubernetes.io/autoupdate: "true"
  name: <cluster_role>access-unauthenticated
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: <cluster_role>
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: Group
  name: system:unauthenticated
```

2. Apply the configuration by running the following command:

```
$ oc apply -f add-<cluster_role>.yaml
```

CHAPTER 12. USING BOUND SERVICE ACCOUNT TOKENS

You can use bound service account tokens, which improves the ability to integrate with cloud provider identity access management (IAM) services, such as OpenShift Dedicated on AWS IAM or Google Cloud IAM.

12.1. ABOUT BOUND SERVICE ACCOUNT TOKENS

You can use bound service account tokens to limit the scope of permissions for a given service account token. These tokens are audience and time-bound. This facilitates the authentication of a service account to an IAM role and the generation of temporary credentials mounted to a pod. You can request bound service account tokens by using volume projection and the TokenRequest API.

12.2. CONFIGURING BOUND SERVICE ACCOUNT TOKENS USING VOLUME PROJECTION

You can configure pods to request bound service account tokens by using volume projection.

Prerequisites

- You have access to the cluster as a user with the **dedicated-admin** role.
- You have created a service account. This procedure assumes that the service account is named **build-robot**.

Procedure

1. Configure a pod to use a bound service account token by using volume projection.
 - a. Create a file called **pod-projected-svc-token.yaml** with the following contents:

```
apiVersion: v1
kind: Pod
metadata:
  name: nginx
spec:
  securityContext:
    runAsNonRoot: true 1
    seccompProfile:
      type: RuntimeDefault 2
  containers:
  - image: nginx
    name: nginx
    volumeMounts:
    - mountPath: /var/run/secrets/tokens
      name: vault-token
      securityContext:
        allowPrivilegeEscalation: false
        capabilities:
          drop: [ALL]
    serviceAccountName: build-robot 3
  volumes:
  - name: vault-token
```

```

projected:
  sources:
  - serviceAccountToken:
      path: vault-token 4
      expirationSeconds: 7200 5
      audience: vault 6

```

- 1** Prevents containers from running as root to minimize compromise risks.
- 2** Sets the default seccomp profile, limiting to essential system calls, to reduce risks.
- 3** A reference to an existing service account.
- 4** The path relative to the mount point of the file to project the token into.
- 5** Optionally set the expiration of the service account token, in seconds. The default value is 3600 seconds (1 hour), and this value must be at least 600 seconds (10 minutes). The kubelet starts trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours.
- 6** Optionally set the intended audience of the token. The recipient of a token should verify that the recipient identity matches the audience claim of the token, and should otherwise reject the token. The audience defaults to the identifier of the API server.



NOTE

In order to prevent unexpected failure, OpenShift Dedicated overrides the **expirationSeconds** value to be one year from the initial token generation with the **--service-account-extend-token-expiration** default of **true**. You cannot change this setting.

- b. Create the pod:

```
$ oc create -f pod-projected-svc-token.yaml
```

The kubelet requests and stores the token on behalf of the pod, makes the token available to the pod at a configurable file path, and refreshes the token as it approaches expiration.

2. The application that uses the bound token must handle reloading the token when it rotates. The kubelet rotates the token if it is older than 80 percent of its time to live, or if the token is older than 24 hours.

12.3. CREATING BOUND SERVICE ACCOUNT TOKENS OUTSIDE THE POD

Prerequisites


- You have created a service account. This procedure assumes that the service account is named **build-robot**.

Procedure

CHAPTER 13. MANAGING SECURITY CONTEXT CONSTRAINTS

In OpenShift Dedicated, you can use security context constraints (SCCs) to control permissions for the pods in your cluster.

Default SCCs are created during installation and when you install some Operators or other components. As a cluster administrator, you can also create your own SCCs by using the OpenShift CLI (**oc**).



IMPORTANT

Do not modify the default SCCs. Customizing the default SCCs can lead to issues when some of the platform pods deploy or OpenShift Dedicated is upgraded. Additionally, the default SCC values are reset to the defaults during some cluster upgrades, which discards all customizations to those SCCs.

Instead of modifying the default SCCs, create and modify your own SCCs as needed. For detailed steps, see [Creating security context constraints](#).



NOTE

In OpenShift Dedicated deployments, you can create your own SCCs only for clusters that use the Customer Cloud Subscription (CCS) model. You cannot create SCCs for OpenShift Dedicated clusters that use a Red Hat cloud account, because SCC resource creation requires **cluster-admin** privileges.

13.1. ABOUT SECURITY CONTEXT CONSTRAINTS

Similar to the way that RBAC resources control user access, administrators can use security context constraints (SCCs) to control permissions for pods. These permissions determine the actions that a pod can perform and what resources it can access. You can use SCCs to define a set of conditions that a pod must run with to be accepted into the system.

Security context constraints allow an administrator to control:

- Whether a pod can run privileged containers with the **allowPrivilegedContainer** flag
- Whether a pod is constrained with the **allowPrivilegeEscalation** flag
- The capabilities that a container can request
- The use of host directories as volumes
- The SELinux context of the container
- The container user ID
- The use of host namespaces and networking
- The allocation of an **FSGroup** that owns the pod volumes
- The configuration of allowable supplemental groups
- Whether a container requires write access to its root file system
- The usage of volume types

- The configuration of allowable **seccomp** profiles



IMPORTANT

Do not set the **openshift.io/run-level** label on any namespaces in OpenShift Dedicated. This label is for use by internal OpenShift Dedicated components to manage the startup of major API groups, such as the Kubernetes API server and OpenShift API server. If the **openshift.io/run-level** label is set, no SCCs are applied to pods in that namespace, causing any workloads running in that namespace to be highly privileged.

13.1.1. Default security context constraints

The cluster contains several default security context constraints (SCCs) as described in the table below. Additional SCCs might be installed when you install Operators or other components to OpenShift Dedicated.



IMPORTANT


Do not modify the default SCCs. Customizing the default SCCs can lead to issues when some of the platform pods deploy or OpenShift Dedicated is upgraded. Additionally, the default SCC values are reset to the defaults during some cluster upgrades, which discards all customizations to those SCCs.

Instead of modifying the default SCCs, create and modify your own SCCs as needed. For detailed steps, see *Creating security context constraints*.

Table 13.1. Default security context constraints

Security context constraint	Description
anyuid	Provides all features of the restricted SCC, but allows users to run with any UID and any GID.
nested-container	<p>Like the restricted-v2 SCC, but with the following differences:</p> <ul style="list-style-type: none"> • seLinuxContext is set to MustRunAs and seLinuxOptions.type is container_engine_t. • runAsUser is set to MustRunAsRange. • requiredDropCapabilities is set to null. • userNamespaceLevel is set to RequirePodLevel, which forces pods to be in a Linux user namespace (hostUsers: false). <p>This SCC allows a user to run a container engine inside of an OpenShift Dedicated pod.</p>
nonroot	Provides all features of the restricted SCC, but allows users to run with any non-root UID. The user must specify the UID or it must be specified in the manifest of the container runtime.

Security context constraint	Description
nonroot-v2	<p>Like the nonroot SCC, but with the following differences:</p> <ul style="list-style-type: none"> ● ALL capabilities are dropped from containers. ● The NET_BIND_SERVICE capability can be added explicitly. ● seccompProfile is set to runtime/default by default. ● allowPrivilegeEscalation must be unset or set to false in security contexts.
restricted	<p>Denies access to all host features and requires pods to be run with a UID, and SELinux context that are allocated to the namespace.</p> <p>The restricted SCC:</p> <ul style="list-style-type: none"> ● Ensures that pods cannot run as privileged ● Ensures that pods cannot mount host directory volumes ● Requires that a pod is run as a user in a pre-allocated range of UIDs ● Requires that a pod is run with a pre-allocated MCS label ● Requires that a pod is run with a preallocated FSGroup ● Allows pods to use any supplemental group <p>In clusters that were upgraded from OpenShift Dedicated 4.10 or earlier, this SCC is available for use by any authenticated user. The restricted SCC is no longer available to users of new OpenShift Dedicated 4.11 or later installations, unless the access is explicitly granted.</p>
restricted-v2	<p>Like the restricted SCC, but with the following differences:</p> <ul style="list-style-type: none"> ● ALL capabilities are dropped from containers. ● The NET_BIND_SERVICE capability can be added explicitly. ● seccompProfile is set to runtime/default by default. ● allowPrivilegeEscalation must be unset or set to false in security contexts. <p>This SCC is used by default for authenticated users.</p>

Security context constraint	Description
restricted-v3	<p>Like the restricted-v2 SCC, but with the following differences:</p> <ul style="list-style-type: none"> • UserNamespaceLevel is set to RequirePodLevel, which forces pods to be in a Linux user namespace (hostUsers: false). <p>This is the most restrictive SCC provided by a new installation and will be used by default for authenticated users.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>The restricted-v3 SCC is the most restrictive of the SCCs that is included by default with the system. However, you can create a custom SCC that is even more restrictive. For example, you can create an SCC that restricts readOnlyRootFilesystem to true.</p> </div> </div>

13.1.2. Security context constraints settings

Security context constraints (SCCs) are composed of settings and strategies that control the security features a pod has access to. These settings fall into three categories:

Category	Description
Controlled by a boolean	Fields of this type default to the most restrictive value. For example, AllowPrivilegedContainer is always set to false if unspecified.
Controlled by an allowable set	Fields of this type are checked against the set to ensure their value is allowed.
Controlled by a strategy	Items that have a strategy to generate a value provide: <ul style="list-style-type: none"> • A mechanism to generate the value, and • A mechanism to ensure that a specified value falls into the set of allowable values.

CRI-O has the following default list of capabilities that are allowed for each container of a pod:

- **CHOWN**
- **DAC_OVERRIDE**

- **FSETID**
- **FOWNER**
- **SETGID**
- **SETUID**
- **SETPCAP**
- **NET_BIND_SERVICE**
- **KILL**

The containers use the capabilities from this default list, but pod manifest authors can alter the list by requesting additional capabilities or removing some of the default behaviors. Use the **allowedCapabilities**, **defaultAddCapabilities**, and **requiredDropCapabilities** parameters to control such requests from the pods. With these parameters you can specify which capabilities can be requested, which ones must be added to each container, and which ones must be forbidden, or dropped, from each container.



NOTE

You can drop all capabilities from containers by setting the **requiredDropCapabilities** parameter to **ALL**. This is what the **restricted-v2** SCC does.

13.1.3. Security context constraints strategies

RunAsUser

- **MustRunAs** - Requires a **runAsUser** to be configured. Uses the configured **runAsUser** as the default. Validates against the configured **runAsUser**.

Example MustRunAs snippet

```
...
runAsUser:
  type: MustRunAs
  uid: <id>
...
```

- **MustRunAsRange** - Requires minimum and maximum values to be defined if not using pre-allocated values. Uses the minimum as the default. Validates against the entire allowable range.

Example MustRunAsRange snippet

```
...
runAsUser:
  type: MustRunAsRange
  uidRangeMax: <maxvalue>
  uidRangeMin: <minvalue>
...
```

- **MustRunAsNonRoot** - Requires that the pod be submitted with a non-zero **runAsUser** or have the **USER** directive defined in the image. No default provided.

Example **MustRunAsNonRoot** snippet

```
...
runAsUser:
  type: MustRunAsNonRoot
...
```

- **RunAsAny** - No default provided. Allows any **runAsUser** to be specified.

Example **RunAsAny** snippet

```
...
runAsUser:
  type: RunAsAny
...
```

SELinuxContext

- **MustRunAs** - Requires **seLinuxOptions** to be configured if not using pre-allocated values. Uses **seLinuxOptions** as the default. Validates against **seLinuxOptions**.
- **RunAsAny** - No default provided. Allows any **seLinuxOptions** to be specified.

SupplementalGroups

- **MustRunAs** - Requires at least one range to be specified if not using pre-allocated values. Uses the minimum value of the first range as the default. Validates against all ranges.
- **RunAsAny** - No default provided. Allows any **supplementalGroups** to be specified.

FSGroup

- **MustRunAs** - Requires at least one range to be specified if not using pre-allocated values. Uses the minimum value of the first range as the default. Validates against the first ID in the first range.
- **RunAsAny** - No default provided. Allows any **fsGroup** ID to be specified.

13.1.4. Controlling volumes for CCS clusters

The usage of specific volume types for OpenShift Dedicated with Customer Cloud Subscription (CCS) clusters can be controlled by setting the **volumes** field of the SCC.

The allowable values of this field correspond to the volume sources that are defined when creating a volume:

- [awsElasticBlockStore](#)
- [azureDisk](#)

- [azureFile](#)
- [cephFS](#)
- [cinder](#)
- [configMap](#)
- [csi](#)
- [downwardAPI](#)
- [emptyDir](#)
- [fc](#)
- [flexVolume](#)
- [flocker](#)
- [gcePersistentDisk](#)
- [ephemeral](#)
- [gitRepo](#)
- [glusterfs](#)
- [hostPath](#)
- [iscsi](#)
- [nfs](#)
- [persistentVolumeClaim](#)
- [photonPersistentDisk](#)
- [portworxVolume](#)
- [projected](#)
- [quobyte](#)
- [rbd](#)
- [scaleIO](#)
- [secret](#)
- [storageos](#)
- [vsphereVolume](#)
- * (A special value to allow the use of all volume types.)

- **none** (A special value to disallow the use of all volumes types. Exists only for backwards compatibility.)

The recommended minimum set of allowed volumes for new SCCs are **configMap**, **downwardAPI**, **emptyDir**, **persistentVolumeClaim**, **secret**, and **projected**.



NOTE

This list of allowable volume types is not exhaustive because new types are added with each release of OpenShift Dedicated.



NOTE

For backwards compatibility, the usage of **allowHostDirVolumePlugin** overrides settings in the **volumes** field. For example, if **allowHostDirVolumePlugin** is set to false but allowed in the **volumes** field, then the **hostPath** value will be removed from **volumes**.

13.1.5. Admission control

Admission control with SCCs allows for control over the creation of resources based on the capabilities granted to a user.

In terms of the SCCs, this means that an admission controller can inspect the user information made available in the context to retrieve an appropriate set of SCCs. Doing so ensures the pod is authorized to make requests about its operating environment or to generate a set of constraints to apply to the pod.

The set of SCCs that admission uses to authorize a pod are determined by the user identity and groups that the user belongs to. Additionally, if the pod specifies a service account, the set of allowable SCCs includes any constraints accessible to the service account.



NOTE

When you create a workload resource, such as deployment, only the service account is used to find the SCCs and admit the pods when they are created.

Admission uses the following approach to create the final security context for the pod:

1. Retrieve all SCCs available for use.
2. Generate field values for security context settings that were not specified on the request.
3. Validate the final settings against the available constraints.

If a matching set of constraints is found, then the pod is accepted. If the request cannot be matched to an SCC, the pod is rejected.

A pod must validate every field against the SCC. The following are examples for just two of the fields that must be validated:



NOTE

These examples are in the context of a strategy using the pre-allocated values.

An FSGroup SCC strategy of MustRunAs

If the pod defines a **fsGroup** ID, then that ID must equal the default **fsGroup** ID. Otherwise, the pod is not validated by that SCC and the next SCC is evaluated.

If the **SecurityContextConstraints.fsGroup** field has value **RunAsAny** and the pod specification omits the **Pod.spec.securityContext.fsGroup**, then this field is considered valid. Note that it is possible that during validation, other SCC settings will reject other pod fields and thus cause the pod to fail.

A SupplementalGroups SCC strategy of MustRunAs

If the pod specification defines one or more **supplementalGroups** IDs, then the pod's IDs must equal one of the IDs in the namespace's **openshift.io/sa.scc.supplemental-groups** annotation. Otherwise, the pod is not validated by that SCC and the next SCC is evaluated.

If the **SecurityContextConstraints.supplementalGroups** field has value **RunAsAny** and the pod specification omits the **Pod.spec.securityContext.supplementalGroups**, then this field is considered valid. Note that it is possible that during validation, other SCC settings will reject other pod fields and thus cause the pod to fail.

13.1.6. Security context constraints prioritization

Security context constraints (SCCs) have a priority field that affects the ordering when attempting to validate a request by the admission controller.

A priority value of **0** is the lowest possible priority. A nil priority is considered a **0**, or lowest, priority. Higher priority SCCs are moved to the front of the set when sorting.

When the complete set of available SCCs is determined, the SCCs are ordered in the following manner:

1. The highest priority SCCs are ordered first.
2. If the priorities are equal, the SCCs are sorted from most restrictive to least restrictive.
3. If both the priorities and restrictions are equal, the SCCs are sorted by name.

By default, the **anyuid** SCC granted to cluster administrators is given priority in their SCC set. This allows cluster administrators to run pods as any user by specifying **RunAsUser** in the pod's **SecurityContext**.

13.2. ABOUT PRE-ALLOCATED SECURITY CONTEXT CONSTRAINTS VALUES

The admission controller is aware of certain conditions in the security context constraints (SCCs) that trigger it to look up pre-allocated values from a namespace and populate the SCC before processing the pod. Each SCC strategy is evaluated independently of other strategies, with the pre-allocated values, where allowed, for each policy aggregated with pod specification values to make the final values for the various IDs defined in the running pod.

The following SCCs cause the admission controller to look for pre-allocated values when no ranges are defined in the pod specification:

1. A **RunAsUser** strategy of **MustRunAsRange** with no minimum or maximum set. Admission looks for the **openshift.io/sa.scc.uid-range** annotation to populate range fields.

2. An **SELinuxContext** strategy of **MustRunAs** with no level set. Admission looks for the **openshift.io/sa.scc.mcs** annotation to populate the level.
3. A **FSGroup** strategy of **MustRunAs**. Admission looks for the **openshift.io/sa.scc.supplemental-groups** annotation.
4. A **SupplementalGroups** strategy of **MustRunAs**. Admission looks for the **openshift.io/sa.scc.supplemental-groups** annotation.

During the generation phase, the security context provider uses default values for any parameter values that are not specifically set in the pod. Default values are based on the selected strategy:

1. **RunAsAny** and **MustRunAsNonRoot** strategies do not provide default values. If the pod needs a parameter value, such as a group ID, you must define the value in the pod specification.
2. **MustRunAs** (single value) strategies provide a default value that is always used. For example, for group IDs, even if the pod specification defines its own ID value, the namespace's default parameter value also appears in the pod's groups.
3. **MustRunAsRange** and **MustRunAs** (range-based) strategies provide the minimum value of the range. As with a single value **MustRunAs** strategy, the namespace's default parameter value appears in the running pod. If a range-based strategy is configurable with multiple ranges, it provides the minimum value of the first configured range.



NOTE

FSGroup and **SupplementalGroups** strategies fall back to the **openshift.io/sa.scc.uid-range** annotation if the **openshift.io/sa.scc.supplemental-groups** annotation does not exist on the namespace. If neither exists, the SCC is not created.



NOTE

By default, the annotation-based **FSGroup** strategy configures itself with a single range based on the minimum value for the annotation. For example, if your annotation reads **1/3**, the **FSGroup** strategy configures itself with a minimum and maximum value of **1**. If you want to allow more groups to be accepted for the **FSGroup** field, you can configure a custom SCC that does not use the annotation.



NOTE

The **openshift.io/sa.scc.supplemental-groups** annotation accepts a comma-delimited list of blocks in the format of **<start>/<length>** or **<start>-<end>**. The **openshift.io/sa.scc.uid-range** annotation accepts only a single block.

13.3. EXAMPLE SECURITY CONTEXT CONSTRAINTS

The following examples show how to define and use security context constraints (SCCs) in your cluster.

Annotated privileged SCC

```
allowHostDirVolumePlugin: true
allowHostIPC: true
allowHostNetwork: true
```

```

allowHostPID: true
allowHostPorts: true
allowPrivilegedContainer: true
allowedCapabilities:
- '*'
apiVersion: security.openshift.io/v1
defaultAddCapabilities: []
fsGroup:
  type: RunAsAny
groups:
- system:cluster-admins
- system:nodes
kind: SecurityContextConstraints
metadata:
  annotations:
    kubernetes.io/description: 'privileged allows access to all privileged and host
      features and the ability to run as any user, any group, any fsGroup, and with
      any SELinux context. WARNING: this is the most relaxed SCC and should be used
      only for cluster administration. Grant with caution.'
  creationTimestamp: null
  name: privileged
priority: null
readOnlyRootFilesystem: false
requiredDropCapabilities: null
runAsUser:
  type: RunAsAny
seLinuxContext:
  type: RunAsAny
seccompProfiles:
- '*'
supplementalGroups:
  type: RunAsAny
users:
- system:serviceaccount:default:registry
- system:serviceaccount:default:router
- system:serviceaccount:openshift-infra:build-controller
volumes:
- '*'

```

where, allowedCapabilities

A list of capabilities that a pod can request. An empty list means that none of capabilities can be requested while the special symbol * allows any capabilities.

defaultAddCapabilities

A list of additional capabilities that are added to any pod.

fsGroup

The **FSGroup** strategy, which dictates the allowable values for the security context.

groups

The groups that can access this SCC.

requiredDropCapabilities

A list of capabilities to drop from a pod. Or, specify **ALL** to drop all capabilities.

runAsUser

The **runAsUser** strategy type, which dictates the allowable values for the security context.

seLinuxContext

The **seLinuxContext** strategy type, which dictates the allowable values for the security context.

supplementalGroups

The **supplementalGroups** strategy, which dictates the allowable supplemental groups for the security context.

users

The users who can access this SCC.

volumes

The allowable volume types for the security context. In the example, * allows the use of all volume types.

The **users** and **groups** fields on the SCC control which users can access the SCC. By default, cluster administrators, nodes, and the build controller are granted access to the privileged SCC. All authenticated users are granted access to the **restricted-v2** SCC.



NOTE

When verifying access to an SCC, be aware of the following command behaviors:

- The **oc adm policy who-can use scc <scs_name>** and **oc auth can-i use scc/<scs_name>** commands evaluate only RBAC policies (**RoleBinding** or **ClusterRoleBinding** resources). Their output does not include users or groups configured directly in the SCC **users** and **groups** fields.
- The **oc describe scc <scs_name>** command displays only the users and groups configured directly within the SCC object. Its output does not include access granted through RBAC policies.

Without explicit runAsUser setting

```
apiVersion: v1
kind: Pod
metadata:
  name: security-context-demo
spec:
  securityContext:
    containers:
      - name: sec-ctx-demo
        image: gcr.io/google-samples/node-hello:1.0
```

When a container or pod does not request a user ID under which it should be run, the effective UID depends on the SCC that emits this pod. Because the **restricted-v2** SCC is granted to all authenticated users by default, it will be available to all users and service accounts and used in most cases. The **restricted-v2** SCC uses **MustRunAsRange** strategy for constraining and defaulting the possible values of the **securityContext.runAsUser** field. The admission plugin will look for the **openshift.io/sa.scc.uid-range** annotation on the current project to populate range fields, as it does not provide this range. In the end, a container will have **runAsUser** equal to the first value of the range that is hard to predict because every project has different ranges.

With explicit runAsUser setting

```

apiVersion: v1
kind: Pod
metadata:
  name: security-context-demo
spec:
  securityContext:
    runAsUser: 1000
  containers:
  - name: sec-ctx-demo
    image: gcr.io/google-samples/node-hello:1.0


```

A container or pod that requests a specific user ID will be accepted by OpenShift Dedicated only when a service account or a user is granted access to a SCC that allows such a user ID. The SCC can allow arbitrary IDs, an ID that falls into a range, or the exact user ID specific to the request.

This configuration is valid for SELinux, fsGroup, and Supplemental Groups.

13.4. CREATING SECURITY CONTEXT CONSTRAINTS FOR CUSTOMER CLOUD SUBSCRIPTION CLUSTERS

If the default security context constraints (SCCs) do not satisfy your application workload requirements, you can create a custom SCC by using the OpenShift CLI (**oc**).



IMPORTANT

Creating and modifying your own SCCs are advanced operations that might cause instability to your cluster. If you have questions about using your own SCCs, contact Red Hat Support. For information about contacting Red Hat support, see *Getting support*.



NOTE

In OpenShift Dedicated deployments, you can create your own SCCs only for clusters that use the Customer Cloud Subscription (CCS) model. You cannot create SCCs for OpenShift Dedicated clusters that use a Red Hat cloud account, because SCC resource creation requires **cluster-admin** privileges.

Prerequisites

- Install the OpenShift CLI (**oc**).
- Log in to the cluster as a user with the **cluster-admin** role.

Procedure

1. Define the SCC in a YAML file named **scc-admin.yaml**:

```

kind: SecurityContextConstraints
apiVersion: security.openshift.io/v1
metadata:
  name: scc-admin

```

```

allowPrivilegedContainer: true
runAsUser:
  type: RunAsAny
seLinuxContext:
  type: RunAsAny
fsGroup:
  type: RunAsAny
supplementalGroups:
  type: RunAsAny
users:
- my-admin-user
groups:
- my-admin-group

```

Optionally, you can drop specific capabilities for an SCC by setting the **requiredDropCapabilities** field with the desired values. Any specified capabilities are dropped from the container. To drop all capabilities, specify **ALL**. For example, to create an SCC that drops the **KILL**, **MKNOD**, and **SYS_CHROOT** capabilities, add the following to the SCC object:

```

requiredDropCapabilities:
- KILL
- MKNOD
- SYS_CHROOT

```



NOTE

You cannot list a capability in both **allowedCapabilities** and **requiredDropCapabilities**.

CRI-O supports the same list of capability values that are found in the [Docker documentation](#).

2. Create the SCC by passing in the file:

```
$ oc create -f scc-admin.yaml
```

Example output

```
securitycontextconstraints "scc-admin" created
```

Verification

- Verify that the SCC was created:

```
$ oc get scc scc-admin
```

Example output

```

NAME      PRIV  CAPS  SELINUX  RUNASUSER  FSGROUP  SUPGROUP
PRIORITY  READONLYROOTFS  VOLUMES
scc-admin true   []    RunAsAny RunAsAny  RunAsAny RunAsAny <none>  false

```

[awsElasticBlockStore azureDisk azureFile cephFS cinder configMap downwardAPI emptyDir fc flexVolume flocker gcePersistentDisk gitRepo glusterfs iscsi nfs persistentVolumeClaim photonPersistentDisk quobyte rbd secret vsphere]

13.5. CONFIGURING A WORKLOAD TO REQUIRE A SPECIFIC SCC

You can configure a workload to require a certain security context constraint (SCC). This is useful in scenarios where you want to pin a specific SCC to the workload or if you want to prevent your required SCC from being preempted by another SCC in the cluster.

To require a specific SCC, set the **openshift.io/required-scc** annotation on your workload. You can set this annotation on any resource that can set a pod manifest template, such as a deployment or daemon set.

The SCC must exist in the cluster and must be applicable to the workload, otherwise pod admission fails. An SCC is considered applicable to the workload if the user creating the pod or the pod's service account has **use** permissions for the SCC in the pod's namespace.



WARNING

Do not change the **openshift.io/required-scc** annotation in the live pod's manifest, because doing so causes the pod admission to fail. To change the required SCC, update the annotation in the underlying pod template, which causes the pod to be deleted and re-created.

Prerequisites

- The SCC must exist in the cluster.

Procedure

1. Create a YAML file for the deployment and specify a required SCC by setting the **openshift.io/required-scc** annotation:

Example deployment.yaml

```
apiVersion: config.openshift.io/v1
kind: Deployment
apiVersion: apps/v1
spec:
# ...
  template:
    metadata:
      annotations:
        openshift.io/required-scc: "my-scc"
# ...
```

2. Create the resource by running the following command:

```
$ oc create -f deployment.yaml
```

Verification

- Verify that the deployment used the specified SCC:
 - a. View the value of the pod's **openshift.io/scc** annotation by running the following command, replacing **<pod_name>** with the name of your deployment pod:

```
$ oc get pod <pod_name> -o jsonpath='{.metadata.annotations.openshift\.io\scc}{"\n"}'
```

1

- b. Examine the output and confirm that the displayed SCC matches the SCC that you defined in the deployment:

Example output

```
my-scc
```

13.6. ROLE-BASED ACCESS TO SECURITY CONTEXT CONSTRAINTS

You can specify SCCs as resources that are handled by RBAC. This allows you to scope access to your SCCs to a certain project or to the entire cluster. Assigning users, groups, or service accounts directly to an SCC retains cluster-wide scope.



IMPORTANT

Do not run workloads in or share access to default projects. Default projects are reserved for running core cluster components.

The following default projects are considered highly privileged: **default**, **kube-public**, **kube-system**, **openshift**, **openshift-infra**, **openshift-node**, and other system-created projects that have the **openshift.io/run-level** label set to **0** or **1**. Functionality that relies on admission plugins, such as pod security admission, security context constraints, cluster resource quotas, and image reference resolution, does not work in highly privileged projects.

To include access to SCCs for your role, specify the **scc** resource when creating a role.

```
$ oc create role <role-name> --verb=use --resource=scc --resource-name=<scc-name> -n
<namespace>
```

This results in the following role definition:

```
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  ...
  name: role-name
  namespace: namespace
  ...
rules:
```

```

- apiGroups:
- security.openshift.io
resourceNames:
- scc-name
resources:
- securitycontextconstraints
verbs:
- use

```

where, name

The name of the role.

namespace

The namespace of the defined role. Defaults to **default** if not specified.

apiGroups

The API group that includes the **SecurityContextConstraints** resource. Automatically defined when **scc** is specified as a resource.

resourceName

An example name for an SCC you want to have access.

resources

The name of the resource group that allows users to specify SCC names in the **resourceNames** field.

verbs

A list of verbs to apply to the role.

A local or cluster role with such a rule allows the subjects that are bound to it with a role binding or a cluster role binding to use the user-defined SCC called **scc-name**.

**NOTE**

Because RBAC is designed to prevent escalation, even project administrators are unable to grant access to an SCC. By default, they are not allowed to use the verb **use** on SCC resources, including the **restricted-v2** SCC.

13.7. REFERENCE OF SECURITY CONTEXT CONSTRAINTS COMMANDS

You can manage security context constraints (SCCs) in your instance as normal API objects by using the OpenShift CLI (**oc**).

13.7.1. Listing security context constraints

To get a current list of SCCs:

```
$ oc get scc
```

Example output

```

NAME                PRIV CAPS                SELINUX  RUNASUSER  FSGROUP
SUPGROUP  PRIORITY  READONLYROOTFS  VOLUMES

```

```

anyuid                false <no value>          MustRunAs RunAsAny      RunAsAny
RunAsAny 10          false
["configMap","downwardAPI","emptyDir","persistentVolumeClaim","projected","secret"]
hostaccess            false <no value>          MustRunAs MustRunAsRange  MustRunAs
RunAsAny <no value> false
["configMap","downwardAPI","emptyDir","hostPath","persistentVolumeClaim","projected","secret"]
hostmount-anyuid     false <no value>          MustRunAs RunAsAny      RunAsAny
RunAsAny <no value> false
["configMap","downwardAPI","emptyDir","hostPath","nfs","persistentVolumeClaim","projected","secret"]

hostnetwork           false <no value>          MustRunAs MustRunAsRange  MustRunAs
MustRunAs <no value> false
["configMap","downwardAPI","emptyDir","persistentVolumeClaim","projected","secret"]
hostnetwork-v2        false ["NET_BIND_SERVICE"] MustRunAs MustRunAsRange
MustRunAs MustRunAs <no value> false
["configMap","downwardAPI","emptyDir","persistentVolumeClaim","projected","secret"]
node-exporter         true  <no value>          RunAsAny RunAsAny      RunAsAny
RunAsAny <no value> false      ["*"]
nonroot               false <no value>          MustRunAs MustRunAsNonRoot RunAsAny
RunAsAny <no value> false
["configMap","downwardAPI","emptyDir","persistentVolumeClaim","projected","secret"]
nonroot-v2            false ["NET_BIND_SERVICE"] MustRunAs MustRunAsNonRoot
RunAsAny RunAsAny <no value> false
["configMap","downwardAPI","emptyDir","persistentVolumeClaim","projected","secret"]
privileged            true  ["*"]                  RunAsAny RunAsAny      RunAsAny RunAsAny
<no value> false      ["*"]
restricted            false <no value>          MustRunAs MustRunAsRange  MustRunAs
RunAsAny <no value> false
["configMap","downwardAPI","emptyDir","persistentVolumeClaim","projected","secret"]
restricted-v2         false ["NET_BIND_SERVICE"] MustRunAs MustRunAsRange
MustRunAs RunAsAny <no value> false
["configMap","downwardAPI","emptyDir","persistentVolumeClaim","projected","secret"]

```

13.7.2. Examining security context constraints

You can view information about a particular SCC, including which users, service accounts, and groups the SCC is applied to.

For example, to examine the **restricted** SCC:

```
$ oc describe scc restricted
```

Example output

```

Name:                restricted
Priority:             <none>
Access:
  Users:              <none>
  Groups:             <none>
Settings:
  Allow Privileged:   false
  Allow Privilege Escalation: true
  Default Add Capabilities: <none>
  Required Drop Capabilities: KILL,MKNOD,SETUID,SETGID
  Allowed Capabilities: <none>

```

```

Allowed Seccomp Profiles:      <none>
Allowed Volume Types:
configMap,downwardAPI,emptyDir,persistentVolumeClaim,projected,secret
Allowed Flexvolumes:         <all>
Allowed Unsafe Sysctls:      <none>
Forbidden Sysctls:           <none>
Allow Host Network:          false
Allow Host Ports:            false
Allow Host PID:              false
Allow Host IPC:              false
Read Only Root Filesystem:   false
Run As User Strategy: MustRunAsRange
  UID:                        <none>
  UID Range Min:              <none>
  UID Range Max:              <none>
SELinux Context Strategy: MustRunAs
  User:                        <none>
  Role:                        <none>
  Type:                        <none>
  Level:                       <none>
FSGroup Strategy: MustRunAs
  Ranges:                      <none>
Supplemental Groups Strategy: RunAsAny
  Ranges:                      <none>

```

where, Users

Lists which users and service accounts the SCC is applied to.

Groups

Lists which groups the SCC is applied to.

13.8. ADDITIONAL RESOURCES

- [Getting support](#)

CHAPTER 14. UNDERSTANDING AND MANAGING POD SECURITY ADMISSION

Pod security admission is an implementation of the [Kubernetes pod security standards](#). Use pod security admission to restrict the behavior of pods.

14.1. ABOUT POD SECURITY ADMISSION

OpenShift Dedicated includes [Kubernetes pod security admission](#). Pods that do not comply with the pod security admission defined globally or at the namespace level are not admitted to the cluster and cannot run.

Globally, the **privileged** profile is enforced, and the **restricted** profile is used for warnings and audits.

You can also configure the pod security admission settings at the namespace level.



IMPORTANT

Do not run workloads in or share access to default projects. Default projects are reserved for running core cluster components.

The following default projects are considered highly privileged: **default**, **kube-public**, **kube-system**, **openshift**, **openshift-infra**, **openshift-node**, and other system-created projects that have the **openshift.io/run-level** label set to **0** or **1**. Functionality that relies on admission plugins, such as pod security admission, security context constraints, cluster resource quotas, and image reference resolution, does not work in highly privileged projects.

14.1.1. Pod security admission modes

You can configure the following pod security admission modes for a namespace:

Table 14.1. Pod security admission modes

Mode	Label	Description
enforce	pod-security.kubernetes.io/enforce	Rejects a pod from admission if it does not comply with the set profile
audit	pod-security.kubernetes.io/audit	Logs audit events if a pod does not comply with the set profile
warn	pod-security.kubernetes.io/warn	Displays warnings if a pod does not comply with the set profile

14.1.2. Pod security admission profiles

You can set each of the pod security admission modes to one of the following profiles:

Table 14.2. Pod security admission profiles

Profile	Description
privileged	Least restrictive policy; allows for known privilege escalation
baseline	Minimally restrictive policy; prevents known privilege escalations
restricted	Most restrictive policy; follows current pod hardening best practices

14.1.3. Privileged namespaces

The following system namespaces are always set to the **privileged** pod security admission profile:

- **default**
- **kube-public**
- **kube-system**

You cannot change the pod security profile for these privileged namespaces.

Example privileged namespace configuration

```
apiVersion: v1
kind: Namespace
metadata:
  labels:
    openshift.io/cluster-monitoring: "true"
    pod-security.kubernetes.io/enforce: privileged
    pod-security.kubernetes.io/audit: privileged
    pod-security.kubernetes.io/warn: privileged
  name: "<mig_namespace>"
# ...
```

14.1.4. Pod security admission and security context constraints

Pod security admission standards and security context constraints are reconciled and enforced by two independent controllers. The two controllers work independently using the following processes to enforce security policies:

1. The security context constraint controller may mutate some security context fields per the pod's assigned SCC. For example, if the `seccomp` profile is empty or not set and if the pod's assigned SCC enforces `seccompProfiles` field to be **runtime/default**, the controller sets the default type to **RuntimeDefault**.
2. The security context constraint controller validates the pod's security context against the matching SCC.
3. The pod security admission controller validates the pod's security context against the pod security standard assigned to the namespace.

14.2. ABOUT POD SECURITY ADMISSION SYNCHRONIZATION

In addition to the global pod security admission control configuration, a controller applies pod security admission control **warn** and **audit** labels to namespaces according to the SCC permissions of the service accounts that are in a given namespace.

The controller examines **ServiceAccount** object permissions to use security context constraints in each namespace. Security context constraints (SCCs) are mapped to pod security profiles based on their field values; the controller uses these translated profiles. Pod security admission **warn** and **audit** labels are set to the most privileged pod security profile in the namespace to prevent displaying warnings and logging audit events when pods are created.

Namespace labeling is based on consideration of namespace-local service account privileges.

Applying pods directly might use the SCC privileges of the user who runs the pod. However, user privileges are not considered during automatic labeling.

14.2.1. Pod security admission synchronization namespace exclusions

Pod security admission synchronization is permanently disabled on system-created namespaces and **openshift-*** prefixed namespaces.

Namespaces that are defined as part of the cluster payload have pod security admission synchronization disabled permanently. The following namespaces are permanently disabled:

- **default**
- **kube-node-lease**
- **kube-system**
- **kube-public**
- **openshift**
- All system-created namespaces that are prefixed with **openshift-**

14.3. CONTROLLING POD SECURITY ADMISSION SYNCHRONIZATION

You can enable or disable automatic pod security admission synchronization for most namespaces.



IMPORTANT

You cannot enable pod security admission synchronization on system-created namespaces. For more information, see *Pod security admission synchronization namespace exclusions*.


Procedure

- For each namespace that you want to configure, set a value for the **security.openshift.io/scc.podSecurityLabelSync** label:
 - To disable pod security admission label synchronization in a namespace, set the value of the **security.openshift.io/scc.podSecurityLabelSync** label to **false**.
Run the following command:

```
$ oc label namespace <namespace>
security.openshift.io/scc.podSecurityLabelSync=false
```

- To enable pod security admission label synchronization in a namespace, set the value of the **security.openshift.io/scc.podSecurityLabelSync** label to **true**.
Run the following command:

```
$ oc label namespace <namespace>
security.openshift.io/scc.podSecurityLabelSync=true
```



NOTE

Use the **--overwrite** flag to overwrite the value if this label is already set on the namespace.

Additional resources

- [Pod security admission synchronization namespace exclusions](#)

14.4. CONFIGURING POD SECURITY ADMISSION FOR A NAMESPACE

You can configure the pod security admission settings at the namespace level. For each of the pod security admission modes on the namespace, you can set which pod security admission profile to use.

Procedure

- For each pod security admission mode that you want to set on a namespace, run the following command:

```
$ oc label namespace <namespace> \
  pod-security.kubernetes.io/<mode>=<profile> \
  --overwrite
```

1 **1** Set **<namespace>** to the namespace to configure.

2 Set **<mode>** to **enforce**, **warn**, or **audit**. Set **<profile>** to **restricted**, **baseline**, or **privileged**.

14.5. ABOUT POD SECURITY ADMISSION ALERTS

A **PodSecurityViolation** alert is triggered when the Kubernetes API server reports that there is a pod denial on the audit level of the pod security admission controller. This alert persists for one day.

View the Kubernetes API server audit logs to investigate alerts that were triggered. As an example, a workload is likely to fail admission if global enforcement is set to the **restricted** pod security level.

For assistance in identifying pod security admission violation audit events, see [Audit annotations](#) in the Kubernetes documentation.

14.6. ADDITIONAL RESOURCES

- [Viewing audit logs](#)
- [Managing security context constraints](#)

CHAPTER 15. SYNCING LDAP GROUPS

As an administrator with the **dedicated-admin** role, you can use groups to manage users, change their permissions, and enhance collaboration. Your organization may have already created user groups and stored them in an LDAP server. OpenShift Dedicated can sync those LDAP records with internal OpenShift Dedicated records, enabling you to manage your groups in one place. OpenShift Dedicated currently supports group sync with LDAP servers using three common schemas for defining group membership: RFC 2307, Active Directory, and augmented Active Directory.

For more information on configuring LDAP, see [Configuring an LDAP identity provider](#).



NOTE

You must have **dedicated-admin** privileges to sync groups.

15.1. ABOUT CONFIGURING LDAP SYNC

Before you can run LDAP sync, you need a sync configuration file. This file contains the following LDAP client configuration details:

- Configuration for connecting to your LDAP server.
- Sync configuration options that are dependent on the schema used in your LDAP server.
- An administrator-defined list of name mappings that maps OpenShift Dedicated group names to groups in your LDAP server.

The format of the configuration file depends upon the schema you are using: RFC 2307, Active Directory, or augmented Active Directory.

LDAP client configuration

The LDAP client configuration section of the configuration defines the connections to your LDAP server.

The LDAP client configuration section of the configuration defines the connections to your LDAP server.

LDAP client configuration

```
url: ldap://10.0.0.0:389 1
bindDN: cn=admin,dc=example,dc=com 2
bindPassword: <password> 3
insecure: false 4
ca: my-ldap-ca-bundle.crt 5
```

- 1 The connection protocol, IP address of the LDAP server hosting your database, and the port to connect to, formatted as **scheme://host:port**.
- 2 Optional distinguished name (DN) to use as the Bind DN. OpenShift Dedicated uses this if elevated privilege is required to retrieve entries for the sync operation.
- 3 Optional password to use to bind. OpenShift Dedicated uses this if elevated privilege is necessary to retrieve entries for the sync operation. This value may also be provided in an environment variable, external file, or encrypted file.

- 4 When **false**, secure LDAP (**ldaps://**) URLs connect using TLS, and insecure LDAP (**ldap://**) URLs are upgraded to TLS. When **true**, no TLS connection is made to the server and you cannot use
- 5 The certificate bundle to use for validating server certificates for the configured URL. If empty, OpenShift Dedicated uses system-trusted roots. This only applies if **insecure** is set to **false**.

LDAP query definition

Sync configurations consist of LDAP query definitions for the entries that are required for synchronization. The specific definition of an LDAP query depends on the schema used to store membership information in the LDAP server.

LDAP query definition

```
baseDN: ou=users,dc=example,dc=com 1
scope: sub 2
derefAliases: never 3
timeout: 0 4
filter: (objectClass=person) 5
pageSize: 0 6
```

- 1 The distinguished name (DN) of the branch of the directory where all searches will start from. It is required that you specify the top of your directory tree, but you can also specify a subtree in the directory.
- 2 The scope of the search. Valid values are **base**, **one**, or **sub**. If this is left undefined, then a scope of **sub** is assumed. Descriptions of the scope options can be found in the table below.
- 3 The behavior of the search with respect to aliases in the LDAP tree. Valid values are **never**, **search**, **base**, or **always**. If this is left undefined, then the default is to **always** dereference aliases. Descriptions of the dereferencing behaviors can be found in the table below.
- 4 The time limit allowed for the search by the client, in seconds. A value of **0** imposes no client-side limit.
- 5 A valid LDAP search filter. If this is left undefined, then the default is **(objectClass=*)**.
- 6 The optional maximum size of response pages from the server, measured in LDAP entries. If set to **0**, no size restrictions will be made on pages of responses. Setting paging sizes is necessary when queries return more entries than the client or server allow by default.

Table 15.1. LDAP search scope options

LDAP search scope	Description
base	Only consider the object specified by the base DN given for the query.
one	Consider all of the objects on the same level in the tree as the base DN for the query.
sub	Consider the entire subtree rooted at the base DN given for the query.

Table 15.2. LDAP dereferencing behaviors

Dereferencing behavior	Description
never	Never dereference any aliases found in the LDAP tree.
search	Only dereference aliases found while searching.
base	Only dereference aliases while finding the base object.
always	Always dereference all aliases found in the LDAP tree.

User-defined name mapping

A user-defined name mapping explicitly maps the names of OpenShift Dedicated groups to unique identifiers that find groups on your LDAP server. The mapping uses normal YAML syntax. A user-defined mapping can contain an entry for every group in your LDAP server or only a subset of those groups. If there are groups on the LDAP server that do not have a user-defined name mapping, the default behavior during sync is to use the attribute specified as the OpenShift Dedicated group's name.


User-defined name mapping

```
groupUIDNameMapping:
  "cn=group1,ou=groups,dc=example,dc=com": firstgroup
  "cn=group2,ou=groups,dc=example,dc=com": secondgroup
  "cn=group3,ou=groups,dc=example,dc=com": thirdgroup
```

15.1.1. About the RFC 2307 configuration file

The RFC 2307 schema requires you to provide an LDAP query definition for both user and group entries, as well as the attributes with which to represent them in the internal OpenShift Dedicated records.

For clarity, the group you create in OpenShift Dedicated should use attributes other than the distinguished name whenever possible for user- or administrator-facing fields. For example, identify the users of an OpenShift Dedicated group by their e-mail, and use the name of the group as the common name. The following configuration file creates these relationships:



NOTE

If using user-defined name mappings, your configuration file will differ.

LDAP sync configuration that uses RFC 2307 schema: `rfc2307_config.yaml`

```
kind: LDAPSyncConfig
apiVersion: v1
url: ldap://LDAP_SERVICE_IP:389 1
insecure: false 2
bindDN: cn=admin,dc=example,dc=com
bindPassword:
```

```

file: "/etc/secrets/bindPassword"
rfc2307:
  groupsQuery:
    baseDN: "ou=groups,dc=example,dc=com"
    scope: sub
    derefAliases: never
    pageSize: 0
  groupUIDAttribute: dn ❸
  groupNameAttributes: [ cn ] ❹
  groupMembershipAttributes: [ member ] ❺
  usersQuery:
    baseDN: "ou=users,dc=example,dc=com"
    scope: sub
    derefAliases: never
    pageSize: 0
  userUIDAttribute: dn ❻
  userNameAttributes: [ mail ] ❼
  tolerateMemberNotFoundErrors: false
  tolerateMemberOutOfScopeErrors: false

```

- ❶ The IP address and host of the LDAP server where this group's record is stored.
- ❷ When **false**, secure LDAP (**ldaps://**) URLs connect using TLS, and insecure LDAP (**ldap://**) URLs are upgraded to TLS. When **true**, no TLS connection is made to the server and you cannot use **ldaps://** URL schemes.
- ❸ The attribute that uniquely identifies a group on the LDAP server. You cannot specify **groupsQuery** filters when using DN for **groupUIDAttribute**. For fine-grained filtering, use the whitelist / blacklist method.
- ❹ The attribute to use as the name of the group.
- ❺ The attribute on the group that stores the membership information.
- ❻ The attribute that uniquely identifies a user on the LDAP server. You cannot specify **usersQuery** filters when using DN for **userUIDAttribute**. For fine-grained filtering, use the whitelist / blacklist method.
- ❼ The attribute to use as the name of the user in the OpenShift Dedicated group record.

15.1.2. About the Active Directory configuration file

The Active Directory schema requires you to provide an LDAP query definition for user entries, as well as the attributes to represent them with in the internal OpenShift Dedicated group records.

For clarity, the group you create in OpenShift Dedicated should use attributes other than the distinguished name whenever possible for user- or administrator-facing fields. For example, identify the users of an OpenShift Dedicated group by their e-mail, but define the name of the group by the name of the group on the LDAP server. The following configuration file creates these relationships:

LDAP sync configuration that uses Active Directory schema: `active_directory_config.yaml`

```

kind: LDAPSvcConfig
apiVersion: v1

```

```
url: ldap://LDAP_SERVICE_IP:389
activeDirectory:
  usersQuery:
    baseDN: "ou=users,dc=example,dc=com"
    scope: sub
    derefAliases: never
    filter: (objectclass=person)
    pageSize: 0
  userNameAttributes: [ mail ] 1
  groupMembershipAttributes: [ memberOf ] 2
```

- 1 The attribute to use as the name of the user in the OpenShift Dedicated group record.
- 2 The attribute on the user that stores the membership information.

15.1.3. About the augmented Active Directory configuration file

The augmented Active Directory schema requires you to provide an LDAP query definition for both user entries and group entries, as well as the attributes with which to represent them in the internal OpenShift Dedicated group records.

For clarity, the group you create in OpenShift Dedicated should use attributes other than the distinguished name whenever possible for user- or administrator-facing fields. For example, identify the users of an OpenShift Dedicated group by their e-mail, and use the name of the group as the common name. The following configuration file creates these relationships.

LDAP sync configuration that uses augmented Active Directory schema: augmented_active_directory_config.yaml

```
kind: LDAPSyncConfig
apiVersion: v1
url: ldap://LDAP_SERVICE_IP:389
augmentedActiveDirectory:
  groupsQuery:
    baseDN: "ou=groups,dc=example,dc=com"
    scope: sub
    derefAliases: never
    pageSize: 0
  groupUIDAttribute: dn 1
  groupNameAttributes: [ cn ] 2
  usersQuery:
    baseDN: "ou=users,dc=example,dc=com"
    scope: sub
    derefAliases: never
    filter: (objectclass=person)
    pageSize: 0
  userNameAttributes: [ mail ] 3
  groupMembershipAttributes: [ memberOf ] 4
```

- 1 The attribute that uniquely identifies a group on the LDAP server. You cannot specify **groupsQuery** filters when using DN for groupUIDAttribute. For fine-grained filtering, use the whitelist / blacklist method.

- 2 The attribute to use as the name of the group.
- 3 The attribute to use as the name of the user in the OpenShift Dedicated group record.
- 4 The attribute on the user that stores the membership information.

15.2. RUNNING LDAP SYNC

Once you have created a sync configuration file, you can begin to sync. OpenShift Dedicated allows administrators to perform a number of different sync types with the same server.

15.2.1. Syncing the LDAP server with OpenShift Dedicated

You can sync all groups from the LDAP server with OpenShift Dedicated.

Prerequisites

- Create a sync configuration file.
- You have access to the cluster as a user with the **dedicated-admin** role.

Procedure

- To sync all groups from the LDAP server with OpenShift Dedicated:

```
$ oc adm groups sync --sync-config=config.yaml --confirm
```



NOTE

By default, all group synchronization operations are dry-run, so you must set the **-confirm** flag on the **oc adm groups sync** command to make changes to OpenShift Dedicated group records.

15.2.2. Syncing OpenShift Dedicated groups with the LDAP server

You can sync all groups already in OpenShift Dedicated that correspond to groups in the LDAP server specified in the configuration file.


Prerequisites

- Create a sync configuration file.
- You have access to the cluster as a user with the **dedicated-admin** role.

Procedure

- To sync OpenShift Dedicated groups with the LDAP server:

```
$ oc adm groups sync --type=openshift --sync-config=config.yaml --confirm
```

**NOTE**

By default, all group synchronization operations are dry-run, so you must set the **-confirm** flag on the **oc adm groups sync** command to make changes to OpenShift Dedicated group records.

15.2.3. Syncing subgroups from the LDAP server with OpenShift Dedicated

You can sync a subset of LDAP groups with OpenShift Dedicated using whitelist files, blacklist files, or both.

**NOTE**

You can use any combination of blacklist files, whitelist files, or whitelist literals. Whitelist and blacklist files must contain one unique group identifier per line, and you can include whitelist literals directly in the command itself. These guidelines apply to groups found on LDAP servers as well as groups already present in OpenShift Dedicated.

Prerequisites

- Create a sync configuration file.
- You have access to the cluster as a user with the **dedicated-admin** role.

Procedure

- To sync a subset of LDAP groups with OpenShift Dedicated, use any the following commands:

```
$ oc adm groups sync --whitelist=<whitelist_file> \
  --sync-config=config.yaml \
  --confirm
```

```
$ oc adm groups sync --blacklist=<blacklist_file> \
  --sync-config=config.yaml \
  --confirm
```

```
$ oc adm groups sync <group_unique_identifier> \
  --sync-config=config.yaml \
  --confirm
```

```
$ oc adm groups sync <group_unique_identifier> \
  --whitelist=<whitelist_file> \
  --blacklist=<blacklist_file> \
  --sync-config=config.yaml \
  --confirm
```

```
$ oc adm groups sync --type=openshift \
  --whitelist=<whitelist_file> \
  --sync-config=config.yaml \
  --confirm
```

**NOTE**

By default, all group synchronization operations are dry-run, so you must set the **-confirm** flag on the **oc adm groups sync** command to make changes to OpenShift Dedicated group records.

15.3. RUNNING A GROUP PRUNING JOB

An administrator can also choose to remove groups from OpenShift Dedicated records if the records on the LDAP server that created them are no longer present. The prune job will accept the same sync configuration file and whitelists or blacklists as used for the sync job.

For example:

```
$ oc adm prune groups --sync-config=/path/to/ldap-sync-config.yaml --confirm
```

```
$ oc adm prune groups --whitelist=/path/to/whitelist.txt --sync-config=/path/to/ldap-sync-config.yaml --confirm
```

```
$ oc adm prune groups --blacklist=/path/to/blacklist.txt --sync-config=/path/to/ldap-sync-config.yaml --confirm
```

15.4. LDAP GROUP SYNC EXAMPLES

This section contains examples for the RFC 2307, Active Directory, and augmented Active Directory schemas.

**NOTE**

These examples assume that all users are direct members of their respective groups. Specifically, no groups have other groups as members. See the Nested Membership Sync Example for information on how to sync nested groups.

15.4.1. Syncing groups using the RFC 2307 schema

For the RFC 2307 schema, the following examples synchronize a group named **admins** that has two members: **Jane** and **Jim**. The examples explain:

- How the group and users are added to the LDAP server.
- What the resulting group record in OpenShift Dedicated will be after synchronization.

**NOTE**

These examples assume that all users are direct members of their respective groups. Specifically, no groups have other groups as members. See the Nested Membership Sync Example for information on how to sync nested groups.

In the RFC 2307 schema, both users (Jane and Jim) and groups exist on the LDAP server as first-class entries, and group membership is stored in attributes on the group. The following snippet of **ldif** defines the users and group for this schema:

LDAP entries that use RFC 2307 schema: `rfc2307.ldif`

```

dn: ou=users,dc=example,dc=com
objectClass: organizationalUnit
ou: users
dn: cn=Jane,ou=users,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
cn: Jane
sn: Smith
displayName: Jane Smith
mail: jane.smith@example.com
dn: cn=Jim,ou=users,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
cn: Jim
sn: Adams
displayName: Jim Adams
mail: jim.adams@example.com
dn: ou=groups,dc=example,dc=com
objectClass: organizationalUnit
ou: groups
dn: cn=admins,ou=groups,dc=example,dc=com 1
objectClass: groupOfNames
cn: admins
owner: cn=admin,dc=example,dc=com
description: System Administrators
member: cn=Jane,ou=users,dc=example,dc=com 2
member: cn=Jim,ou=users,dc=example,dc=com

```

- 1 The group is a first-class entry in the LDAP server.
- 2 Members of a group are listed with an identifying reference as attributes on the group.

Prerequisites

- Create the configuration file.
- You have access to the cluster as a user with the **dedicated-admin** role.

Procedure

- Run the sync with the **`rfc2307_config.yaml`** file:

```
$ oc adm groups sync --sync-config=rfc2307_config.yaml --confirm
```

OpenShift Dedicated creates the following group record as a result of the above sync operation:

OpenShift Dedicated group created by using the `rfc2307_config.yaml` file

```

apiVersion: user.openshift.io/v1
kind: Group
metadata:
  annotations:
    openshift.io/ldap.sync-time: 2015-10-13T10:08:38-0400 1
    openshift.io/ldap.uid: cn=admins,ou=groups,dc=example,dc=com 2
    openshift.io/ldap.url: LDAP_SERVER_IP:389 3
  creationTimestamp:
    name: admins 4
  users: 5
  - jane.smith@example.com
  - jim.adams@example.com

```

- 1 The last time this OpenShift Dedicated group was synchronized with the LDAP server, in ISO 6801 format.
- 2 The unique identifier for the group on the LDAP server.
- 3 The IP address and host of the LDAP server where this group's record is stored.
- 4 The name of the group as specified by the sync file.
- 5 The users that are members of the group, named as specified by the sync file.

15.4.2. Syncing groups using the RFC2307 schema with user-defined name mappings

When syncing groups with user-defined name mappings, the configuration file changes to contain these mappings as shown below.

LDAP sync configuration that uses RFC 2307 schema with user-defined name mappings: rfc2307_config_user_defined.yaml

```

kind: LDAPSyncConfig
apiVersion: v1
groupUIDNameMapping:
  "cn=admins,ou=groups,dc=example,dc=com": Administrators 1
rfc2307:
  groupsQuery:
    baseDN: "ou=groups,dc=example,dc=com"
    scope: sub
    derefAliases: never
    pageSize: 0
  groupUIDAttribute: dn 2
  groupNameAttributes: [ cn ] 3
  groupMembershipAttributes: [ member ]
  usersQuery:
    baseDN: "ou=users,dc=example,dc=com"
    scope: sub
    derefAliases: never
    pageSize: 0
  userUIDAttribute: dn 4

```

```

userNameAttributes: [ mail ]
tolerateMemberNotFoundErrors: false
tolerateMemberOutOfScopeErrors: false

```

- 1 The user-defined name mapping.
- 2 The unique identifier attribute that is used for the keys in the user-defined name mapping. You cannot specify **groupsQuery** filters when using DN for groupUIDAttribute. For fine-grained filtering, use the whitelist / blacklist method.
- 3 The attribute to name OpenShift Dedicated groups with if their unique identifier is not in the user-defined name mapping.
- 4 The attribute that uniquely identifies a user on the LDAP server. You cannot specify **usersQuery** filters when using DN for userUIDAttribute. For fine-grained filtering, use the whitelist / blacklist method.

Prerequisites

- Create the configuration file.
- You have access to the cluster as a user with the **dedicated-admin** role.

Procedure

- Run the sync with the **rfc2307_config_user_defined.yaml** file:

```
$ oc adm groups sync --sync-config=rfc2307_config_user_defined.yaml --confirm
```

OpenShift Dedicated creates the following group record as a result of the above sync operation:

OpenShift Dedicated group created by using the **rfc2307_config_user_defined.yaml** file

```

apiVersion: user.openshift.io/v1
kind: Group
metadata:
  annotations:
    openshift.io/ldap.sync-time: 2015-10-13T10:08:38-0400
    openshift.io/ldap.uid: cn=admins,ou=groups,dc=example,dc=com
    openshift.io/ldap.url: LDAP_SERVER_IP:389
  creationTimestamp:
  name: Administrators 1
users:
- jane.smith@example.com
- jim.adams@example.com

```

- 1 The name of the group as specified by the user-defined name mapping.

15.4.3. Syncing groups using RFC 2307 with user-defined error tolerances

By default, if the groups being synced contain members whose entries are outside of the scope defined in the member query, the group sync fails with an error:

```
Error determining LDAP group membership for "<group>": membership lookup for user "<user>" in group "<group>" failed because of "search for entry with dn="<user-dn>" would search outside of the base dn specified (dn="<base-dn>")".
```

This often indicates a misconfigured **baseDN** in the **usersQuery** field. However, in cases where the **baseDN** intentionally does not contain some of the members of the group, setting **tolerateMemberOutOfScopeErrors: true** allows the group sync to continue. Out of scope members will be ignored.

Similarly, when the group sync process fails to locate a member for a group, it fails outright with errors:

```
Error determining LDAP group membership for "<group>": membership lookup for user "<user>" in group "<group>" failed because of "search for entry with base dn="<user-dn>" refers to a non-existent entry".
Error determining LDAP group membership for "<group>": membership lookup for user "<user>" in group "<group>" failed because of "search for entry with base dn="<user-dn>" and filter "<filter>" did not return any results".
```

This often indicates a misconfigured **usersQuery** field. However, in cases where the group contains member entries that are known to be missing, setting **tolerateMemberNotFoundErrors: true** allows the group sync to continue. Problematic members will be ignored.



WARNING

Enabling error tolerances for the LDAP group sync causes the sync process to ignore problematic member entries. If the LDAP group sync is not configured correctly, this could result in synced OpenShift Dedicated groups missing members.

LDAP entries that use RFC 2307 schema with problematic group membership: `rfc2307_problematic_users.ldif`

```
dn: ou=users,dc=example,dc=com
objectClass: organizationalUnit
ou: users
dn: cn=Jane,ou=users,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
cn: Jane
sn: Smith
displayName: Jane Smith
mail: jane.smith@example.com
dn: cn=Jim,ou=users,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
```

```

cn: Jim
sn: Adams
displayName: Jim Adams
mail: jim.adams@example.com
dn: ou=groups,dc=example,dc=com
objectClass: organizationalUnit
ou: groups
dn: cn=admins,ou=groups,dc=example,dc=com
objectClass: groupOfNames
cn: admins
owner: cn=admin,dc=example,dc=com
description: System Administrators
member: cn=Jane,ou=users,dc=example,dc=com
member: cn=Jim,ou=users,dc=example,dc=com
member: cn=INVALID,ou=users,dc=example,dc=com 1
member: cn=Jim,ou=OUTOFSCOPE,dc=example,dc=com 2

```

- 1 A member that does not exist on the LDAP server.
- 2 A member that may exist, but is not under the **baseDN** in the user query for the sync job.

To tolerate the errors in the above example, the following additions to your sync configuration file must be made:

LDAP sync configuration that uses RFC 2307 schema tolerating errors: rfc2307_config_tolerating.yamll

```

kind: LDAPSvcConfig
apiVersion: v1
url: ldap://LDAP_SERVICE_IP:389
rfc2307:
  groupsQuery:
    baseDN: "ou=groups,dc=example,dc=com"
    scope: sub
    derefAliases: never
  groupUIDAttribute: dn
  groupNameAttributes: [ cn ]
  groupMembershipAttributes: [ member ]
  usersQuery:
    baseDN: "ou=users,dc=example,dc=com"
    scope: sub
    derefAliases: never
  userUIDAttribute: dn 1
  userNameAttributes: [ mail ]
  tolerateMemberNotFoundErrors: true 2
  tolerateMemberOutOfScopeErrors: true 3

```

- 1 The attribute that uniquely identifies a user on the LDAP server. You cannot specify **usersQuery** filters when using DN for userUIDAttribute. For fine-grained filtering, use the whitelist / blacklist method.
- 2 When **true**, the sync job tolerates groups for which some members were not found, and members whose LDAP entries are not found are ignored. The default behavior for the sync job is to fail if a member of a group is not found.

- 3 When **true**, the sync job tolerates groups for which some members are outside the user scope given in the **usersQuery** base DN, and members outside the member query scope are ignored.

Prerequisites

- Create the configuration file.
- You have access to the cluster as a user with the **dedicated-admin** role.

Procedure

- Run the sync with the **rfc2307_config_tolerating.yaml** file:

```
$ oc adm groups sync --sync-config=rfc2307_config_tolerating.yaml --confirm
```

OpenShift Dedicated creates the following group record as a result of the above sync operation:

OpenShift Dedicated group created by using the **rfc2307_config.yaml** file

```
apiVersion: user.openshift.io/v1
kind: Group
metadata:
  annotations:
    openshift.io/ldap.sync-time: 2015-10-13T10:08:38-0400
    openshift.io/ldap.uid: cn=admins,ou=groups,dc=example,dc=com
    openshift.io/ldap.url: LDAP_SERVER_IP:389
  creationTimestamp:
    name: admins
  users: 1
  - jane.smith@example.com
  - jim.adams@example.com
```

- 1 The users that are members of the group, as specified by the sync file. Members for which lookup encountered tolerated errors are absent.

15.4.4. Syncing groups using the Active Directory schema

In the Active Directory schema, both users (Jane and Jim) exist in the LDAP server as first-class entries, and group membership is stored in attributes on the user. The following snippet of **ldif** defines the users and group for this schema:

LDAP entries that use Active Directory schema: **active_directory.ldif**

```
dn: ou=users,dc=example,dc=com
objectClass: organizationalUnit
ou: users

dn: cn=Jane,ou=users,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
```

```
objectClass: testPerson
cn: Jane
sn: Smith
displayName: Jane Smith
mail: jane.smith@example.com
memberOf: admins ❶

dn: cn=Jim,ou=users,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
objectClass: testPerson
cn: Jim
sn: Adams
displayName: Jim Adams
mail: jim.adams@example.com
memberOf: admins
```

- ❶ The user's group memberships are listed as attributes on the user, and the group does not exist as an entry on the server. The **memberOf** attribute does not have to be a literal attribute on the user; in some LDAP servers, it is created during search and returned to the client, but not committed to the database.

Prerequisites

- Create the configuration file.
- You have access to the cluster as a user with the **dedicated-admin** role.

Procedure

- Run the sync with the **active_directory_config.yaml** file:

```
$ oc adm groups sync --sync-config=active_directory_config.yaml --confirm
```

OpenShift Dedicated creates the following group record as a result of the above sync operation:

OpenShift Dedicated group created by using the **active_directory_config.yaml** file

```
apiVersion: user.openshift.io/v1
kind: Group
metadata:
  annotations:
    openshift.io/ldap.sync-time: 2015-10-13T10:08:38-0400 ❶
    openshift.io/ldap.uid: admins ❷
    openshift.io/ldap.url: LDAP_SERVER_IP:389 ❸
  creationTimestamp:
    name: admins ❹
  users: ❺
  - jane.smith@example.com
  - jim.adams@example.com
```

- 1 The last time this OpenShift Dedicated group was synchronized with the LDAP server, in ISO 6801 format.
- 2 The unique identifier for the group on the LDAP server.
- 3 The IP address and host of the LDAP server where this group's record is stored.
- 4 The name of the group as listed in the LDAP server.
- 5 The users that are members of the group, named as specified by the sync file.

15.4.5. Syncing groups using the augmented Active Directory schema

In the augmented Active Directory schema, both users (Jane and Jim) and groups exist in the LDAP server as first-class entries, and group membership is stored in attributes on the user. The following snippet of **ldif** defines the users and group for this schema:

LDAP entries that use augmented Active Directory schema: `augmented_active_directory.ldif`

```
dn: ou=users,dc=example,dc=com
objectClass: organizationalUnit
ou: users

dn: cn=Jane,ou=users,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
objectClass: testPerson
cn: Jane
sn: Smith
displayName: Jane Smith
mail: jane.smith@example.com
memberOf: cn=admin,ou=groups,dc=example,dc=com 1

dn: cn=Jim,ou=users,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
objectClass: testPerson
cn: Jim
sn: Adams
displayName: Jim Adams
mail: jim.adams@example.com
memberOf: cn=admin,ou=groups,dc=example,dc=com

dn: ou=groups,dc=example,dc=com
objectClass: organizationalUnit
ou: groups

dn: cn=admin,ou=groups,dc=example,dc=com 2
objectClass: groupOfNames
cn: admin
owner: cn=admin,dc=example,dc=com
```

```
description: System Administrators
member: cn=Jane,ou=users,dc=example,dc=com
member: cn=Jim,ou=users,dc=example,dc=com
```

- 1 The user's group memberships are listed as attributes on the user.
- 2 The group is a first-class entry on the LDAP server.

Prerequisites

- Create the configuration file.
- You have access to the cluster as a user with the **dedicated-admin** role.

Procedure

- Run the sync with the **augmented_active_directory_config.yaml** file:

```
$ oc adm groups sync --sync-config=augmented_active_directory_config.yaml --confirm
```

OpenShift Dedicated creates the following group record as a result of the above sync operation:

OpenShift Dedicated group created by using the **augmented_active_directory_config.yaml** file

```
apiVersion: user.openshift.io/v1
kind: Group
metadata:
  annotations:
    openshift.io/ldap.sync-time: 2015-10-13T10:08:38-0400 1
    openshift.io/ldap.uid: cn=admins,ou=groups,dc=example,dc=com 2
    openshift.io/ldap.url: LDAP_SERVER_IP:389 3
  creationTimestamp:
  name: admins 4
users: 5
- jane.smith@example.com
- jim.adams@example.com
```

- 1 The last time this OpenShift Dedicated group was synchronized with the LDAP server, in ISO 6801 format.
- 2 The unique identifier for the group on the LDAP server.
- 3 The IP address and host of the LDAP server where this group's record is stored.
- 4 The name of the group as specified by the sync file.
- 5 The users that are members of the group, named as specified by the sync file.

15.4.5.1. LDAP nested membership sync example

Groups in OpenShift Dedicated do not nest. The LDAP server must flatten group membership before the data can be consumed. Microsoft's Active Directory Server supports this feature via the **LDAP_MATCHING_RULE_IN_CHAIN** rule, which has the OID **1.2.840.113556.1.4.1941**. Furthermore, only explicitly whitelisted groups can be synced when using this matching rule.

This section has an example for the augmented Active Directory schema, which synchronizes a group named **admins** that has one user **Jane** and one group **otheradmins** as members. The **otheradmins** group has one user member: **Jim**. This example explains:

- How the group and users are added to the LDAP server.
- What the LDAP sync configuration file looks like.
- What the resulting group record in OpenShift Dedicated will be after synchronization.

In the augmented Active Directory schema, both users (**Jane** and **Jim**) and groups exist in the LDAP server as first-class entries, and group membership is stored in attributes on the user or the group. The following snippet of **ldif** defines the users and groups for this schema:

LDAP entries that use augmented Active Directory schema with nested members: **augmented_active_directory_nested.ldif**

```
dn: ou=users,dc=example,dc=com
objectClass: organizationalUnit
ou: users

dn: cn=Jane,ou=users,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
objectClass: testPerson
cn: Jane
sn: Smith
displayName: Jane Smith
mail: jane.smith@example.com
memberOf: cn=admins,ou=groups,dc=example,dc=com 1

dn: cn=Jim,ou=users,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
objectClass: testPerson
cn: Jim
sn: Adams
displayName: Jim Adams
mail: jim.adams@example.com
memberOf: cn=otheradmins,ou=groups,dc=example,dc=com 2

dn: ou=groups,dc=example,dc=com
objectClass: organizationalUnit
ou: groups

dn: cn=admins,ou=groups,dc=example,dc=com 3
objectClass: group
cn: admins
```

```

owner: cn=admin,dc=example,dc=com
description: System Administrators
member: cn=Jane,ou=users,dc=example,dc=com
member: cn=otheradmins,ou=groups,dc=example,dc=com

dn: cn=otheradmins,ou=groups,dc=example,dc=com 4
objectClass: group
cn: otheradmins
owner: cn=admin,dc=example,dc=com
description: Other System Administrators
memberOf: cn=admins,ou=groups,dc=example,dc=com 5 6
member: cn=Jim,ou=users,dc=example,dc=com

```

1 2 5 The user's and group's memberships are listed as attributes on the object.

3 4 The groups are first-class entries on the LDAP server.

6 The **otheradmins** group is a member of the **admins** group.

When syncing nested groups with Active Directory, you must provide an LDAP query definition for both user entries and group entries, as well as the attributes with which to represent them in the internal OpenShift Dedicated group records. Furthermore, certain changes are required in this configuration:

- The **oc adm groups sync** command must explicitly whitelist groups.
- The user's **groupMembershipAttributes** must include "**memberOf:1.2.840.113556.1.4.1941:**" to comply with the **LDAP_MATCHING_RULE_IN_CHAIN** rule.
- The **groupUIDAttribute** must be set to **dn**.
- The **groupsQuery**:
 - Must not set **filter**.
 - Must set a valid **derefAliases**.
 - Should not set **baseDN** as that value is ignored.
 - Should not set **scope** as that value is ignored.

For clarity, the group you create in OpenShift Dedicated should use attributes other than the distinguished name whenever possible for user- or administrator-facing fields. For example, identify the users of an OpenShift Dedicated group by their e-mail, and use the name of the group as the common name. The following configuration file creates these relationships:

LDAP sync configuration that uses augmented Active Directory schema with nested members: **augmented_active_directory_config_nested.yaml**

```

kind: LDAPSyncConfig
apiVersion: v1
url: ldap://LDAP_SERVICE_IP:389
augmentedActiveDirectory:
  groupsQuery: 1
  derefAliases: never
  pageSize: 0


```

```

groupUIDAttribute: dn 2
groupNameAttributes: [ cn ] 3
usersQuery:
  baseDN: "ou=users,dc=example,dc=com"
  scope: sub
  derefAliases: never
  filter: (objectclass=person)
  pageSize: 0
userNameAttributes: [ mail ] 4
groupMembershipAttributes: [ "memberOf:1.2.840.113556.1.4.1941:" ] 5

```

- 1** **groupsQuery** filters cannot be specified. The **groupsQuery** base DN and scope values are ignored. **groupsQuery** must set a valid **derefAliases**.
- 2** The attribute that uniquely identifies a group on the LDAP server. It must be set to **dn**.
- 3** The attribute to use as the name of the group.
- 4** The attribute to use as the name of the user in the OpenShift Dedicated group record.



NOTE

mail or **sAMAccountName** are preferred choices in most installations.

- 5** The attribute on the user that stores the membership information. Note the use of **LDAP_MATCHING_RULE_IN_CHAIN**.

Prerequisites

- Create the configuration file.
- You have access to the cluster as a user with the **dedicated-admin** role.


Procedure

- Run the sync with the **augmented_active_directory_config_nested.yaml** file:

```

$ oc adm groups sync \
  'cn=admins,ou=groups,dc=example,dc=com' \
  --sync-config=augmented_active_directory_config_nested.yaml \
  --confirm

```



NOTE

You must explicitly whitelist the **cn=admins,ou=groups,dc=example,dc=com** group.

OpenShift Dedicated creates the following group record as a result of the above sync operation:

OpenShift Dedicated group created by using the `augmented_active_directory_config_nested.yaml` file

```

apiVersion: user.openshift.io/v1
kind: Group
metadata:
  annotations:
    openshift.io/ldap.sync-time: 2015-10-13T10:08:38-0400 1
    openshift.io/ldap.uid: cn=admins,ou=groups,dc=example,dc=com 2
    openshift.io/ldap.url: LDAP_SERVER_IP:389 3
  creationTimestamp:
    name: admins 4
  users: 5
- jane.smith@example.com
- jim.adams@example.com

```

- 1** The last time this OpenShift Dedicated group was synchronized with the LDAP server, in ISO 6801 format.
- 2** The unique identifier for the group on the LDAP server.
- 3** The IP address and host of the LDAP server where this group's record is stored.
- 4** The name of the group as specified by the sync file.
- 5** The users that are members of the group, named as specified by the sync file. Note that members of nested groups are included since the group membership was flattened by the Microsoft Active Directory Server.

15.5. LDAP SYNC CONFIGURATION SPECIFICATION

The object specification for the configuration file is below. Note that the different schema objects have different fields. For example, `v1.ActiveDirectoryConfig` has no **groupsQuery** field whereas `v1.RFC2307Config` and `v1.AugmentedActiveDirectoryConfig` both do.



IMPORTANT

There is no support for binary attributes. All attribute data coming from the LDAP server must be in the format of a UTF-8 encoded string. For example, never use a binary attribute, such as **objectGUID**, as an ID attribute. You must use string attributes, such as **sAMAccountName** or **userPrincipalName**, instead.

15.5.1. v1.LDAPSyncConfig

LDAPSyncConfig holds the necessary configuration options to define an LDAP group sync.

Name	Description	Schema
------	-------------	--------

Name	Description	Schema
kind	String value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://github.com/kubernetes/community/blob/master/contributors/devel/sig-architecture/api-conventions.md#types-kinds	string
apiVersion	Defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://github.com/kubernetes/community/blob/master/contributors/devel/sig-architecture/api-conventions.md#resources	string
url	Host is the scheme, host and port of the LDAP server to connect to: scheme://host:port	string
bindDN	Optional DN to bind to the LDAP server with.	string
bindPassword	Optional password to bind with during the search phase.	v1.StringSource
insecure	If true , indicates the connection should not use TLS. If false , ldaps:// URLs connect using TLS, and ldap:// URLs are upgraded to a TLS connection using StartTLS as specified in https://tools.ietf.org/html/rfc2830 . If you set insecure to true , you cannot use ldaps:// URL schemes.	boolean
ca	Optional trusted certificate authority bundle to use when making requests to the server. If empty, the default system roots are used.	string

Name	Description	Schema
groupUIDNameMapping	Optional direct mapping of LDAP group UIDs to OpenShift Dedicated group names.	object
rfc2307	Holds the configuration for extracting data from an LDAP server set up in a fashion similar to RFC2307: first-class group and user entries, with group membership determined by a multi-valued attribute on the group entry listing its members.	v1.RFC2307Config
activeDirectory	Holds the configuration for extracting data from an LDAP server set up in a fashion similar to that used in Active Directory: first-class user entries, with group membership determined by a multi-valued attribute on members listing groups they are a member of.	v1.ActiveDirectoryConfig
augmentedActiveDirectory	Holds the configuration for extracting data from an LDAP server set up in a fashion similar to that used in Active Directory as described above, with one addition: first-class group entries exist and are used to hold metadata but not group membership.	v1.AugmentedActiveDirectoryConfig

15.5.2. v1.StringSource

StringSource allows specifying a string inline, or externally via environment variable or file. When it contains only a string value, it marshals to a simple JSON string.

Name	Description	Schema
value	Specifies the cleartext value, or an encrypted value if keyFile is specified.	string
env	Specifies an environment variable containing the cleartext value, or an encrypted value if the keyFile is specified.	string

Name	Description	Schema
file	References a file containing the cleartext value, or an encrypted value if a keyFile is specified.	string
keyFile	References a file containing the key to use to decrypt the value.	string

15.5.3. v1.LDAPQuery

LDAPQuery holds the options necessary to build an LDAP query.

Name	Description	Schema
baseDN	DN of the branch of the directory where all searches should start from.	string
scope	The optional scope of the search. Can be base : only the base object, one : all objects on the base level, sub : the entire subtree. Defaults to sub if not set.	string
derefAliases	The optional behavior of the search with regards to aliases. Can be never : never dereference aliases, search : only dereference in searching, base : only dereference in finding the base object, always : always dereference. Defaults to always if not set.	string
timeout	Holds the limit of time in seconds that any request to the server can remain outstanding before the wait for a response is given up. If this is 0 , no client-side limit is imposed.	integer
filter	A valid LDAP search filter that retrieves all relevant entries from the LDAP server with the base DN.	string

Name	Description	Schema
pageSize	Maximum preferred page size, measured in LDAP entries. A page size of 0 means no paging will be done.	integer

15.5.4. v1.RFC2307Config

RFC2307Config holds the necessary configuration options to define how an LDAP group sync interacts with an LDAP server using the RFC2307 schema.

Name	Description	Schema
groupsQuery	Holds the template for an LDAP query that returns group entries.	v1.LDAPQuery
groupUIDAttribute	Defines which attribute on an LDAP group entry will be interpreted as its unique identifier. (IdapGroupUID)	string
groupNameAttributes	Defines which attributes on an LDAP group entry will be interpreted as its name to use for an OpenShift Dedicated group.	string array
groupMembershipAttributes	Defines which attributes on an LDAP group entry will be interpreted as its members. The values contained in those attributes must be queryable by your UserUIDAttribute .	string array
usersQuery	Holds the template for an LDAP query that returns user entries.	v1.LDAPQuery
userUIDAttribute	Defines which attribute on an LDAP user entry will be interpreted as its unique identifier. It must correspond to values that will be found from the GroupMembershipAttributes .	string

Name	Description	Schema
userNameAttributes	Defines which attributes on an LDAP user entry will be used, in order, as its OpenShift Dedicated user name. The first attribute with a non-empty value is used. This should match your PreferredUsername setting for your LDAPPasswordIdentityProvider . The attribute to use as the name of the user in the OpenShift Dedicated group record. mail or sAMAccountName are preferred choices in most installations.	string array
tolerateMemberNotFoundErrors	Determines the behavior of the LDAP sync job when missing user entries are encountered. If true , an LDAP query for users that does not find any will be tolerated and an only and error will be logged. If false , the LDAP sync job will fail if a query for users does not find any. The default value is false . Misconfigured LDAP sync jobs with this flag set to true can cause group membership to be removed, so it is recommended to use this flag with caution.	boolean
tolerateMemberOutOfScopeErrors	Determines the behavior of the LDAP sync job when out-of-scope user entries are encountered. If true , an LDAP query for a user that falls outside of the base DN given for the all user query will be tolerated and only an error will be logged. If false , the LDAP sync job will fail if a user query would search outside of the base DN specified by the all user query. Misconfigured LDAP sync jobs with this flag set to true can result in groups missing users, so it is recommended to use this flag with caution.	boolean

15.5.5. v1.ActiveDirectoryConfig

ActiveDirectoryConfig holds the necessary configuration options to define how an LDAP group sync interacts with an LDAP server using the Active Directory schema.

Name	Description	Schema
usersQuery	Holds the template for an LDAP query that returns user entries.	v1.LDAPQuery
userNameAttributes	Defines which attributes on an LDAP user entry will be interpreted as its OpenShift Dedicated user name. The attribute to use as the name of the user in the OpenShift Dedicated group record. mail or sAMAccountName are preferred choices in most installations.	string array
groupMembershipAttributes	Defines which attributes on an LDAP user entry will be interpreted as the groups it is a member of.	string array

15.5.6. v1.AugmentedActiveDirectoryConfig

AugmentedActiveDirectoryConfig holds the necessary configuration options to define how an LDAP group sync interacts with an LDAP server using the augmented Active Directory schema.

Name	Description	Schema
usersQuery	Holds the template for an LDAP query that returns user entries.	v1.LDAPQuery
userNameAttributes	Defines which attributes on an LDAP user entry will be interpreted as its OpenShift Dedicated user name. The attribute to use as the name of the user in the OpenShift Dedicated group record. mail or sAMAccountName are preferred choices in most installations.	string array
groupMembershipAttributes	Defines which attributes on an LDAP user entry will be interpreted as the groups it is a member of.	string array

Name	Description	Schema
groupsQuery	Holds the template for an LDAP query that returns group entries.	v1.LDAPQuery
groupUIDAttribute	Defines which attribute on an LDAP group entry will be interpreted as its unique identifier. (ldapGroupUID)	string
groupNameAttributes	Defines which attributes on an LDAP group entry will be interpreted as its name to use for an OpenShift Dedicated group.	string array