



**Hewlett Packard**  
Enterprise

# **HPE Ezmeral Data Fabric 7.9 Documentation**

## **HPE Ezmeral Data Fabric**

# Contents

<b>Get Started</b> .....	<b>7</b>
Release Notes.....	7
What's New in Release 7.9.0.....	7
Known Issues (Release 7.9.0).....	10
Installation.....	18
Fabric Deployment Using a Seed Node.....	18
Prerequisites for On-Premises Installation.....	27
AWS Fabric Configuration Parameters.....	29
Azure Fabric Configuration Parameters.....	31
GCP Fabric Configuration Parameters.....	32
On-Premises Fabric Configuration Parameters.....	33
Troubleshooting Seed Node Installation.....	34
Downloading Installation Logs.....	36
Creating a Local Repository for an Air-Gapped Installation.....	38
Planning Worksheet for Cloud Deployments.....	40
Help for datafabric_container_setup.sh.....	41
Service Activation and Billing.....	42
Adding an Activation Key.....	42
Registering a Fabric.....	43
Setting the Billing Model.....	43
Viewing Activation Information.....	44
Displaying the Fabric ID.....	44
Obtaining a License.....	45
Billing in Connected Environments.....	46
Billing in Air-Gapped Environments.....	46
Restoring a Disabled Fabric.....	49
Displaying a maprccli Prompt.....	49
SSO Using Keycloak.....	50
Accessing the Keycloak Administration Console.....	50
Changing the Keycloak admin Password.....	51
Adding New Users to Keycloak.....	54
Adding a Group to Keycloak.....	60
Integrating Your LDAP Directory with Keycloak.....	61
Using LDAP Mappers.....	62
Completing SSO Setup Using the Data Fabric UI.....	67
Resetting the SSO Configuration.....	69
Identifying All CLDB Nodes.....	69
Setting Up Clients.....	70
Installing Clients on a Linux Host.....	70
Setting up the Data Fabric Repository.....	71
Installing the Data Fabric Client on RHEL.....	75
Installing the Data Fabric Client on SLES.....	76
Installing the Data Fabric Client on Ubuntu.....	76
Installing Client Libraries.....	77
About Access and Refresh Tokens.....	77
Upgrade.....	80
Upgrading a Data Fabric.....	80
Upgrade Fabric Parameters.....	84
User Assistance.....	85

Troubleshooting Online Help for the Data Fabric UI.....	85
<b>Platform.....</b>	<b>86</b>
Data Fabric UI.....	86
Global Namespace (GNS).....	88
S3 Federation in Global Namespace.....	89
S3 Global Namespace.....	90
Network File System in Global Namespace.....	91
Single Sign-On (SSO) Support.....	92
Iceberg Support.....	92
Fabric Resources.....	93
Volumes.....	93
Buckets.....	94
Topics.....	94
Data Storage Management.....	94
Data Tiering.....	94
Cold Tier.....	95
Data Read, Write, and Recall.....	95
Data Storage Policy.....	96
AWS Architecture Notes.....	96
Azure Architecture Notes.....	99
HPE Ezmeral Unified Analytics.....	102
GCP Architecture Notes.....	102
<b>Administration.....</b>	<b>104</b>
IPv6 Support in Data Fabric.....	104
Enabling IPv6.....	105
Administering Fabrics.....	107
Configuring a Proxy Server for Data Fabric Access to the Internet.....	108
Creating a Fabric.....	109
Importing a Fabric.....	110
Importing an as-a-Service Fabric.....	110
Viewing the Fabric Status.....	112
Viewing Fabric Settings.....	113
Viewing the Fabric Endpoint.....	114
Viewing the Software Version.....	115
Generating S3 Access Keys for the Global Namespace.....	116
Setting a Quota for a User.....	118
Setting a Quota for a Group.....	119
Viewing Fabric-Related Metrics.....	120
View Storage Consumption by User.....	120
View System Resource Utilization by Fabric.....	121
View Top Fabrics by Storage Capacity.....	122
View Billing Data by Fabric.....	123
Setting Default Quotas for Users/Groups.....	124
Viewing the Fabric Service Status.....	125
View Capacity Usage by User on Fabric.....	125
SSH Access to a Cloud-Based Fabric.....	126
Deleting a Fabric.....	127
Administering Identities.....	127
About Roles.....	128
Pre-defined Roles and Associated Permissions.....	129
Creating a Role.....	131
Viewing Roles.....	132

Editing a Role.....	132
Deleting a Role.....	133
Assigning a Role to a User.....	134
Assigning a Role to a Group.....	134
Viewing a List of Users.....	135
Viewing a List of Groups.....	135
Administering IAM Policies.....	136
About IAM Policy.....	136
Identity Access Management Policy Life Cycle.....	137
Resource-level Permissions in an IAM Policy.....	137
Creating an IAM Policy.....	139
Editing an IAM Policy.....	140
Deleting an IAM Policy.....	141
Viewing All IAM Policies.....	142
Assigning an IAM Policy.....	142
Configuring Email Notifications.....	143
Viewing and Editing Access Control Information.....	144
Access Control Expression Syntax.....	145
Administering Buckets.....	146
Creating a Bucket.....	147
Creating a Folder on a Bucket.....	148
Uploading Objects to a Bucket.....	148
Downloading an Object from a Bucket.....	149
Deleting an Object from a Bucket.....	150
Deleting a Folder from a Bucket.....	151
Deleting a Bucket.....	151
Administering Tables.....	152
Managing Tables.....	152
Creating a Table.....	152
Deleting a Table.....	153
Managing Column Families and Columns.....	153
Creating a Column Family.....	153
Configuring Column Family Permissions.....	155
Deleting a Column Family.....	156
Viewing Table Information.....	156
Viewing the List of Tables.....	156
Viewing Column Families.....	157
Managing Table Replication.....	158
Adding a Table Replica.....	158
Viewing Table Replicas.....	159
Administering Access Controls for Tables.....	159
Administering Topics.....	160
Creating a Topic.....	160
Editing a Topic.....	161
Deleting a Topic.....	162
Viewing or Downloading Topic Connection Properties.....	163
Administering Volumes.....	164
Creating a Standard Volume.....	164
Creating a Mirror Volume.....	166
Converting Standard Volume to Mirror Volume.....	167
Editing a Volume.....	168
Setting a Volume Quota.....	168
Configuring Data Access Control for Volume.....	169
Configuring Volume Administration Settings.....	170
Renaming a Volume.....	171
Viewing Volume Endpoint Info.....	172

Viewing Object Endpoint Info to Remotely Access Files as Objects.....	172
Downloading Volume Endpoint Information.....	173
Deleting a Volume.....	174
Administering Volume Snapshots.....	174
Schedules for Volume Snapshots.....	175
Creating a Volume Snapshot.....	176
Scheduling Volume Snapshots.....	177
Preserving a Volume Snapshot.....	178
Restoring a Volume from Volume Snapshot.....	179
Deleting a Volume Snapshot.....	179
Data Tiering.....	180
Schedules for Volume Data Tiering.....	181
Manually Offloading Data to a Cold Tier.....	182
Recalling Data to the Data Fabric File System.....	183
Administering Storage Policies.....	184
Administering Remote Targets.....	189
Administering Schedules.....	192
Mirroring.....	195
Local Mirroring.....	195
Remote Mirroring.....	196
Starting Volume Mirroring.....	196
Stopping Volume Mirroring.....	197
Scheduling Volume Mirroring.....	197
Administering Nodes.....	198
Viewing Node Information.....	198
Adding Nodes (On-premises Deployment).....	198
Removing Nodes (On-premises Deployment).....	201
Adding Nodes (Cloud Deployment).....	202
Auditing Fabric and Fabric Data.....	205
Enabling/Disabling Fabric Auditing.....	205
Configuring Auditing for Data Access Operation.....	206
Administering Security Policies.....	206
About Security Policy Domain.....	209
Security Policy Implementation Workflow.....	210
Security Policy Enforcement Process.....	212
Understanding Access Control in a Security Policy.....	215
Managing File and Directory ACEs.....	217
Security Policy Permissions.....	218
Designating a Fabric as Global Policy Master.....	222
Creating a Security Policy.....	222
Viewing a Security Policy.....	223
Viewing All Security Policies.....	224
Editing a Security Policy.....	224
Assigning a Security Policy to One or More Volumes.....	225
Assigning Multiple Security Policies to One or More Volumes.....	225
Unassigning One or More Security Policies from a Volume.....	226
Disabling a Security Policy.....	227
Enabling a Security Policy.....	228
Administering Bucket Policies.....	228
Creating a Bucket Policy using JSON.....	229
Sample Bucket Policy.....	230
Creating a Bucket Policy using Policy Generator.....	231
Sample Bucket Policy using Policy Builder.....	232
Editing a Bucket Policy.....	235
Delete a Bucket Policy from a Bucket.....	236
Working with an External NFS Server.....	236

- Importing an External Network File System Server..... 236
- Viewing the IP Address/Hostname for External NFS Server.....237
- Deleting an External NFS Server..... 237
- Working with an External S3 Object Store.....238
  - Importing an External S3 Object Store..... 238
  - Deleting an External S3 Object Store..... 241
  - Viewing Object Store Details.....242
  - Sharing External S3 Object Store with User or Group..... 242
  - Removing Share for External S3 Object Store.....243
- Integrating the AWS Security Token Service (STS) with Data Fabric..... 243
  - Configuring STS for Data Fabric..... 245
- Administering Alarms..... 250
  - Viewing Alarms..... 250
  - Muting/Dismissing Alarms.....251
- Monitoring..... 252
  - Adding an OTel Endpoint.....252
- Getting Started with Iceberg..... 252
- Configuring Data Fabric to Track User Behavior..... 253

**Reference.....255**

- Release History.....255
- Cloud Instance Specifications..... 256
- Third-Party Storage Solutions..... 257
- Port Information.....258
- maprcli Commands in This Guide..... 258
- Operating System Support Matrix .....259
- Container Image Vulnerabilities and CVE Reports..... 260
- Doc Site Available as a PDF..... 261
- Product Licensing..... 261
  - Additional License Authorizations (ALA)..... 261
  - Open-Source Software Acknowledgements (Release 7.9.x)..... 261
- Other Resources..... 338
- Contact HPE..... 338

**Glossary.....338**

- access policy..... 339
- Domain..... 342
- domain user..... 343
- IAM users..... 344
- MOSS..... 345
- object..... 345
- Object Store..... 345

# Get Started

This section describes how you can get started learning about, installing, and using the HPE Ezmeral Data Fabric.

## HPE Ezmeral Data Fabric 7.9.0 Release Notes

These notes contain information about release 7.9.0 of the HPE Ezmeral Data Fabric as-a-service platform.

### What's New in Release 7.9.0

HPE Ezmeral Data Fabric 7.9.0 is a data-storage platform from Hewlett Packard Enterprise that offers a multimodal, subscription-based, hybrid-cloud experience for the enterprise.

**!** **IMPORTANT:** The predecessor of the HPE Ezmeral Data Fabric is the HPE Ezmeral Data Fabric – Customer Managed platform. See [Comparing the Data Fabric Platforms](#) on page 9. To view information for the customer-managed Data Fabric platform, see [this website](#).

### New Features for Release 7.9.0

Following are some new features and capabilities that distinguish release 7.9.0 from the previous release (7.8.0):

New Feature or Capability	Supported on		See for more information . . .
	DF SaaS?	Customer Managed?	
Support for user-defined roles	Yes	No	<a href="#">Administering Identities</a> on page 127** <a href="#">acl*</a>
Runbook for Data Fabric cloud configurations	Yes	No	Coming soon
Support for C API Librdkafka for HPE Ezmeral Data Fabric Streams	No	Yes	<a href="#">HPE Ezmeral Data Fabric Streams C Applications*</a> <a href="#">HPE Ezmeral Data Fabric Streams C Client 2.0.2 Release Notes*</a>
Updates to Insight services	Yes	Yes	<a href="#">Configuring Data Fabric to Track User Behavior</a> on page 253** <a href="#">insight*</a> <a href="#">setloglevel insight*</a>
Support for RHEL 9.4	Yes	Yes	<a href="#">Operating System Support Matrix</a> on page 259 <a href="#">Operating System Support Matrix*</a>
Remove nodes from a fabric	Yes	No	<a href="#">Removing Nodes (On-premises Deployment)</a> on page 201** <a href="#">installer clusterremovenode*</a>
Identity access management	Yes	No	<a href="#">Administering IAM Policies</a> on page 136**

\*Indicates a link to the HPE Ezmeral Data Fabric – Customer Managed documentation.

\*\*Requires using the Data Fabric UI. With release 7.3.0 and later, you can use the Data Fabric UI on customer-managed clusters. To understand the limitations and benefits of doing so, see [Data Fabric UI](#).

### **Welcome to the HPE Ezmeral Data Fabric**

Release 7.9.0 introduces an HPE Ezmeral Data Fabric that is:

- Managed as a service
- Subscription-based
- Designed to make data accessible

The single logical view integrates files, objects, and streaming data and features consumption-based pricing. A [global namespace](#) enables deployments to join a single fabric no matter where the data is located.

### **Managed As a Service**

With the HPE Ezmeral Data Fabric platform, HPE manages the configuration, upgrade, and lifecycle of the platform – as long as the platform is connected.

### **Deploy to the Cloud or On-Premises**

Release 7.8.0 of the HPE Ezmeral Data Fabric provides a high-performance file system for files, binary tables, objects, and streaming files that can be deployed quickly on:

- Amazon Web Services (AWS)
- Microsoft Azure
- Google Cloud Platform (GCP)
- Your on-premises infrastructure

The Data Fabric simplifies how you manage petabyte-scale data, enabling on-premises files to be transferred as native S3 objects to the cloud, where upstream apps can transform the data. Results can be stored either on-premises or on the cloud as objects. The Data Fabric makes it easy to migrate workloads across clouds or on-premise installations without the need for format changes.

### **Storage Tiers and Consumption**

With the HPE Ezmeral Data Fabric, the amount of storage you use determines your monthly charges for the service. HPE offers licenses for the following storage tiers:

- 1 TB
- 10 TB
- 100 TB
- 1 PB

The storage tier is a baseline for consumption. You select a storage tier when you purchase a license and confirm the tier when you create a fabric. You can consume more storage than your specified storage tier, but rates are higher when your consumption exceeds the specified tier. For current rate information, contact your HPE sales representative.

The [Data Fabric UI](#) on page 86 enables you to monitor your consumption and provides estimates of your monthly charges.



## Comparing the Data Fabric Platforms

The as-a-service HPE Ezmeral Data Fabric leverages the strengths of its predecessor, the HPE Ezmeral Data Fabric – Customer Managed platform. The as-a-service platform also improves on its predecessor in many ways. The following table compares the platforms:

Feature	HPE Ezmeral Data Fabric	HPE Ezmeral Data Fabric – Customer Managed
Distributed File System	Yes	Yes
Global Namespace (GNS)	Yes	No
Object Support	Yes	Yes
Table Support	Yes	Yes
Event Stream Support	Yes	Yes
NFSv4 Support*	Yes	Yes
Container Storage Interface (CSI) Support	Yes	Yes
Database Support	Yes	Yes
Client Support	Yes	Yes
Single Sign-On (SSO) Support	Yes	Yes (release 7.3.0 and later)
Rolling Upgrades	Managed by HPE	User managed
Billing and Licensing	Consumption-only and some form of term (term-only not supported)	Consumption and Term
Air-Gap Support	Yes	Yes
Graphical User Interface	Data Fabric UI or Control System	Data Fabric UI or Control System
maprcli Command Line	Yes	Yes
Scale (number of nodes per fabric / cluster)	See note**	Thousands of nodes
EEP (HPE Ezmeral Ecosystem Pack)	No	Yes
OpenTelemetry (OTel)	Yes	Yes

\*NFSv3 is not supported.

\*\*For cloud deployments, the nodes and instances are predetermined based on the storage tier that you select during installation. While new fabrics can be added at any time, adding nodes after fabric creation is not currently supported for cloud deployments. For on-premises deployments, you determine the number of nodes at create time. You can add nodes later by using the steps in [Adding Nodes \(On-premises Deployment\)](#) on page 198.

## Terminology


Some terms used to describe the new HPE Ezmeral Data Fabric are different from terms used to describe the HPE Ezmeral Data Fabric – Customer Managed platform. The following table highlights some terminology differences:

Term	Definition	Comparable HPE Ezmeral Data Fabric – Customer Managed Term
Data Fabric UI	The new user interface for managing the HPE Ezmeral Data Fabric. See <a href="#">Data Fabric UI</a> on page 86.	Control System

Term	Definition	Comparable HPE Ezmeral Data Fabric – Customer Managed Term
fabric	A software-defined storage system that provides multi-modal access to data. Fabrics help you manage your data, making it possible to access, integrate, model, analyze, and provision your data seamlessly. Multiple fabrics share a common global namespace.	cluster
global namespace	A data plane that spans all of your Data Fabric deployments and includes all of the as-a-service fabrics in your enterprise. See <a href="#">Global Namespace (GNS)</a> on page 88.	cross-cluster security
HPE Ezmeral Data Fabric	Now refers to the consumption-based, as-a-service Data Fabric platform. Data Fabric deployments that are not consumption based and not provided as-a-service are now referred to as "customer managed."	HPE Ezmeral Data Fabric – Customer Managed

## Known Issues (Release 7.9.0)

You might encounter the following known issues after upgrading to release 7.9.0. This list is current as of the release date.

 **IMPORTANT:** The "Support notices of known issues" tool is no longer available, but you can obtain the same information by logging on to the [HPE Support Center](#).

Where available, the workaround for an issue is also documented. HPE regularly releases maintenance releases and patches to fix issues. We recommend checking the release notes for any subsequent maintenance releases to see if one or more of these issues are fixed.

### HPE Ezmeral Data Fabric Streams

#### MS-1511

When messages are produced for the second time on a stream, CopyStream fails with an exception `com.mapr.db.exceptions.DBException: flush() failed with err code = 22`.

**Workaround:** None.

### Client Libraries

#### MFS-20211

User unable to use standalone mapr-client on Ubuntu machines.

**Workaround:** Install the `libcurl3-gnutls` package before using mapr-client on Ubuntu machines.

#### MFS-18258

When you add a new cluster to a cluster group, the FUSE-based POSIX client and the loopbacknfs POSIX client take about five minutes to load or list the newly added cluster.

**Workaround:** None.

### Data Fabric UI

#### Node Removal

#### DFUI-2751

Data Fabric UI hangs when there is removal of node operation is repeated by adding and removing the same node repeatedly.

**Workaround:** Use the `maprccli` for this operation.

**Sign-in Issues****DFUI-2743**

SSO (Keycloak SSO) user unable to log in to Data Fabric UI when the user is assigned a pre-defined role from Data Fabric UI.

**Workaround:** Assign the pre-defined role using the Keycloak console.

**DFUI-2734**

A user that is assigned a user-defined role is unable to log in to Data Fabric UI. The 'no login permission' error is displayed when such a user attempts to log in to Data Fabric UI.

**Workaround:** None

**DFUI-2701**

Even after being assigned full control permission on all fabrics, the user is unable to log in to a non-primary fabric using Data Fabric UI

**Workaround:** None.

**DFUI-160**

If you sign in to the Data Fabric UI as an SSO user but you do not have fabric-level login permission, a sign-in page for the "Managed Control System" (MCS) is displayed. The "Managed Control System" sign-in is not usable for the consumption-based HPE Ezmeral Data Fabric.

**Workaround:** Use one of the following workarounds:

- Edit the MCS URL, and retry logging in. For example, change the boldface characters in the following URL:

```
https://
<host-name>:8443/app/mcs/#/app/
overview
```

To this:

```
https://<host-name>:8443/app/dfui
```

- Try signing in as a user who has fabric-level login permission.
- Dismiss the MCS page, clear your browser cache, and retry signing in.

**DFUI-437**

If you sign in to the Data Fabric UI as a non-SSO user and then sign out and try to sign in as an SSO user, a sign-in page for the "Managed Control System" (MCS) is displayed. The "Managed Control System" sign-in is not usable for the consumption-based HPE Ezmeral Data Fabric.

**Workaround:** Use one of the following workarounds:

- Edit the MCS URL, and retry logging in. For example, change the boldface characters in the following URL:

```
https://
<host-name>:8443/app/mcs/#/app/
overview
```

To this:

```
https://<host-name>:8443/app/dfui
```

- Dismiss the "Managed Control System" sign-in screen, and retry signing in as a non-SSO user.
- Dismiss the MCS page, clear your browser cache, and retry signing in.

#### DFUI-811

If you launch the Data Fabric UI and then sign out and wait for 5-10 minutes and then attempt to sign in, a sign-in page for the "Managed Control System" (MCS) is displayed.

**Workaround:** See the workaround for DFUI-437.

#### DFUI-826

In a cloud fabric, an empty page is displayed after a session expires and you subsequently click on a fabric name. The browser can display the following URL:

```
https://<hostname>:8443/oath/login
```

**Workaround:** None.

#### DFUI-874

Sometimes when you attempt to sign in to the Data Fabric UI, the "Managed Control System" (MCS) is displayed, or the Object Store UI is displayed.

**Workaround:** See the workaround for DFUI-437.

#### DFUI-897

A user with no assigned role cannot sign in to the Data Fabric UI.

**Workaround:** Using your SSO provider software, assign a role to the user, and retry the sign-in operation.

#### DFUI-1123

Attempting to sign in to the Data Fabric UI as a group results in a login error message in the browser. For example:

```
https://<hostname>:8443/login?error
```

**Workaround:** None.

### Mirroring Issues

#### MFS-17538

During PBS validation, a primary cluster with automatic mirroring of a PBS volume to a non-primary cluster might give the following error:

```
Failed to fetch fabric cluster-151-B
401 Unauthorized: "HTTP ERROR 401 JWT
validation failed: null<EOL>URI:
/rest/dashboard/info/<EOL>STATUS:
```

```
401<EOL>MESSAGE:
JWT validation failed:
null<EOL>SERVLET: mapr-apiserver<EOL>"
```

**Workaround:** Restart the API server.

#### DFUI-1227

If you create a mirror volume with a security policy, an error is generated when you try to remove the security policy.

**Workaround:** None.

#### DFUI-1229

Data aces on a mirror volume cannot be edited.

**Workaround:** None.

### Display Issues

#### DFUI-2691

Existing user-defined roles are not visible on the Data Fabric UI when you wish to assign roles to users.

**Workaround:** Use the respective `maprccli` command to perform the operation.

#### DFUI-2708

A user that is provided admin permissions by way of IAM policy assignment is unable to log in to the Data Fabric UI.

**Workaround:** Use the `maprccli` to login and perform any actions related to Data Fabric.

#### DFUI-2719

Permissions assigned a user by way of assigning IAM policy to user are not visible on the Data Fabric UI.

**Workaround:** None.

#### DFUI-2703

A user is unable to edit cluster settings and ACLs though the user has been granted permissions equivalent to fabric manager by way of assigning an with IAM policy with fabric management actions.

**Workaround:** None.

#### DFUI-2749

User-defined roles assigned to a user are not reflecting on the Data Fabric UI.

**Workaround:** Use the `security iam role mapping maprccli` command to view the user-defined role assigned to the user.

#### DFUI-1186

After you complete the SSO setup for a new fabric, fabric resources such as volumes and mirrors are not immediately displayed in the Data Fabric UI.

**Workaround:** Wait at least 20 minutes or more for the Data Fabric UI to display the fabric details.

#### DFUI-1221

If a fabric includes a large number of resources, loading the resources to display in the **Resources** card on the home page can take a long time.

**Workaround:** None.

#### DFUI-2102

When you create a table replica on a primary cluster with the source table on a secondary cluster, the replication operation times out. However, the table replica is successfully created on the primary cluster. The table replica appears in the **Replication** tab, but does not appear in the Data Fabric UI **Graph** or **Table** view for the primary cluster.

This behavior is the same for both a source table on the primary cluster and the replica on the secondary cluster.

**Workaround:** None.

## External S3

### MFS-20148

A Keycloak user is unable connect to an external S3 server with the access key and secret key, by using an S3 client.

**Workaround:** None.

### DFUI-2157

Editing buckets on external S3 servers is not supported.

**Workaround:** None.

## Installation or Fabric Creation

### MFS-18972

RHEL-based default keycloak that is shipped with Data Fabric cannot be used to configure STS.

**Workaround:** Set up external Keycloak to run on port 443 for STS to work.

### MFS-18734

Release 7.7.0 of the HPE Ezmeral Data Fabric has a dependency on the `libssl1.1` package, which is not included in Ubuntu 22.04. As a result, you must apply the package manually to Ubuntu 22.04 nodes before installing Data Fabric software.

**Workaround:** On every node in the fabric or cluster:



**NOTE:** The following steps are required for cluster nodes but are not required for client nodes.

1. Download the `libssl1.1` package:

```
wget http://archive.ubuntu.com/ubuntu/pool/main/o/openssl/libssl1.1_1.1.0g-2ubuntu4_amd64.deb
```

2. Use the following command to install the package:

```
sudo dpkg -i libssl1.1_1.1.0g-2ubuntu4_amd64.deb
```

### IN-3482

Fabric creation can fail if host-name resolution takes more than 300 ms.

**Workaround:** Check your host-name resolution time, and take steps to improve it. See [Troubleshoot Fabric Creation](#). Then retry fabric deployment.

### DFUI-565, EZINDEFAAS-169

Installation or fabric creation can fail if a proxy is used for internet traffic with the HPE Ezmeral Data Fabric.

**Workaround:** Export the following proxy settings, and retry the operation:

```
# cat /etc/environment
export http_proxy=http://
<proxy_server_hostname_or_IP>:<proxy_p
ort>
export https_proxy=http://
<proxy_server_hostname_or_IP>:<proxy_p
ort>
export HTTP_PROXY=http://
<proxy_server_hostname_or_IP>:<proxy_p
ort>
export HTTPS_PROXY=http://
<proxy_server_hostname_or_IP>:<proxy_p
ort>
```

## Object Store

### DFUI-519

An SSO user is unable to create buckets on the Data Fabric UI and the Object Store. This is applicable to an SSO user with any role such as infrastructure administrator, fabric manager or developer.

**Workaround:** Create an IAM policy with all permissions in the user account. This has to be done via minIO client or the Object Store UI. Assign the IAM policy to the SSO user. Login to the Data Fabric UI and create a bucket/view bucket.

### DFUI-577

Downloading a large file (1 GB or larger) can fail with the following error:

```
Unable to download file "<filename>":
Request failed with status code 500
```

**Workaround:** Instead of using the Data Fabric UI to download a large file, use a MinIO Client (mc) command. For more information about mc commands, see [MinIO Client \(mc\) Commands](#).

## Online Help

### DFUI-459

If a proxy is used for internet traffic with the HPE Ezmeral Data Fabric, online help screens can time out or fail to fetch help content.

**Workaround:** Add the following proxy servers to the `/opt/mapr/apiserver/conf/properties.cfg` file:

- `http.proxy=<proxyServer>:<proxyPort>`
- `https.proxy=<proxyServer>:<proxyPort>`

## Security Policies

### DFUI-2736

A fabric user is able to create a security policy from the command line but unable to create security policy from the Data Fabric UI.

**Workaround:** None.

**MFS-18154/EZINDFAAS-674**

A security policy created on a cloud-based primary fabric (such as AWS) is not replicated on to a secondary fabric created on another cloud provider (such as GCP).

**Workaround:** None.

**Topics****DFUI-637**

Non-LDAP SSO user authenticating to Keycloak cannot create topic on the Data Fabric UI.

**Workaround:** None.

**DFUI-639**

A non-LDAP SSO user authenticating to Keycloak cannot create a volume or stream using the Data Fabric UI.

**Workaround:** None. Non-LDAP and SSO local users are not currently supported.

**Upgrade****COMSECURE-615**

Upgrading directly from release 6.1.x to release 7.x.x can fail because the upgrade process reads password information from the default Hadoop `ssl-server.xml` and `ssl-client.xml` files rather than the original `.xml` files. Note that upgrades from release 6.2.0 to 7.x.x are not affected by this issue.

The issue does not occur, and the upgrade succeeds, if either of the following conditions is true:

- The existing password is `mapr123` (the default value) when the EEP upgrade is initiated.
- You upgrade the cluster first to release 6.2.0 and then subsequently to release 7.x.x.

**Understanding the Upgrade Process and**

**Workaround:** The workaround in this section modifies the release 6.1.x-to-7.x.x upgrade so that it works like the 6.2.0-to-7.x.x upgrade.

Upgrading to core 7.x.x requires installing the `mapr-hadoop-util` package. Before the upgrade, Hadoop files are stored in a subdirectory such as `hadoop-2.7.0`. Installation of the `mapr-hadoop-util` package:

- Creates a subdirectory to preserve the original `.xml` files. This subdirectory has the same name as the original Hadoop directory and a timestamp suffix (for example, `hadoop-2.7.0.20210324131839.GA`).
- Creates a subdirectory for the new Hadoop version (`hadoop-2.7.6`).
- Deletes the original `hadoop-2.7.0` directory.

During the upgrade, a special file called `/opt/mapr/hadoop/prior_hadoop_dir` needs to be created to store the location of the prior Hadoop directory. The `configure.sh` script uses this location to copy the `ssl-server.xml` and `ssl-client.xml` files to the new `hadoop-2.7.6` subdirectory.



In a release 6.1.x-to-7.x.x upgrade, the `prior_hadoop_dir` file does not get created, and `configure.sh` uses the default `ssl-server.xml` and `ssl-client.xml` files provided with Hadoop 2.7.6. In this scenario, any customization in the original `.xml` files is not applied.

The following workaround restores the missing `prior_hadoop_dir` file. With the file restored, `configure.sh -R` consumes the `prior_hadoop_dir` file and copies the original `ssl-server.xml` and `ssl-client.xml` files into the `hadoop-2.7.6` directory, replacing the files that contain the default `mapr123` password.

**Workaround:** After upgrading the ecosystem packages, *but before running* `configure.sh -R`:

1. Create a file named `prior_hadoop_dir` that contains the Hadoop directory path. For example:

```
# cat /opt/mapr/hadoop/
prior_hadoop_dir
/opt/mapr/hadoop/
hadoop-2.7.0.20210324131839.GA
```

If multiple directories are present, specify the directory with the most recent timestamp.

2. Run the `configure.sh -R` command as instructed to complete the EEP upgrade.

## EZINFAAS-811

Upgrading from release 7.6.1 to 7.7.0 fails if you initiate the upgrade from a Data Fabric UI URL that is not the URL provided by the seed node when you created the fabric. The seed node indicates the API server node that is the primary installer host.

**Workaround:** Use either of the following workarounds:

- Initiate the upgrade from the Data Fabric UI URL provided by the seed node when the fabric was created. This URL uses the API server node with the running installer service.
- If you must use an API server node other than the primary installer host:
  1. Copy the `.pem` file from the `/infrastructure/terraform/` directory of the primary installer host to the `/tmp` directory of the secondary installer host where you want to initiate the upgrade.
  2. Restart the installer service on the secondary installer host:

```
sudo service mapr-installer
restart
```

3. Initiate the upgrade as described in [Upgrading a Data Fabric](#).

**MFS-17624**

An upgrade from release 7.5.0 or earlier to 7.6.0 or later can terminate with a fatal error detected by the Java Runtime Environment.

**Workaround:** None.

**OTSDB-147**

After upgrading OpenTSDB from version 2.4.0 to version 2.4.1, the Crontab on each OpenTSDB node is not updated and continues to point to the previous OpenTSDB version.

**Workaround:** To fix the Crontab, run the following commands on each OpenTSDB node, replacing \$MAPR\_USER with the name of the cluster admin (typically mapr) :

- **RHEL**

```
export CRONTAB="/var/spool/cron/
$MAPR_USER"
sed -i 's/2.4.0/2.4.1/' $CRONTAB
```

- **SLES**

```
export CRONTAB="/var/spool/cron/
tabs/$MAPR_USER"
sed -i 's/2.4.0/2.4.1/' $CRONTAB
```

- **Ubuntu**

```
export CRONTAB="/var/spool/cron/
crontabs/$MAPR_USER"
sed -i 's/2.4.0/2.4.1/' $CRONTAB
```

**DFUI-2163**

SSO authentication is not enabled for Data Fabric UI, after upgrading from HPE Ezmeral Data Fabric release version 7.5 to release version 7.6.

**Workaround:** Restart the API server after upgrade.

**Volumes****DFUI-638**

Non-LDAP SSO user authenticating to Keycloak cannot create volume on the Data Fabric UI.

**Workaround:** Create a volume via the Data Fabric minIO client.

## Installation

---

This section contains information about installing the HPE Ezmeral Data Fabric as-a-service platform.



**IMPORTANT:** To install the HPE Ezmeral Data Fabric – Customer Managed platform, see [this website](#).

### Fabric Deployment Using a Seed Node

Describes how to install the platform using a seed node and the Create Fabric interface.

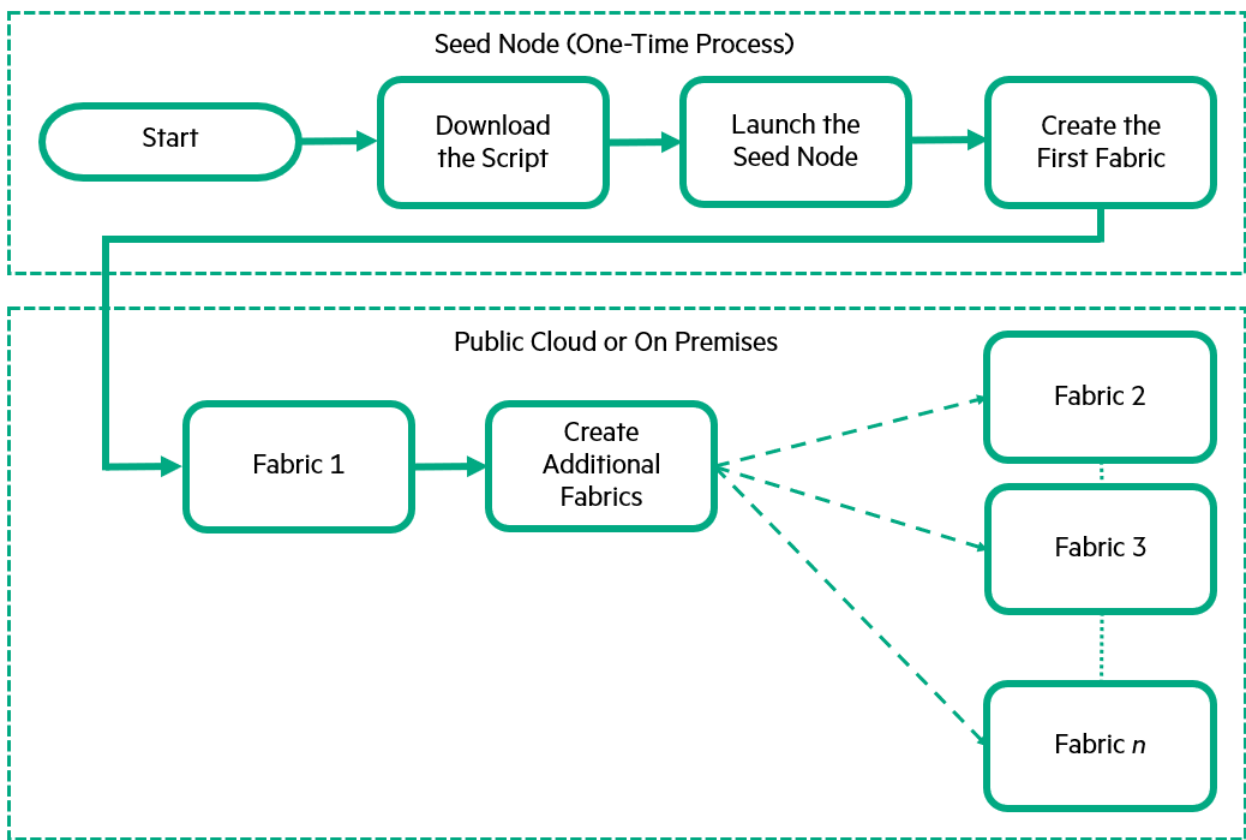
The first fabric you deploy – whether it is in the cloud or on-premises – must be deployed using a seed node. The seed node creates a lightweight, temporary fabric. The installation sequence uses this temporary fabric to display the interface for fabric creation.

The seed node is needed only for creation of the first fabric. Any additional fabric must be created from your cloud-based or on-premises fabric. The seed node can only be used to create fabrics. Creating volumes, buckets, or other resources from the seed node is not supported. For information about creating additional fabrics, see [Creating a Fabric](#) on page 109.

### Seed Node Deployment Process

To deploy the HPE Ezmeral Data Fabric in your environment, you run a script that starts a Docker container. The Docker container emulates the behavior of a Data Fabric node. This emulated node is the seed node. The seed node enables access to the **Create Fabric** interface used to create fabrics.

Once the seed node is created and started, the system prints a URL that you use to access the **Create Fabric** interface. The interface enables you to complete the steps required to deploy a fabric in the cloud or on-premises. After the fabric is created, you can spawn additional fabrics from any installed fabric:



### Prerequisites for Fabric Deployment

Before you deploy a fabric, review all of the following requirements:

#### Seed Node Prerequisites

Verify that the node you plan to use as the seed node meets the following prerequisites:

Prerequisite	Notes
Users	To install using a seed node, you must be <code>root</code> or a user that can run <code>sudo</code> commands without being prompted for a password.

Prerequisite	Notes
OS	<p>The seed node has been tested on the following operating systems, but it can work on other operating systems and in other environments where Docker containers are hosted:</p> <ul style="list-style-type: none"> <li>• Mac OS 10</li> <li>• Ubuntu 20.04</li> </ul>
Connectivity	<p>The seed node that you use to host the Docker container can be a server or laptop that supports the following, but must have connectivity to all the subnets and the VPN for the cloud provider:</p> <ul style="list-style-type: none"> <li>• Docker*</li> <li>• Bash</li> <li>• SSH</li> </ul> <p>*On Ubuntu, RHEL, and Fedora hosts where Docker is not present, Docker is installed automatically. Installation of Docker can take 5-7 minutes. On Mac hosts where Docker is not present, you must install Docker before proceeding with the seed-node deployment.</p>
CPU	64-bit x86 with a minimum of 16 cores per node.
Memory	<p>On the seed node, enough memory must be allocated to Docker to enable the container to come up and run. Docker must be installed and have the following memory allocated to it:</p> <ul style="list-style-type: none"> <li>• Mac: At least 32 GB</li> <li>• Other platforms: At least 8 GB</li> </ul> <p>This is a general recommendation. Sometimes the container can run with less memory. On a seed node with many containers competing for memory, this recommendation might not be sufficient.</p>
Disk Space	If you use a Linux server for the seed node, allocate at least 50 GB of disk space to run the container.
Proxy	<p>If the seed node is behind a proxy, update the proxy configuration in <code>/etc/environment</code> using the following commands. You must do this before attempting to create a fabric:</p> <pre>export http_proxy=&lt;http://myproxy.net&gt; export https_proxy=&lt;http://myproxy.net&gt; export HTTP_PROXY=&lt;http://myproxy.net&gt; export HTTPS_PROXY=&lt;http://myproxy.net&gt;</pre> <p>To specify the proxy information in a file when you run the setup script, see <a href="#">Help for datafabric_container_setup.sh</a> on page 41.</p>

### On-Premises Deployment Prerequisites

See [Prerequisites for On-Premises Installation](#) on page 27.

### Cloud Provider Prerequisites

You must have sufficient permissions to perform tasks in the cloud environment you are using. Hewlett Packard Enterprise recommends the following minimum permissions for installers:

Cloud Provider	Minimum Permissions
AWS	<a href="#">AmazonEBSCSIDriverPolicy</a> and <a href="#">AmazonEC2FullAccess</a>
Azure	<a href="#">Contributor role</a>
GCP	<a href="#">Editor role</a>

Before starting the deployment, gather the information that you will need to fill out the fabric-creation form for your cloud service provider. You can record the necessary information in the [Planning Worksheet for Cloud Deployments](#) on page 40.

### Docker Image Prerequisites

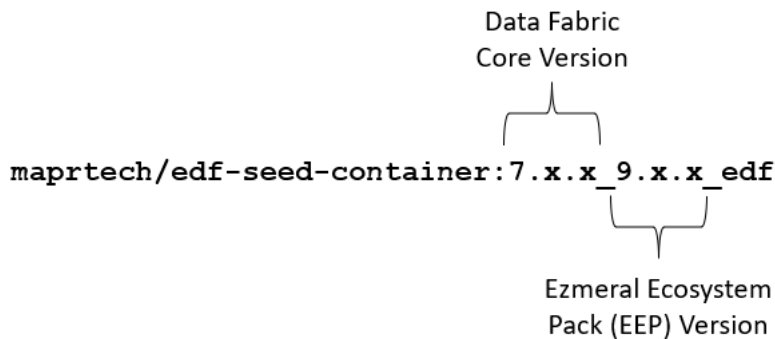
The Docker image that must be used for HPE Ezmeral Data Fabric deployments is:

```
maprtech/edf-seed-container:7.9.0_9.3.1_edf
```

Users can also specify the `maprtech/edf-seed-container:latest` tag, which is always the latest version of the image.

**!** **IMPORTANT:** Only an image that has the `_edf` suffix in the image tag can be used to install the consumption-based, software-as-a-service platform.

The Docker image name includes the Data Fabric core and ecosystem pack version information:



You can download the image from the Docker hub at:

```
https://hub.docker.com/r/maprtech/edf-seed-container/tags/
```

Or you can use the optional `docker pull` command as described in the following steps.

After you verify that your system and environment meet the listed prerequisites, complete the following steps to deploy the HPE Ezmeral Data Fabric.

### Run the Script to Bring up the Container Image on the Seed Node

Use the following steps to bring up the container image and launch the **Create Fabric** interface:

1. Sign in to the seed node host as `root` or a user that can run `sudo` commands without being prompted for a password.
2. Download the `datafabric_container_setup.sh` script from GitHub. For example, download the script in its raw form by using the following `wget` command:

```
wget https://raw.githubusercontent.com/mapr-demos/edf-seednode-790-getting-started/main/datafabric_container_setup.sh
```

3. **Optional:** Use a `docker pull` command to pre-download a copy of the image from <https://hub.docker.com/r/maprtech/edf-seed-container/tags>. Using `docker pull` requires `sudo` privileges:

```
docker pull maprtech/edf-seed-container:latest
```

**IMPORTANT:** Note these considerations:

- Pre-downloading is optional, but it makes the script run faster and prevents download issues when you run the script. The script checks to see if the image is already present on your system. If the image is present, the script uses the image. If it is not present, the script tries to download it.
- If you want to use the `-i <image>` option, you must specify:

```
maprtech/edf-seed-container:7.9.0_9.3.1_edf
```

**4.** Modify the script so it is executable:

```
chmod +x datafabric_container_setup.sh
```

**5.** Before running the script, review the following considerations:

- Running the `datafabric_container_setup.sh` script requires sudo privileges.
- The script can take 5-10 minutes to run the first time you run it. Unless you pre-downloaded the image, the script downloads the latest Docker image from the Docker repository. For a list of available tags, see <https://hub.docker.com/r/maprtech/edf-seed-container/tags>.
- The script supports the `-i` option for specifying an image other than the latest image. The script also supports a `-p` option that must be used if the seed node is a cloud instance. To view the command line help for the script, see [Help for datafabric\\_container\\_setup.sh](#) on page 41.

6. Run the script to deploy the container for the Data Fabric image:

```
./datafabric_container_setup.sh
```

The script downloads the latest image. When the Docker image is running, you see the following output:

```
./datafabric_container_setup.sh
      RAM NEEDED : AVAILABLE
      DOCKER STATUS : RUNNING
      PORTS NEEDED : AVAILABLE
      PROCEEDING FORWARD WITH DEPLOYING SEED NODE
Please enter the local sudo password for root

latest: Pulling from maprtech/edf-seed-container
Digest:
sha256:b863f487de7eaa809b66f923aaec297aelcab4fdb9441950fbe5c68328235a7a
Status: Downloaded newer image for maprtech/edf-seed-container:latest
docker.io/maprtech/edf-seed-container:latest
Developer Sandbox Container 8935a11fb7d6 is running..
services required for Ezmeral Data fabric are coming up
services required for Ezmeral Data fabric are coming up
services required for Ezmeral Data fabric are coming up

Client has been configured with the docker container.

Please click on the link https://<hostname>:8443/app/dfui to deploy data
fabric
For user documentation, see https://docs.ezmeral.hpe.com/datafabric/home/
installation/installation_main.html
```

If the services do not come up, see [Troubleshooting Seed Node Installation](#) on page 34.

7. Navigate to the link specified in the Docker output message in step 6. The Data Fabric UI displays the **Create your first fabric** form.

### Create the Fabric

Bringing up the seed node automatically displays the fabric-creation interface. To create your first permanent fabric, complete the following steps.

1. On the **Create your first fabric** form, fill in the configuration parameters for the type of fabric you want to create. For more information, see the following topics:
  - [AWS Fabric Configuration Parameters](#) on page 29
  - [Azure Fabric Configuration Parameters](#) on page 31
  - [GCP Fabric Configuration Parameters](#) on page 32
  - [On-Premises Fabric Configuration Parameters](#) on page 33

## Create your first fabric

**Name\***

### Fabric guide

Submit the form to create your first fabric.

**i** Your first fabric must be created from a seed node. Once the fabric is created, you will be redirected to your fabric. For details, see [Fabric Deployment Using a Seed Node](#)

To submit a fabric order, visit [My HPE Software Center](#)

**Provider**

Amazon Web Services (AWS) ▼

See [AWS Fabric Configuration Parameters](#) for details about each field.

**AWS access credentials**  
User must have "AmazonEC2FullAccess" permission.

**Access key ID\***

**Secret key\***

**Fabric details**

**Region\***

US East (Ohio) ▼

- Click **Create**. The **Fabric details** dialog box is displayed. For example:

### Fabric details

<b>Fabric name</b>	testcluster
<b>Status</b>	<div style="width: 20px; height: 10px; background-color: #ccc; display: inline-block;"></div> <b>0%</b> <a href="#">See details</a>
<b>Provider</b>	On-premises

---

**i** Fabric setup is in progress. You will be redirected to Data Fabric UI once fabric setup is completed.

- To monitor the progress of fabric creation, check the status bar in the **Fabric details** dialog box, or click **See details**. Fabric creation can take 20 minutes or more.

While the installation is in progress, you can **Cancel** the installation at any time. If you click **Cancel**, a confirmation box asks for your credentials based on the provider type (AWS, GCP, Azure, or on-premises). Once you provide your credentials and click **Continue**, you cannot resume the current installation. However, you can use the **Re-initiate** button later to trigger a new installation request.

If fabric creation is successful, the Data Fabric UI displays a message with the endpoint link for the new fabric. For example:



## Fabric details

i Installing fabric.

<b>Fabric name</b>	newfabric
<b>Status</b>	<span style="color: green;">●</span> Complete
<b>Provider</b>	On-premises

---

Click the link to go to fabric  
[https://\[redacted\]:8443/app/dfui](https://[redacted]:8443/app/dfui) ↗

4. Using a browser, navigate to the endpoint link, and sign in to the Data Fabric UI for the newly created fabric using user name `admin` and password `p@ssw0rd`. These credentials are the default user name and password for Keycloak.



**IMPORTANT:** Hewlett Packard Enterprise recommends that you change the default Keycloak password immediately after installation. See [Changing the Keycloak admin Password](#) on page 51.

**HPE**  
Ezmeral Data Fabric

Username

Password

Sign In

Sign In with SSO

5. If fabric creation is successful, use the following steps to activate and register the fabric.



**NOTE:** After you have successfully created a cloud-based or on-premises fabric, you can kill the container that hosts the `edf-installer.hpe.com` fabric. For example, at the Docker command line type:

```
% docker kill <seed-node-container-ID>
```

6. If an error occurs during fabric creation, see [Troubleshoot Fabric Creation](#) on page 26 later on this page.

### Activate and Register the Fabric

After the first fabric is created, perform these steps:

1. Add the activation key for the fabric. See [Adding an Activation Key](#) on page 42.
2. Register the fabric. See [Registering a Fabric](#) on page 43.
3. Set the billing model. See [Setting the Billing Model](#) on page 43.
4. Configure single sign-on. See [SSO Using Keycloak](#) on page 50.

### Troubleshoot Fabric Creation

Fabric creation can fail for various reasons. You might see a message like this:

#### Fabric details

◆ Unable to create fabric. Delete the fabric to retry the installation.

<b>Fabric name</b>	testcluster
<b>Status</b>	◆ Failed <a href="#">See details</a>
<b>Provider</b>	On-premises

Delete

To view the log information, click **See details**. For example:

#### Fabric log

Fabric name: testcluster

◆ Unable to complete the installation.

```

'Ansible execution was not successful; 2: ansible-playbook [core 2.11.12] \n config file = None\n configured
module search path = [\/home/mapr/ansible/plugins/modules', \usr/share/ansible/plugins/modules]\n
ansible python module location = /opt/mapr/installer/build/installer/lib/python3.10/site-packages/ansible\n
ansible collection location = /home/mapr/ansible/collections:usr/share/ansible/collections\n executable
location = /opt/mapr/installer/build/installer/bin/ansible-playbook\n python version = 3.10.9
(tags/v3.10.9:1dd9be6, Oct 2 2023, 21:20:17) [GCC 7.5.0]\n jinja version = 3.1.2\n libyaml = False\nNo
config file found; using defaults\nSkipping callback 'default', as we already have a stdout callback.\nSkipping
callback 'minimal', as we already have a stdout callback.\nSkipping callback 'oneline', as we already have a
stdout callback.\n\nPLAYBOOK: 1-install-stanza-onprem.yml *****\n1 plays in
/opt/mapr/installer/ezdfaas/deployments/testcluster/installer/playbooks/1-install-stanza-onprem.yml\n\nPLAY
[setup machine] *****\nMETA: ran handlers\n\nTASK [Wait until
the installer host can be sshd to] *****\n\nTASK path:
/opt/mapr/installer/ezdfaas/deployments/testcluster/installer/playbooks/1-install-stanza-onprem.yml:10\nok:
                    
```

Cancel

Review the log information to determine if the failure is correctable. If you can resolve the failure condition, you can retry creating the fabric.

The **Reinitiate** button allows you to clean up failed fabric resources and trigger a fresh fabric request. To reinitiate fabric creation:

1. Click **Reinitiate**. The **Reinitiate fabric** form is displayed.
2. Fill in the required fields in the form.
3. Click **Create**. The **Fabric details** dialog box is displayed.

4. Monitor the progress of fabric creation as described in [Create the Fabric](#) on page 23.

To delete a failed fabric without reinitiating fabric creation, click **Delete** on the **Fabric details** dialog box.

If you are not able to resolve the issue that caused fabric creation to fail, contact [HPE Support](#).

### Prerequisites for On-Premises Installation

Describes fabric node and user prerequisites for on-premises installation of the HPE Ezmeral Data Fabric.

### Node Requirements for On-Premises Installation

Before deploying on-premises, you must provide the nodes that will host the on-premises fabric. Nodes that you want to include in the fabric must meet the following criteria:

#### One-Node Minimum Fabric

At least one node is required for an on-premises installation. You can provide as many additional nodes as you need.

#### Operating System

The following OS versions are supported:

- RHEL 9.4, 9.0, 8.8, 8.6, 8.5, 8.4, 8.3, 8.2, 8.1
- Rocky 8.5, 8.4
- Ubuntu 22.04, 20.04, 18.04
- SLES 15 SP3, 15 SP2
- OEL 8.4, 8.3, 8.2

#### Storage

For on-premises nodes, the HPE Ezmeral Data Fabric uses all the storage that is available to it. The platform does not impose limits based on the storage tier you select when creating the fabric. HPE recommends the following minimum number of nodes for each storage tier, but you do not need to follow this guideline precisely. The software does not currently check for a specific number of nodes or a specific amount of storage.

Storage Tier	Recommended Minimum Number of Nodes
1 TB	3
10 TB	5
100 TB	7
1 PG	12

#### Fully Qualified Domain Names (FQDNs)

The nodes must be expressed as fully-qualified domain names (FQDNs). DO NOT specify hostnames as aliases or IP addresses.

## Disk Space and Software Requirements

Nodes must meet the requirements in the following table. The Installer verifies the requirements prior to installation.

**Table**

Component	Requirements
CPU	64-bit x86.
CPU Cores	Minimum of 16 per node
OS	RHEL, Oracle Linux, Rocky, SLES, or Ubuntu.
Memory	32 GB minimum for nodes in production.
Disk	Raw, unformatted drives and no partitions.
DNS	Hostname, reaches all other nodes.
Users	Common users across all nodes; passwordless ssh (optional).
Java	Must run Java 11 or 17. Also, the Java and Java C versions must be the same on all nodes.
Python	The default Python version must be set to Python 3 on all nodes.
Other	NTP, Syslog, PAM.

Provide at least 10 GB of free disk space on the operating system partition. Provide 10 GB of free disk space in the `/tmp` directory and 128 GB of free disk space in the `/opt` directory. Services such as the ResourceManager and NodeManager use the `/tmp` directory. Files, such as logs and cores, use the `/opt` directory.

For data disks, the Installer requires a minimum disk size that is equal to the physical memory on the node. If a data disk does not meet the minimum disk size requirement, a verification error is generated.

## Proxy Server Requirements

If nodes in the fabric use an HTTP proxy server, the nodes must also meet the following requirements:

- The `no_proxy` environment variable must be set. Nodes in the fabric need to be able to communicate without the use of a proxy. If the `https_proxy` and `http_proxy` environment variable is set for nodes in the fabric, you must also set the `no_proxy` environment variable for the fabric admin user and the `root` user on each node. Configure the `no_proxy` environment variable to the IP range of the nodes or to the sub-domain that contains the nodes.

In addition, you must follow this guideline from the [Python documentation](#): "The `no_proxy` environment variable can be used to specify hosts which shouldn't be reached via proxy; if set, it should be a comma-separated list of hostname suffixes, optionally with `:port` appended, for example `cern.ch,ncsa.uiuc.edu,some.host:8080`."

For cloud-based fabrics (Amazon EC2, Google Compute Engine (GCE), and Microsoft Azure), you must include this entry in the no-proxy configuration:

```
169.254.169.254
```

- The global proxy for package repositories must be set. The Installer creates repository files. However, the proxy setting is not configured for each repository. Therefore, configure global proxy settings on each node in the fabric.
  - On CentOS/RedHat, set global proxy settings in `/etc/yum.conf`.
  - On Ubuntu, set global proxy settings in `/etc/apt/apt.conf`.

## User Requirements for On-Premises Installation

On-premises nodes must meet the following requirements for users:

- For all users, the numeric user and group IDs (UID and GID) must be configured, and these values must match on all nodes in all fabrics.
- The `mapr` user and `root` user must be configured to use bash. Other shells are not supported.
- The user that initiates fabric creation for an on-premises deployment must be present and have the same user name and password on all fabric nodes.
- Using an SSH key is not supported during on-premises fabric creation.

## AWS Fabric Configuration Parameters

This page describes the configuration values that you need to specify to create a new fabric using Amazon Web Services (AWS).

Parameters with an asterisk (\*) are required. Before you can initiate the **Create** process, you must specify all required parameters.

<b>Name*</b>	<p>Name of the fabric. Use a name that is unique across all of your fabrics and is from 1 to 40 characters. The name:</p> <ul style="list-style-type: none"> <li>• Must start with a letter (either lowercase or uppercase).</li> <li>• Can contain lowercase letters, uppercase letters, numbers, and hyphens.</li> <li>• Must not contain consecutive hyphens.</li> <li>• Must include a letter or a number as the final character.</li> </ul>
<b>Provider</b>	The cloud provider on which to create the fabric. Select <b>Amazon Web Services (AWS)</b> .
<b>Access key ID*</b>	<p>AWS credential Access Key.</p> <p>The user must have "AmazonEC2FullAccess" permission.</p>
<b>Secret key*</b>	AWS credential Secret Access Key.
<b>Region*</b>	The AWS region in which to provision the fabric.
<b>Storage Tier*</b>	<p>The consumption baseline that you specified in your license for the fabric. Your actual storage consumption can exceed this level. Select from these tiers:</p> <ul style="list-style-type: none"> <li>• 1 TB</li> <li>• 10 TB</li> <li>• 100 TB</li> <li>• 1 PB</li> </ul>
<b>Data-at-rest encryption</b>	Data on disk (or data at rest) on a secure fabric can be encrypted, enabling you to protect the data if a disk is compromised. Encryption of data at rest not only prevents unauthorized users from accessing sensitive data, but it also protects against data theft via sector-level disk access. Data-at-rest encryption is ON by default.
<b>Nodes</b>	The number of nodes allocated based on the Storage tier you selected. You do not need to specify a number. The nodes are populated automatically.
<b>Virtual Private Cloud (VPC) ID*</b>	The AWS Virtual Private Cloud (VPC) ID to use in the selected region. For example: <code>vpc-0b5177b19511ee301</code> . You must provide a VPC, and the VPC must have an internet gateway attached.
<b>Public subnet ID*</b>	<p>The subnet ID to use in the selected VPC. For example: <code>subnet-0445a49217546b101</code>.</p> <p>The public subnet must be accessible from the internet.</p>

## Azure Fabric Configuration Parameters

This page describes the configuration values that you need to specify to create a new fabric using Microsoft Azure.

Parameters with an asterisk (\*) are required. Before you can initiate the **Create** process, you must specify all required parameters.

<b>Name*</b>	<p>Name of the fabric. Use a name that is unique across all of your fabrics and is from 1 to 40 characters. The name:</p> <ul style="list-style-type: none"> <li>• Must start with a letter (either lowercase or uppercase).</li> <li>• Can contain lowercase letters, uppercase letters, numbers, and hyphens.</li> <li>• Must not contain consecutive hyphens.</li> <li>• Must include a letter or a number as the final character.</li> </ul>
<b>Provider</b>	<p>The cloud provider on which to create the fabric. Select <b>Azure</b>.</p>
<b>Azure tenant ID*</b>	<p>The ID of the Azure tenant. For information about how to obtain the ID, see <a href="#">this website</a>.</p> <p>The tenant must be accessible from the internet.</p>
<b>Subscription ID*</b>	<p>The Azure subscription ID. Azure tenants can have one or more subscriptions, which are agreements with Microsoft to use Azure services. Every Azure resource is associated with a subscription. For information about how to obtain the ID, see <a href="#">this website</a>.</p> <p>The subscription must have an attached internet gateway.</p>
<b>Client ID*</b>	<p>The ID of the Azure client (application) in the Active Directory. For information about how to obtain the ID, see <a href="#">this website</a>.</p> <p>The client must be accessible from the internet.</p>
<b>Client secret*</b>	<p>The Azure client (application) secret in the Active Directory. For information about how to obtain the secret, see <a href="#">this website</a>.</p> <p>The client secret must be accessible from the internet.</p>
<b>Region*</b>	<p>The Azure region in which to provision the fabric.</p>
<b>Storage Tier*</b>	<p>The consumption baseline that you specified in your license for the fabric. Your actual storage consumption can exceed this level. Select from these tiers:</p> <ul style="list-style-type: none"> <li>• 1 TB</li> <li>• 10 TB</li> <li>• 100 TB</li> <li>• 1 PB</li> </ul>

<b>Data-at-rest encryption</b>	Data on disk (or data at rest) on a secure fabric can be encrypted, enabling you to protect the data if a disk is compromised. Encryption of data at rest not only prevents unauthorized users from accessing sensitive data, but it also protects against data theft via sector-level disk access. Data-at-rest encryption is ON by default.
<b>Nodes</b>	The number of nodes allocated based on the Storage tier you selected. You do not need to specify a number. The nodes are populated automatically.
<b>Resource group name*</b>	The name of the Azure resource group. The resource group is a container that comprises multiple resources and facilitates the management of those resources.
<b>Virtual network*</b>	The name of the Azure Virtual Network (VNet).
<b>Subnetwork*</b>	The name of the subnet in your virtual network to be used for the fabric.

### GCP Fabric Configuration Parameters

This page describes the configuration values that you need to specify to create a new fabric using Google Cloud Platform (GCP).

Parameters with an asterisk (\*) are required. Before you can initiate the **Create** process, you must specify all required parameters.

<b>Name*</b>	<p>Name of the fabric. Use a name that is unique across all of your fabrics and is from 1 to 40 characters. The name:</p> <ul style="list-style-type: none"> <li>• Must start with a lowercase letter.</li> <li>• Can contain lowercase letters, numbers, and hyphens.</li> <li>• Must not contain consecutive hyphens.</li> <li>• Must include a lowercase letter or a number as the final character.</li> </ul>
<b>Provider</b>	The cloud provider on which to create the fabric. Select <b>Google Cloud Platform (GCP)</b> .
<b>Service account key file*</b>	A file containing your GCP service account credentials. For more information, see <a href="#">Create and delete service account keys</a> .
<b>Zone*</b>	The GCP zone in which to provision the fabric.
<b>Storage Tier*</b>	<p>The consumption baseline that you specified in your license for the fabric. Your actual storage consumption can exceed this level. Select from these tiers:</p> <ul style="list-style-type: none"> <li>• 1 TB</li> <li>• 10 TB</li> <li>• 100 TB</li> <li>• 1 PB</li> </ul>



<b>Data-at-rest encryption</b>	Data on disk (or data at rest) on a secure fabric can be encrypted, enabling you to protect the data if a disk is compromised. Encryption of data at rest not only prevents unauthorized users from accessing sensitive data, but it also protects against data theft via sector-level disk access. Data-at-rest encryption is ON by default.
<b>Nodes</b>	The number of nodes allocated based on the Storage tier you selected. You do not need to specify a number. The nodes are populated automatically.
<b>VPC network*</b>	The identifier for the VPC. The VPC must have an internet gateway attached.
<b>Subnetwork*</b>	The identifier for the public subnet.

### On-Premises Fabric Configuration Parameters

This page describes the configuration values that you need to specify to create a new fabric that is hosted on-site.

Parameters with an asterisk (\*) are required. Before you can initiate the **Create** process, you must specify all required parameters.

Creating an on-premises fabric requires you to provide host nodes *before* starting fabric creation. These nodes must meet certain prerequisites. Before creating an on-premises fabric, review [Prerequisites for On-Premises Installation](#) on page 27.

<b>Name*</b>	Name of the fabric. Use a name that is unique across all of your fabrics and is from 1 to 40 characters. The name: <ul style="list-style-type: none"> <li>• Must start with a letter (either lowercase or uppercase).</li> <li>• Can contain lowercase letters, uppercase letters, numbers, and hyphens.</li> <li>• Must not contain consecutive hyphens.</li> <li>• Must include a letter or a number as the final character.</li> </ul>
<b>Provider</b>	The cloud or on-premises platform where you want to create the fabric. Select <b>On-premises</b> .
<b>Username*</b>	The SSH username.
<b>Password*</b>	The SSH password.
<b>Airgap repository</b>	The repository for the Installer to use if your installation cannot access the internet. The repository must contain nested folders. For example: <code>./installer/redhat</code> . You must create this repository before installing an air-gapped fabric.
<b>Data-at-rest encryption</b>	Data on disk (or data at rest) on a secure fabric can be encrypted, enabling you to protect the data if a disk is compromised. Encryption of data at rest not only prevents unauthorized users from accessing sensitive data, but it also protects against data theft via sector-level disk access. Data-at-rest encryption is ON by default.

**IPv6 support**

The setting that enables or disables support for IPv6 network addresses. For more information about IPv6, see [IPv6 Support in Data Fabric](#) on page 104.

IPv6 support is OFF by default. You can enable IPv6 support only for on-premises deployments (not for cloud deployments). Once IPv6 support is enabled, it cannot be disabled.

**Nodes\***

The recommended minimum number of nodes that should be allocated. The form provides this information based on the Storage tier you selected.

**Node FQDN**

The fully qualified domain name of a node that will host the fabric. This is a required field. Specify the FQDN of a node that you provided, as described in [Prerequisites for On-Premises Installation](#) on page 27. For example: `mynode.lab.mycompany.net`. Use fully qualified domain names (FQDNs). DO NOT specify hostnames as aliases or IP addresses.

If you are using multiple nodes, click the **Add node** button to add as many additional nodes as you provisioned.

**EDF subnet**

The Data Fabric subnet. This parameter is optional. This parameter allows you to set a subnet mask to restrict fabric services to a subset of network interface cards (NICs). Specify one or more comma-separated subnet masks. For example:

```
10.10.15.0/24,10.10.16.0/24
```

**EDF external**

The Data Fabric external IP addresses for the CLDB, file system, and MAST Gateway nodes. This parameter is optional. This parameter allows you to designate a specific IP of the host as a public IP address to handle the external traffic targeted to the host. Specify a comma-separated list of tuples using this format: `<hostname>:<host_external_IP_address>`. For example:

```
host1.corp.net:1.1.1.1,host2.corp.net:1.1.1.2,host3.corp.net:1.1.1.3
```

**Troubleshooting Seed Node Installation**

Describes some common issues that can interfere with seed node installation.

**Issue: Fabric Creation Fails If Host-Name Resolution Takes Too Long**

Fabric creation requires host-name resolution to complete in less than 300 ms. If host-name resolution takes longer than 300 ms, fabric creation can fail. To check the speed of host-name resolution on a node, enter one of these commands:

- `time nslookup <hostname>`
- `time getent ahosts <hostname>`

Host-name resolution depends on the name servers that you are using. You can find them in the `/etc/resolv.conf` file. If you cannot use faster name servers, try improving the resolution time by providing

local resolution. You can do this by updating the `/etc/hosts` file on each of the fabric nodes with the details of the other fabric nodes. For example:

```
<node1IP> <node1fqdn> <node1hostname>
<node2IP> <node2fqdn> <node2hostname>
<node3IP> <node3fqdn> <node3hostname>
```

### Issue: Seed Node Container Services Do Not Come Up

Use these steps to resolve the problem:

1. After running the script to deploy the container for the Data Fabric image, wait for at least 10 minutes for services to come up. If the services do not come up, the screen can display a message such as the following:

```
7.6.0-mapr-devdocker-container % ./datafabric_container_setup.sh
Please enter the local sudo password for <username>
Password:

latest: Pulling from maprtech/edf-seed-container
Digest:
sha256:052f461d98b1d0b8251cd47bab71b42103e61aaaa33d31335d3ca60182f4a87e
Status: Image is up to date for maprtech/edf-seed-container:latest
docker.io/maprtech/edf-seed-container:latest
Developer Sandbox Container b4be66858760 is running..
services required for Ezmeral Data fabric are coming up
services required for Ezmeral Data fabric are coming up
services required for Ezmeral Data fabric are coming up
services required for Ezmeral Data fabric are coming up
services required for Ezmeral Data fabric are coming up
services required for Ezmeral Data fabric are coming up
services required for Ezmeral Data fabric are coming up
services required for Ezmeral Data fabric are coming up
services required for Ezmeral Data fabric are coming up
services required for Ezmeral Data fabric are coming up
services required for Ezmeral Data fabric are coming up

services required for Ezmeral Data fabric are coming up

services didnt come up in stipulated 10 mins time
please login to the container using ssh root@localhost -p 2222 with mapr
as password and check further
For documentation on steps to debug, see https://docs.ezmeral.hpe.com/
datafabric/home/installation/troubleshooting_seed_node_installation.html
once all services are up fabric UI is available at https://
<hostname>:8443/app/dfui and fabrics can be deployed from that page
```

2. Sign in to the Docker container using `mapr` as the password:

```
ssh root@localhost -p 2222
```

3. Enter the `jps` command and check the output. Continue entering the `jps` command until the command shows the `AdminApplication` java process, which indicates that all the services are started:

```
root@edf-installer:~# jps
71136 FsShell
71315 Jps
19349 AdminApplication
10024 WardenMain
14236 CLDB
13213 QuorumPeerMain
```

4. If the services do not start, check that sufficient resources have been allocated to the seed node. See the "Seed Node Prerequisites" in [Fabric Deployment Using a Seed Node](#) on page 18. You might need to allocate more resources and retry installing the seed node.
5. Check to see if the Warden and ZooKeeper services are up and running:

```
systemctl status mapr-warden
systemctl status mapr-zookeeper
```

6. If the services did not start within 10 minutes, check the following logs for errors or exceptions. If the logs contain errors or exceptions, contact [HPE Support](#). If there are no errors or exceptions, start the services (see step 7).

- `/opt/mapr/logs/clldb.log`
- `/opt/mapr/logs/configure.log`
- `/opt/mapr/logs/warden.log`
- `/opt/mapr/apiserver/logs/apiserver.log`
- `/opt/mapr/zookeeper/zookeeper-3.5.6/logs/zookeeper.log`

7. If the Warden and ZooKeeper services are not up and running, try restarting the services manually:

```
systemctl start mapr-zookeeper
systemctl start mapr-warden
```

### Downloading Installation Logs


This section describes how to download a zip archive of the installation logs for a specified fabric by using the Data Fabric UI.

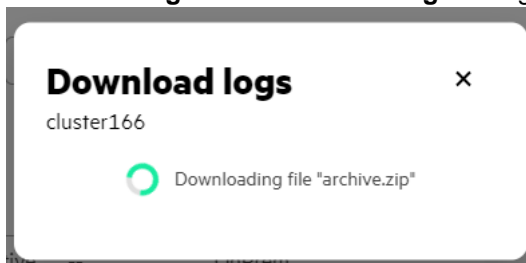
You can use the Data Fabric UI **Download logs** button only on fabrics created using the seed node installer or the **Create fabric** button. The **Download logs** button is not available on fabrics installed manually.

### Steps for Downloading the Log Archive


To download all the logs for a fabric:

1. Log on to the Data Fabric UI.
2. Select the **Fabric manager** option.
3. Click **Global namespace**, and check the table view.

- For any fabric listed in the table view, click the ellipsis (  ) under the **Action** column, and select **Download logs**. The **Download logs** dialog shows the status of the download:



The download operation copies a zip archive to the **Downloads** folder on your workstation. For example, in the Chrome browser, you can view the contents of the **Downloads** folder by clicking the **Downloads** icon in the upper right corner of the browser:



- To view the logs, un-zip the archive to a directory of your choosing.

### Contents of Log Archive

The log archive includes `\data` and `\log` directories that contain the following files:

Archive Directory	Log Files	Number of Files
\data	installer.mv.db	1 per fabric
	properties.json	1 per fabric
\logs	create-ssl-keys.log	1 per fabric
	installer.json[.x]	1 per node
	installer.log	1 per fabric
	installer_cli.log	1 per fabric
	installer_cli_root.log	1 per fabric
	installer-process.log	1 per fabric
	<FQDN>.log	1 per node
	<FQDN>_cfg_r_<timestamp>.log	1 per node
mapr-installer.log[.x]	1 per node	

### Related maprcli Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. A link to this command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- [installer logs](#)

**Creating a Local Repository for an Air-Gapped Installation**

Describes how to make installation packages available through a local repository for an air-gapped installation.

You can set up a local repository on each node to provide access to installation packages. With this method, nodes do not require internet connectivity. The package manager on each node installs from packages in the local repository. To set up a local repository, nodes need access to a running web server to download the packages.

**Creating a Local Repository on RHEL, Rocky, or Oracle Linux**

Describes how to create and use a local repository for RHEL, Rocky, or Oracle Linux.

1. Ensure that you have access to the HPE internet repository so that you can download package files. For more information, see [Accessing the HPE Ezmeral Token-Authenticated Internet Repository](#) on page 74.
2. On the machine where you will set up the repository, log in as `root` or use `sudo`.
3. Create the following directory if it does not exist: `/var/www/html/yum/base`
4. On a computer that is connected to the internet, download the following files, substituting the appropriate `<version>` number and `<datestamp>`:

```
https://package.ezmeral.hpe.com/releases/v7.x.x/redhat/
mapr-<version>GA.rpm.tgz
https://package.ezmeral.hpe.com/releases/MEP/MEP-<version>/redhat/
mapr-mep-<version>-<datestamp>.rpm.tgz
```

5. Copy the files to `/var/www/html/yum/base` on the node, and extract them there:

```
tar -xvzf mapr-v<version>GA.rpm.tgz
tar -xvzf mapr-mep-v<version>.<datestamp>.rpm.tgz
```

6. Create the base repository headers by using the following command:

```
createrepo /var/www/html/yum/base
```

HPE software assumes that the web server root directory is: `/var/www/html`.

7. When finished, verify the content of the new `/var/www/html/yum/base/repodata` directory:

```
filelists.xml.gz, other.xml.gz, primary.xml.gz, repomd.xml
```

**Creating a Local Repository on SLES**

Describes how to create and use a local repository for SLES.

1. Ensure that you have access to the HPE internet repository so that you can download package files. For more information, see [Accessing the HPE Ezmeral Token-Authenticated Internet Repository](#) on page 74.
2. On the machine where you will set up the repository, log in as `root` or use `sudo`.
3. Create the following directory if it does not exist: `/var/www/html/zypper/base`

4. On a computer that is connected to the internet, download the following files, substituting the appropriate `<version>` and `<datestamp>`:

```
https://package.ezmeral.hpe.com/releases/v<version>/suse/
mapr-<version>GA.rpm.tgz
https://package.ezmeral.hpe.com/releases/MEP/MEP-<version>/suse/
mapr-mep-<version>-<datestamp>.rpm.tgz
```

5. Copy the files to `/var/www/html/zypper/base` on the node, and extract them there:

```
tar -xvzf mapr-<version>GA.rpm.tgz
tar -xvzf mapr-mep-<version>-<datestamp>.rpm.tgz
```

6. Create the base repository headers:

```
createrepo /var/www/html/zypper/base
```

7. When finished, verify the content of the new `/var/www/html/zypper/base/repodata` directory:

```
filelists.xml.gz, other.xml.gz, primary.xml.gz, repomd.xml
```

### Creating a Local Repository on Ubuntu

Describes how to create and use a local repository for Ubuntu.

1. Ensure that you have access to the HPE internet repository so that you can download package files. For more information, see [Accessing the HPE Ezmeral Token-Authenticated Internet Repository](#) on page 74.
2. On the machine where you will set up the repository, log in as `root`.
3. Change to the directory `/root`, and create the following directories within it:

```
~/mapr
|---dists
|-----binary
|-----optional
|-----binary-amd64
|---mapr
```

4. On a computer that is connected to the internet, download the following files, substituting the appropriate `<version>` and `<datestamp>`:

```
https://package.ezmeral.hpe.com/releases/v7.x.x/ubuntu/
mapr-<version>GA.deb.tgz
https://package.ezmeral.hpe.com/releases/MEP/MEP-<version>/ubuntu/
mapr-mep-<version>-<datestamp>.deb.tgz
```

5. Copy the files to `/root/mapr/mapr` on the node, and extract them there:

```
tar -xvzf mapr-<version>GA.deb.tgz
tar -xvzf mapr-mep-<version>-<datestamp>.deb.tgz
```

6. Navigate to the `/root/mapr` directory.

- Use `dpkg-scanpackages` to create `Packages.gz` in the `binary-amd64` directory:

```
dpkg-scanpackages . /dev/null | gzip -9c > ./dists/binary/optional/binary-amd64/Packages.gz
```

- Move the entire `/root/mapr/mapr` directory to the default directory served by the HTTP server (for example, `/var/www`), and make sure the HTTP server is running.

### Planning Worksheet for Cloud Deployments

Print this worksheet, and use it to record configuration information for your cloud deployment.

For installation information, see [Fabric Deployment Using a Seed Node](#) on page 18.

### Host and Networking Configuration Information

Description	Value
Seed node host name	
Network Interface Card (NIC) name	
Proxy settings	<code>export http_proxy=</code>
	<code>export https_proxy=</code>
	<code>export HTTP_PROXY=</code>
	<code>export HTTP_PROXY=</code>

### AWS Information

AWS Parameter	Value
Name	
Provider	Amazon Web Services (AWS)
Access key ID	
Secret key	
Region	
Storage Tier	
Data-at-rest encryption	
Virtual Private Cloud (VPC) ID	
Public subnet ID	

### Azure Information

Azure Parameter	Value
Name	
Provider	Azure
Azure tenant ID	
Subscription ID	
Client ID	



Azure Parameter	Value
Client secret	
Region	
Storage Tier	
Data-at-rest encryption	
Resource group name	
Virtual network	
Subnetwork	

### GCP Information

GCP Parameter	Value
Name	
Provider	Google Cloud Platform (GCP)
Service account key file	
Zone	
Storage Tier	
Data-at-rest encryption	
VPC network	
Subnetwork	

### Help for `datafabric_container_setup.sh`

From the Docker command line, you can access the help text for the `datafabric_container_setup.sh` script.

To view the help for the setup script, use the `./datafabric_container_setup.sh -h` command:

```
% ./datafabric_container_setup.sh -h
This script will take of deploying edf on seed node.

Syntax: ./datafabric_container_setup.sh [-i|--image] [-p|--publicip4dns]
[-f|--proxyfiledetails]
options:
-i|--image this is optional,By default it will pull image having latest tag,
we can also provide image which has custom tag example:maprtech/
edf-seed-container:7.4.0_9.1.2
-p|--publicip4dns is the public IPv4 DNS and needed for cloud deployed
seed nodes. Note that both inbound and outbound traffic on port 8443
needs to be enabled on the cloud instance. Otherwise, the Data
Fabric UI cannot be acessible
-f|--proxyfiledetails is the location of file from where proxy details
provided by user are copied to docker container.
```

Normally, using the `-i|--image` option is not needed. You may provide the `-i` option if you want to specify a specific image. If you do not provide the option, the latest available image is downloaded.

The `-f|--proxyfiledetails` option allows you to specify proxy information in a file. On Linux nodes, if you do not provide the `-f` option, the contents of `/etc/profile.d/proxy.sh` and `/etc/environment` are appended and copied to the container. On a Mac seed node, if you do not provide the `-f` option, no proxy details are copied.

## Service Activation and Billing

Describes how to activate and register a new fabric to take advantage of automated billing.

When you [install](#) HPE Ezmeral Data Fabric software, you have the option to install in a connected environment or air-gapped environment:

Environment	Description
Connected	An environment that has continuous internet access.
Air-gapped	An environment that has no internet access, usually for the purpose of increasing security.

The activation and billing processes differ for each type of installation. In a connected environment, billing is an automated process. In an air-gapped environment, the billing process is manual and requires an activation code.

The following sections describe what you need to do to activate your Data Fabric and enable billing to keep the fabric operational.

### Adding an Activation Key

You must add an activation key after installing the HPE Ezmeral Data Fabric or adding a new fabric.

An activation key is provided when you purchase a consumption-based license. Adding the activation key enables the Data Fabric UI to display important details about your license, including the start date and expiration dates. You must be a fabric manager to add an activation key.

To add an activation key:

1. Sign in to the Data Fabric UI, and switch to the **Fabric manager experience**.
2. Click the **Fabric administration** button.
3. On the **Activation** card, click **Add activation key**. The **Add activation key** form appears.
4. For the **Input type**, select **File** or **Text**.
5. If you selected **File** as the **Input type**, drag and drop the activation key file into the box on the form. Or, click **Select File** to navigate to the file and select it.
6. If you selected **Text** as the **Input type**, copy and paste your activation key information into the box on the form.
7. Click **Add**. Your activation details become visible on the **Activation** card.

#### Related concepts

[Obtaining a License](#) on page 45

Describes the process of obtaining a consumption-based license from the My HPE Software Center.

[Registering a Fabric](#) on page 43

You register the HPE Ezmeral Data Fabric after installing the fabric or adding a new fabric. Registration provides HPE with information about your internet connection and determines the billing process that you will use.

[Setting the Billing Model](#) on page 43

Setting the billing model enables the Data Fabric UI to display estimated billing charges for each fabric.

## Registering a Fabric

You register the HPE Ezmeral Data Fabric after installing the fabric or adding a new fabric. Registration provides HPE with information about your internet connection and determines the billing process that you will use.

You must be a fabric manager to register a fabric. If your fabric is in an air-gapped environment, you must obtain an activation code, which is provided when you purchase a consumption-based license. To obtain a license, see [Obtaining a License](#) on page 45.

If your fabric is behind a proxy, registration for the `Connected` mode can fail unless you first configure the proxy server. See [Configuring a Proxy Server for Data Fabric Access to the Internet](#) on page 108.

To register a fabric:

1. Sign in to the Data Fabric UI, and switch to the **Fabric manager experience**.
2. Click the **Fabric administration** button.
3. On the **Activation** card, click **Actions**, and select **Register fabric**. The **Register fabric** form appears.
4. Select the **Operational mode** as `Air-gapped` or `Connected`:

Mode	Description
Air-gapped	The fabric is not connected to the Internet. In the air-gapped mode, you must provide an activation code to the Data Fabric UI.
Connected	The fabric is connected to the Internet and can communicate with the HPE billing service. In the connected mode, the system automatically communicates an activation code to HPE.

5. For an air-gapped deployment, upload the activation code file that you received when you purchased your license. You can drag and drop the file into the box on the form, or click **Select File** to navigate to the code file.
6. Click **Register**. Your registration details become visible on the **Activation** card.

### Related concepts

[Setting the Billing Model](#) on page 43

Setting the billing model enables the Data Fabric UI to display estimated billing charges for each fabric.

[Viewing Activation Information](#) on page 44

Use the Data Fabric UI to view important activation information, such as the status of your activation key and activation code (for air-gapped installations).

## Setting the Billing Model

Setting the billing model enables the Data Fabric UI to display estimated billing charges for each fabric.

Setting the billing model is optional. However, if you set the billing model, the Data Fabric UI can show estimates of your aggregated and on-demand billing charges on the **Billing and Storage Consumption** card. See [View Billing Data by Fabric](#) on page 123.

To set the billing model, you must enter information provided to you by HPE when you purchased your license for the fabric. Use these steps to provide the information:

1. Sign in to the Data Fabric UI, and switch to the **Fabric manager experience**.
2. Click the **Fabric administration** button.
3. On the **Activation** card, click **Set billing model**. The **Set billing model** form appears.
4. Fill in the form as follows (an asterisk (\*) denotes required fields):

<b>Commit amount*</b>	The minimum amount of storage that you committed to purchase in your Data Fabric license.
<b>Unit</b>	The units for the commit amount. Available units are: <ul style="list-style-type: none"> <li>• TB (terabytes)</li> <li>• PB (petabytes)</li> <li>• EB (exabytes)</li> <li>• ZB (zettabytes)</li> </ul>
<b>Commit rate*</b>	The monthly storage charge in dollars (\$) per GB hour that you committed to when you purchased your license. To ensure that the UI provides accurate estimates, be sure to factor in any discount in the <b>Commit rate</b> that you received from HPE.
<b>On-demand rate*</b>	The rate in dollars (\$) per GB hour for storage use in excess of your commit amount. To ensure that the UI provides accurate estimates, be sure to factor in any discount in the <b>On-demand rate</b> that you received from HPE.

5. Click **Save**.

## Viewing Activation Information

Use the Data Fabric UI to view important activation information, such as the status of your activation key and activation code (for air-gapped installations).

To view activation information:

1. Sign in to the Data Fabric UI, and switch to the **Fabric manager experience**.
2. Click the **Fabric administration** button.
3. Locate the **Activation** card, which displays detailed information about your Activation code (for air-gapped installations) and Activation key.

## Displaying the Fabric ID

Describes how to display the fabric ID.

You must supply your fabric ID in order to [obtain a license](#). Each fabric has a unique fabric ID (sometimes also referred to as a “cluster ID”).

Before you apply an activation key to your fabric, you can display the fabric ID by using the following command at a Linux command prompt:

```
$ cat /opt/mapr/conf/clusterid
4626587677795940777
```

After you add a valid activation key, the Data Fabric UI displays the fabric ID in the **Fabric administration** page. For example:

Welcome, admin Fabric manager

Global namesp... Fabric metrics **Fabric adminis...** Security admi...

**Fabric administration**

Select fabric

venkat-cluster214

**Fabric ID** 5747617797603936219  
**Build version** 7.5.0.0.20231020123031.GA

Monitoring **Fabric administra...** Administration

### Fabric administration

**Fabric name** venkat\_cluster214  
**Fabric ID** 5534713928724326568  
**Build version** 7.4.0.0.20230721023422.GA

To view **Fabric administration** information:

1. Sign in to the Data Fabric UI, and switch to the **Fabric manager experience**.
2. Click the **Fabric administration** button.

## Obtaining a License

Describes the process of obtaining a consumption-based license from the My HPE Software Center.

To obtain a license:

1. For new deployments, install the fabric as described in [Installation](#) on page 18. Or, for existing deployments, create a new fabric as described in [Creating a Fabric](#) on page 109, or import a fabric as described in [Importing a Fabric](#) on page 110.
2. Note the ID of the new fabric. To obtain a license, you must supply the fabric ID. See [Displaying the Fabric ID](#) on page 44.
3. After purchasing HPE Ezmeral Data Fabric software, a license key is made available to you through the **Access Your Products** button in the **HPE Subscription Electronic Receipt** email that you receive from HPE. This receipt will direct you to [MY HPE SOFTWARE CENTER](#) where you can activate your product
4. Log in to the [MY HPE SOFTWARE CENTER](#) with your HPE Passport user ID and password. You should see an **Activate EON:** page.
5. In the **Qty to Activate** field, specify 1.
6. Click **Confirm Selection**.

7. In Step 2: Designate Activatee, click **Next**.
8. In Step 3, enter your cluster ID (fabric ID).
9. Click **Activate**. The activation process can take several minutes to complete. Eventually, the HPE Ezmeral Data Fab SW Base SaaS page is displayed.
10. Click the box to accept the license terms and authorizations.
11. Click the box for **Licenses Keys (3)**.
12. Click **Download**. The licenses are downloaded as .DAT files.

You can now add an activation key to your fabric. See [Adding an Activation Key](#) on page 42.

## Billing in Connected Environments

Describes how billing is enabled in a connected environment.

### Automatic Billing

A connected environment is an installation in which the fabric is connected to the internet and can communicate with the HPE billing service. In a connected environment, billing and activation are initiated automatically after you [add an activation key](#) and [register a fabric](#).

### Viewing the Billing and Consumption Information for a Connected Fabric

For a connected environment, the Data Fabric UI displays billing and consumption information on the **Billing and Storage Consumption** card. Note that cost information is estimated. This information is based on the rates you specified in [Setting the Billing Model](#) on page 43. The actual cost reflected on the billing portal might be different.

To view the billing information:

1. Sign in to the Data Fabric UI, and switch to the **Infrastructure admin** view or **Fabric manager** view.
2. Click **Fabric metrics**.
3. Scroll down to see the **Billing and Storage Consumption** card.

### Viewing the Operational Mode for a Connected Fabric

The operational mode of a fabric refers to the internet connection status, which can be *Connected* or *Air-gapped*. To view the operational mode:

1. Sign in to the Data Fabric UI, and switch to the **Fabric manager** view.
2. Click **Fabric administration**.
3. Locate the **Activation** card. The **Activation code** section of the card displays the operational mode.

### If You Forget to Pay Your Invoice

If you forget to pay your invoice or fail to renew an expired license, a connected fabric can be disabled by HPE. If you suspect that the fabric has been disabled, contact HPE Support. To restore the fabric, see [Restoring a Disabled Fabric](#) on page 49.

## Billing in Air-Gapped Environments

Describes how billing is enabled in an air-gapped environment.

In an air-gapped environment, manual steps are needed to support billing and activation for a consumption-based fabric. This section describes how to activate an air-gapped fabric and keep the fabric operational.

### Using `maprcli` Commands in an Air-Gapped environment

Some tasks for keeping an air-gapped fabric operational require you to use `maprcli` commands. This is because certain operations are not currently available in the Data Fabric UI. The `maprcli` commands you need are provided on this page.

To run `maprcli` commands, use an `ssh` connection to any node in the fabric.

### Understanding the Activation Code and Billing Cycle

When you place an order for the HPE Ezmeral Data Fabric and specify an air-gapped environment, HPE provides you with an activation code. The activation code allows you to register the product and sign usage records for one billing cycle. The billing cycle is one month with a 15-day grace period.

The activation code has two important dates:

- **Start Date** – The first day of the one-month billing cycle.
- **End Date** – The end of the month-long billing cycle and the start of the grace period. This date is usually 30 days after the Start Date.

You can view these dates by using the Data Fabric UI:

1. Sign in to the Data Fabric UI, and switch to the **Fabric manager** view.
2. Click **Fabric administration**.
3. Locate the **Activation** card. The **Activation code** section of the card displays the start date, end date, and current month charges.

Once the activation code is applied, the code is valid (and the fabric is operational) until the **End Date**. After the **End Date**, a short grace period is applied to allow you to perform the steps to maintain activation.

### Steps for Maintaining Activation

For an air-gapped environment, the fabric administrator must perform the following steps to keep the fabric activated:

Step	Task	When
1.	Add an activation key.	Installation time
2.	Register the fabric.	Installation time
3.	Collect usage records.	Monthly
4.	Send a usage record file to HPE.	Monthly
5.	Pay the monthly HPE invoice.	Monthly
6.	Renew your activation to keep the fabric operational.	Monthly

#### 1. Add an Activation Key

Regardless of your fabric's operational mode, you must obtain a license and add the activation key from the license by using the Data Fabric UI. See [Adding an Activation Key](#) on page 42.

## 2. Register a New Air-Gapped Fabric

For a new air-gapped fabric, you must register the fabric by selecting the operational mode and uploading the activation code provided by HPE Support when you ordered the product. See [Registering a Fabric](#) on page 43.

After registration, the fabric is usable for a month with a 15-day grace period. The fabric continues to be usable as long as you continue to pay your monthly bill and reapply new activation keys.

## 3. Collect Usage Records

On or near the first day of each month, the fabric administrator should collect usage records for the previous month. The following `maprcli` command collects your usage records in a file named `usage_file.txt`:

```
maprcli cluster getbillingusage -fileName usage_file.txt -clearText true
```

Licensing for the HPE Ezmeral Data Fabric is consumption based, meaning that you are charged based on actual usage. Usage is measured in storage-hour units.

For both connected and air-gapped Data Fabrics, the [container location database \(CLDB\)](#) collects usage metrics. The CLDB analyzes the logical data size of all volumes in the fabric and records the average consumption for each hour at the end of the hour. Even before the fabric is activated and billing is enabled, you can view the recorded metrics by using the following `maprcli` command:

```
maprcli cluster getmeteringusage -from <date-time> -till <date-time> -cleartext true
```

For example:

```
$ maprcli cluster getmeteringusage -from 2021-01-01,00:00 -till 2023-05-13,00:00 -clearText true
userdata  metadata  total  epoch  timestamp
1 Mb      0 Mb      1 Mb   1683716400000  Wed May 10 11:00:00 UTC 2023
1 Mb      0 Mb      1 Mb   1683720000000  Wed May 10 12:00:00 UTC 2023
1 Mb      0 Mb      1 Mb   1683723600000  Wed May 10 13:00:00 UTC 2023
1 Mb      0 Mb      1 Mb   1683727200000  Wed May 10 14:00:00 UTC 2023
1 Mb      0 Mb      1 Mb   1683730800000  Wed May 10 15:00:00 UTC 2023
```

## 4. Send the Usage Record File to HPE

To share your usage record file and obtain a new activation code (every 30 days), complete the following steps:

1. Open a support case at <https://support.hpe.com> using the account you have on the HPE Support Center customer portal, and include the following information:
  - Fabric ID (cluster ID)
  - Current activation code
  - Usage record file
2. When HPE Support updates the ticket, go to your customer portal to get the new activation key.

## 5. Pay Your Monthly Invoice

Each month you must pay the HPE invoice before the 15-day grace period ends. Otherwise, the fabric can be disabled, as described in [Restoring a Disabled Fabric](#) on page 49.



## 6. Renew Your Activation

As long as you continue to provide usage records and pay your monthly invoice within the billing grace period, HPE will continue to provide an activation code that allows you to renew your activation.

After obtaining the new activation code from the customer portal, use the following `maprcli` command to renew your activation:

```
maprcli cluster startup set -activationkey <path-to-file>
```

For example:

```
maprcli cluster startup set -activationkey /tmp/
Renew_key_mycluster.text -is_file true -json
```

### Log Information

To view log information for service activation and billing, see the main CLDB log:

```
/opt/mapr/conf/cldb.log
```

## Restoring a Disabled Fabric

Describes how to obtain an activation key to restore a disabled fabric.

If you forget to pay your invoice, a fabric can be disabled by HPE. If your contract terms are not met, HPE activates a "kill switch" that causes the CLDBs to restart, eventually causing the fabric to enter a non-functional state.

To check the fabric status, use the `maprcli cluster services status` command. For example:

```
$ maprcli cluster services status -json
{
  "timestamp":1691100826677,
  "timeofday":"2023-08-03 03:13:46.677 GMT-0700 PM",
  "status":"OK",
  "total":1,
  "data":[
    {
      "status":"ENABLED"
    }
  ]
}
```

If you suspect that the fabric has been disabled, contact HPE Support. HPE Support can supply a special activation key that you can use to restore the fabric. With the activation key, you can restore the fabric.

## Displaying a maprcli Prompt

You can use `maprcli` commands to register the fabric and perform certain configuration tasks. The steps for displaying a `maprcli` prompt are the same for all cloud-based deployments but are different for on-premises deployments.

### maprcli Prompt for an On-Premises Fabric

To run `maprcli` commands, use an `ssh` connection to any node in the fabric.

### maprcli Prompt for an AWS, Azure, or GCP Fabric

See [SSH Access to a Cloud-Based Fabric](#) on page 126.

## SSO Using Keycloak

Describes how single sign-on (SSO) is implemented by using Keycloak.

### Keycloak Is Preinstalled and Preconfigured

Keycloak is the identity and access management (IAM) solution that provides single-sign-on (SSO) support for the Data Fabric. Starting with release 7.5.0, Keycloak is preinstalled and preconfigured whenever you create a new fabric.

During fabric creation, Keycloak is installed on all the nodes in the fabric. However, the Keycloak server is started on only one node. If new fabrics are created from the first fabric, Keycloak is installed on all the new fabric nodes, but the primary Keycloak node continues to serve the new fabrics.

At installation, Keycloak is preconfigured with users, groups, and roles that enable integration of Keycloak with the Data Fabric. The following table describes the preconfigured items:

Keycloak Preconfigured Items	How Many?	Names	Notes
Users	1	admin	Any additional users that are added must be created with <code>uid</code> and <code>gid</code> attributes, as described in <a href="#">Adding New Users to Keycloak</a> on page 54.
Groups	1	fabric-manager	Any additional groups that are added must be created with the <code>gidNumber</code> attribute, as described in <a href="#">Adding a Group to Keycloak</a> on page 60.
Roles	3	fabric-manager infrastructure-admin developer	These are the only supported roles. The <code>developer</code> role is sometimes referred to as the "fabric user" role.
Clients	1	edf-client	This is the dedicated client for the Data Fabric. In Keycloak, a client is an application or service that can request authentication for a user.

Keycloak installation also gives you access to the Keycloak admin portal.

### Accessing the Keycloak Administration Console

Describes how to start the Keycloak administration console so you can manage Keycloak and your SSO users.

If a new fabric has been created, you can access the Keycloak administration console by using these steps:

1. In a browser, specify the URL for your first fabric. This is the URL provided by the seed node procedure following successful fabric creation. For example:

```
https://<FQDN_for_host_node_first_fabric>
```

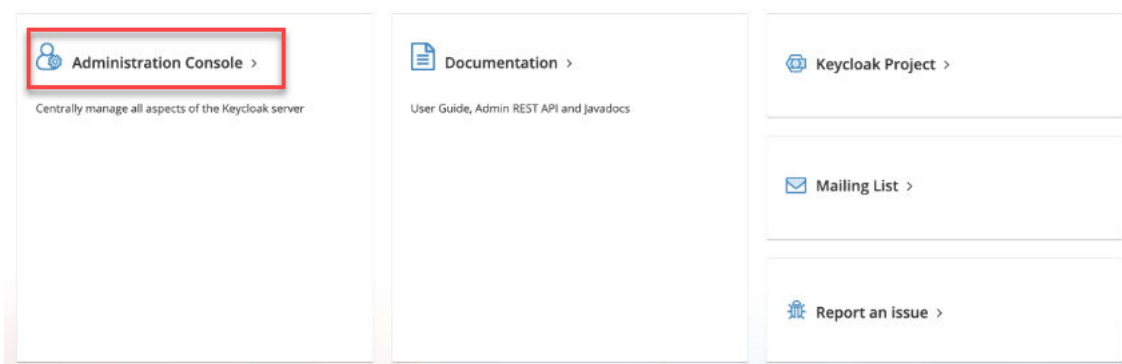


**NOTE:** If the cluster was upgraded from release 7.6.0 or a previous release, append the port (`:6443`) to the URL. Newer deployments use port 443, which does not need to be appended because it is the default HTTPS port.

2. Click **Administration Console**:



Welcome to **Keycloak**



The **Sign In** page is displayed:

3. Sign in using the default credentials:

**Username:** admin

**Password:** p@ssw0rd

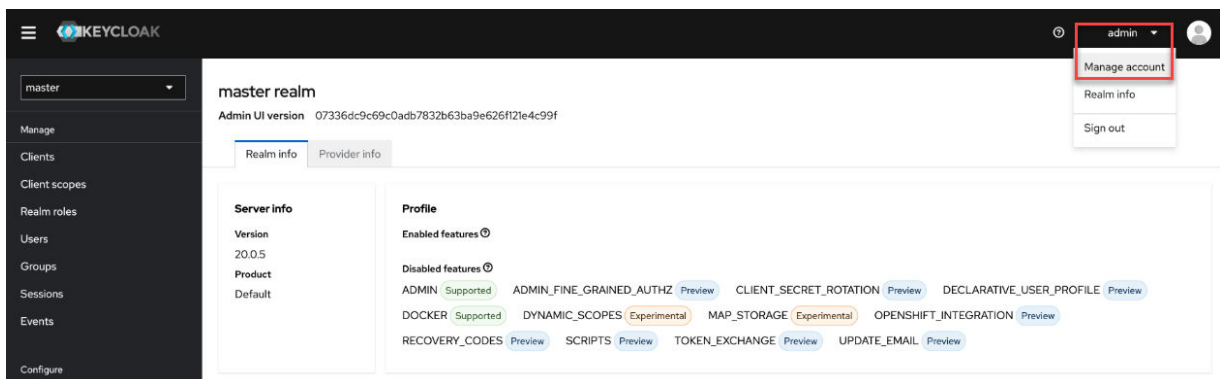
- !** **IMPORTANT:** HPE recommends that you change the password for the `admin` user soon after sign in. See [Changing the Keycloak admin Password](#) on page 51.

## Changing the Keycloak admin Password

Describes how to change the default Keycloak `admin` password to prevent unauthorized access to Keycloak and your Data Fabric user information.

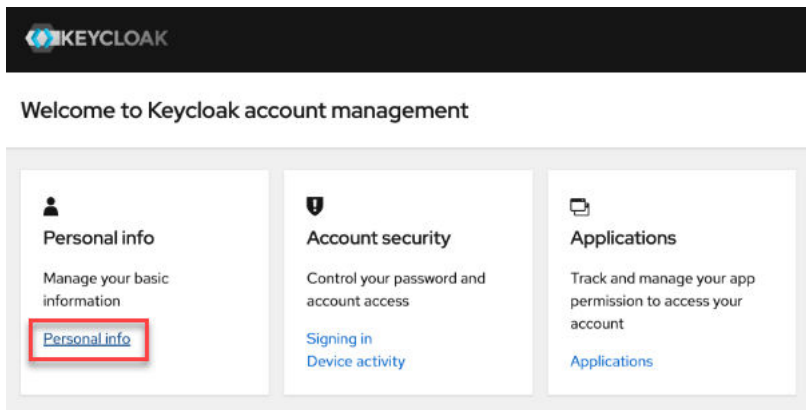
The default `admin` password provided in the bundled version of Keycloak is a well-known password that must be changed immediately after installation. Use these steps to change the password:

1. Sign in to the Keycloak administration console as described in [Accessing the Keycloak Administration Console](#) on page 50. The master realm information is displayed:
2. In the top right corner of the page, click the down arrow for the **admin** user, and select **Manage account**:



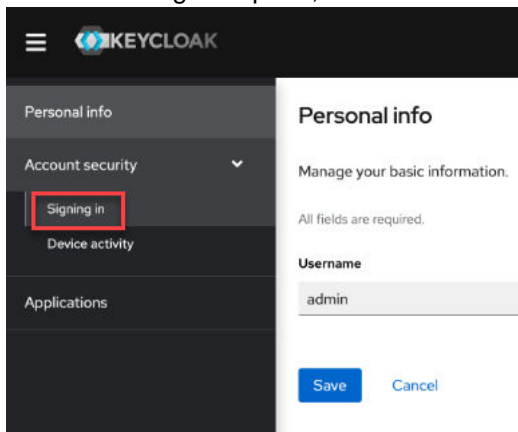
The account management information is displayed.

3. Click the **Personal Info**:



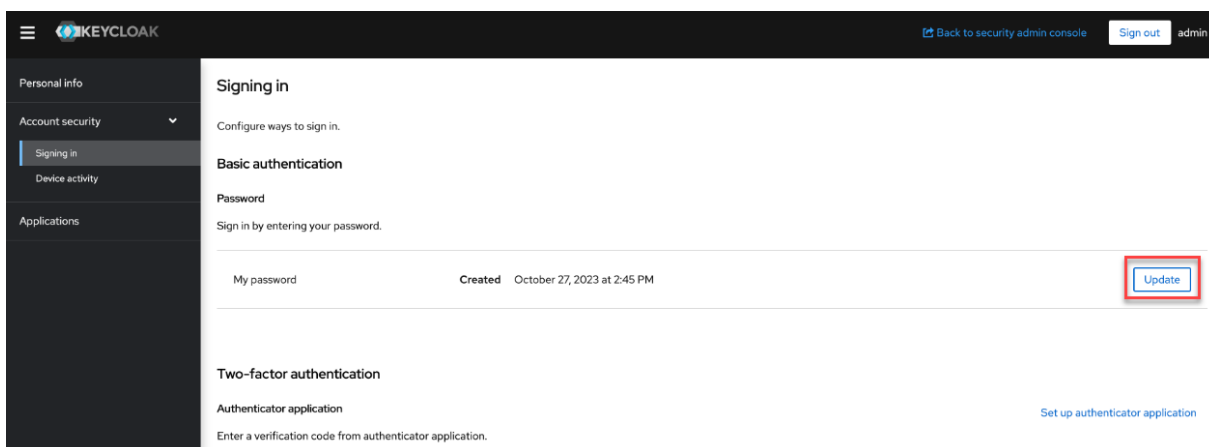
The **Personal Info** page is displayed.

4. In the left navigation pane, under **Account security**, click **Signing in**:



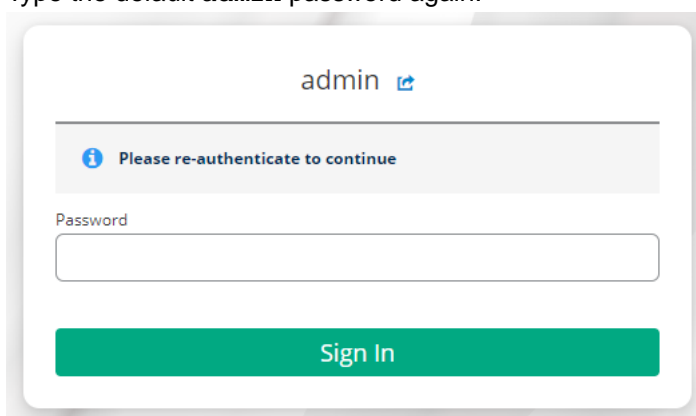
The **Signing in** page is displayed.

5. On the **Signing In** page, click **Update**:



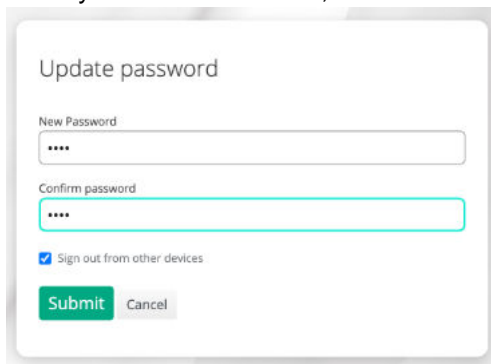
Keycloak asks you to re-authenticate.

6. Type the default `admin` password again:

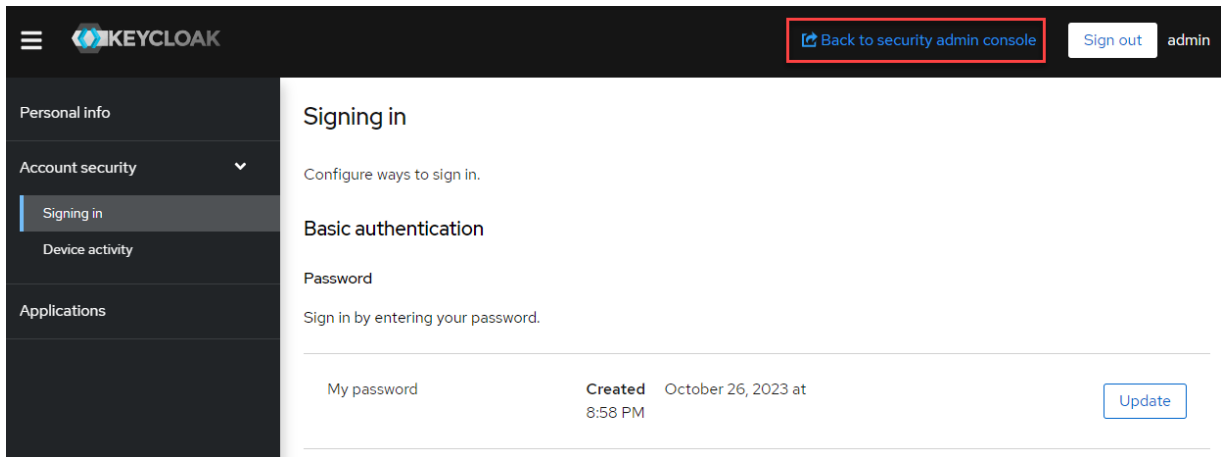


The **Update password** page is displayed.

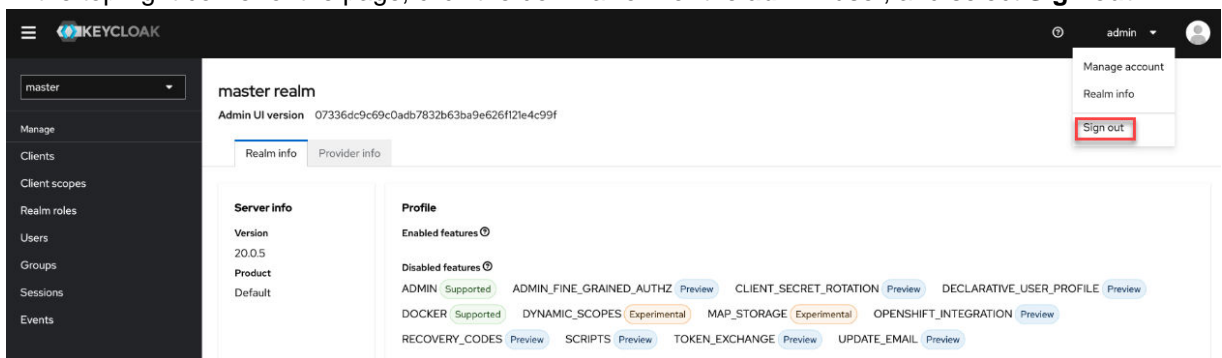
7. Enter your new credentials, and click **Submit**:



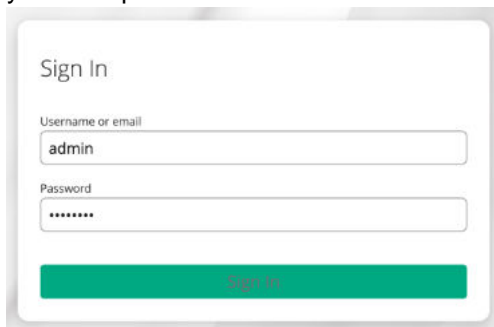
8. Click **Back to security admin console** to return to the administration console:



- In the top right corner of the page, click the down arrow for the `admin` user, and select **Sign out**:



- Repeat step 1, signing in to the Keycloak administration console as described in [Accessing the Keycloak Administration Console](#) on page 50. On the **Sign In** page, sign in as the `admin` user with your *new* password:

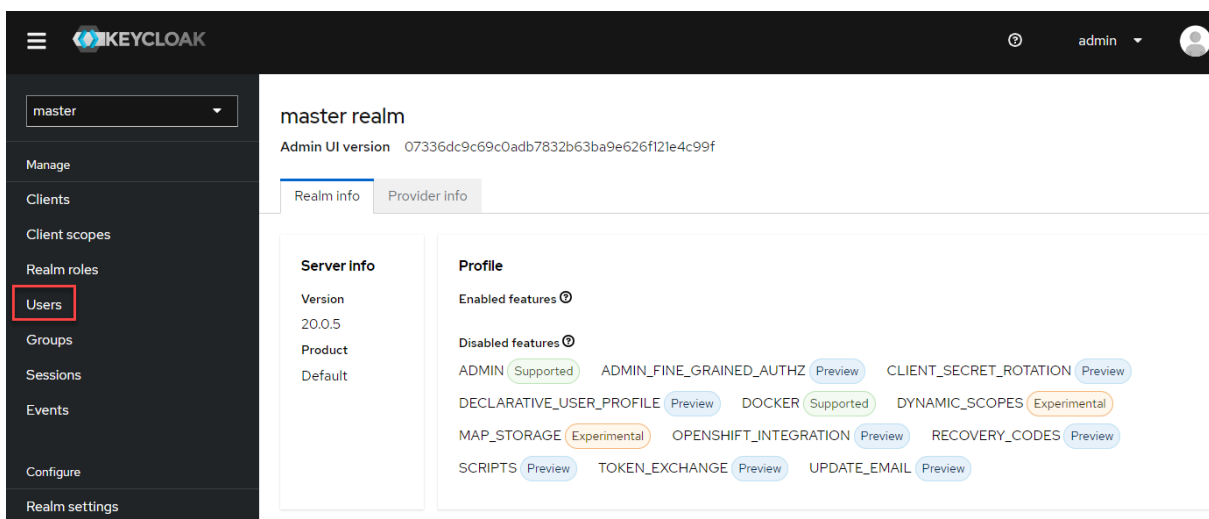


## Adding New Users to Keycloak

Describes how to add new users in Keycloak so you can use them to sign in to the Data Fabric UI.


By default, the Keycloak software provided with release 7.5.0 and later is preconfigured with only one user (the `admin` user). To add new users:

- Sign in to the Keycloak administration console as described in [Accessing the Keycloak Administration Console](#) on page 50. The master realm information is displayed.
- In the left navigation pane, click **Users**:



The **Users** page is displayed, showing the preconfigured `admin` user.

3. Click **Add user**:



The **Create user** page is displayed.

4. In the **Username\*** field, type the name of a new user, and click **Create**:

Users > Create user

### Create user

Enabled Action ▾

Username \*

Email

Email verified ⓘ  Off

First name

Last name

Required user actions  ⓘ

Groups ⓘ

The **User details** page for the new user is displayed.

5. Click the **Attributes** tab:

Users > User details

### catherine

Enabled Action ▾

Details **Attributes** Credentials Role mapping Groups Consents Identity provider links Sessions

ID \*

Created at \*

Username \*

Email

Email verified ⓘ  Off

First name

Last name

Required user actions  ⓘ

The **Attributes** page is displayed.



6. Enter `uid` and `gid` values for the new user:
  - a. In the **Key** field, type `uid`, then specify a `uid` value, such as `12345`, in the **Value** field.
  - b. Click **Add an attribute**.
  - c. In the second **Key** field, type `gid`, then specify a `gid` value, such as `12345`, in the **Value** field:

Users > User details

catherine Enabled Action

Details **Attributes** Credentials Role mapping Groups Consents Identity provider links Sessions

Key	Value
uid	12345
gid	12345

[+ Add an attribute](#)

**Save** [Revert](#)

7. Click **Save**.
8. Click the **Credentials** tab. The **Credentials** page shows **No credentials**.
9. Click **Set password**:

Users > User details

catherine Enabled Action

Details Attributes **Credentials** Role mapping Groups Consents Identity provider links Sessions

+

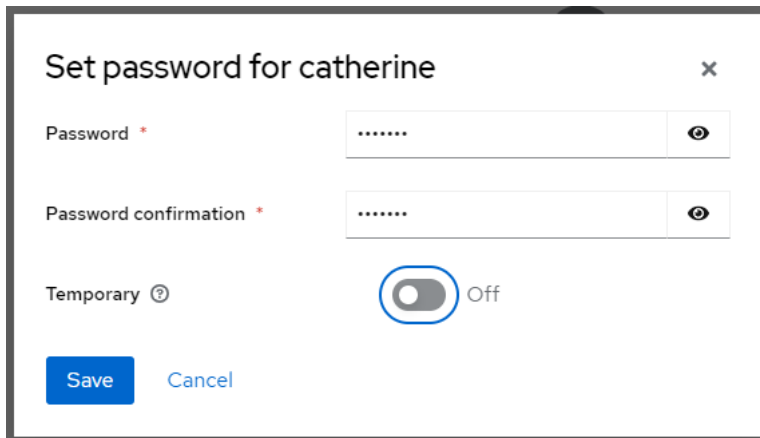
**No credentials**

This user does not have any credentials. You can set password for this user.

**Set password**

The **Set password for <new\_user>** dialog box is displayed.

10. Enter a password for the new user, and confirm the password.
11. Move the **Temporary** slider to the **Off** position:




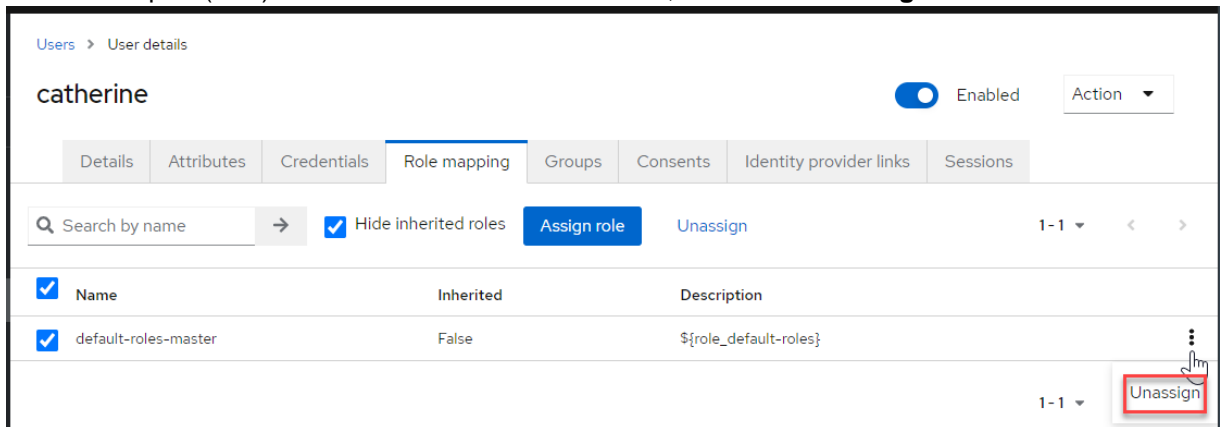
Set password for catherine

Password \*

Password confirmation \*

Temporary  Off

12. Click **Save**. The **Set password?** confirmation dialog box is displayed.
13. On the **Set password?** confirmation screen, click **Save password**. The **Credentials** tab of the **User details** page is displayed.
14. Click the **Role mapping** tab. The **Role mapping details** are displayed.
15. Click the **default-roles-master** role.
16. Click the ellipsis (  ) for the **default-roles-master** role, and select **Unassign**:



Users > User details

catherine Enabled Action

Details Attributes Credentials **Role mapping** Groups Consents Identity provider links Sessions

Search by name  Hide inherited roles  Unassign 1-1

<input checked="" type="checkbox"/>	Name	Inherited	Description
<input checked="" type="checkbox"/>	default-roles-master	False	`\${role_default-roles}`

1-1 Unassign

The **Remove mapping?** dialog box is displayed.

17. Click **Remove**. The **Role mapping details** page shows **No roles for this user**.
18. Click **Assign role**:

Users > User details

**catherine** Enabled Action

Details | Attributes | Credentials | **Role mapping** | Groups | Consents | Identity provider links | Sessions

**+**

**No roles for this user**

You haven't assigned any roles to this user. Assign a role to get started.

**Assign role**

The **Assign roles to <new\_user> account** is displayed.

19. In the **Name** column, click one of the preconfigured roles to assign it to the new user:

Assign roles to catherine account

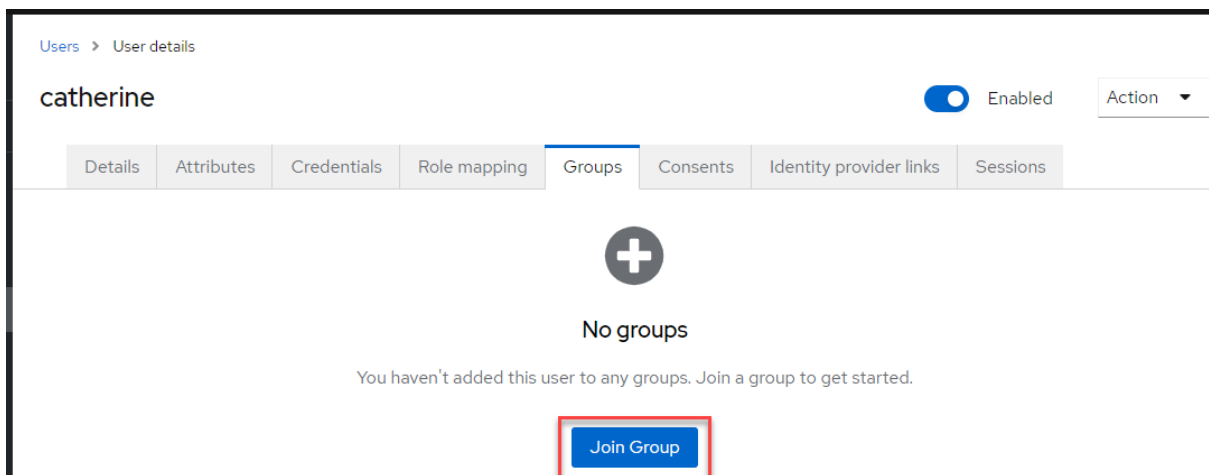
Filter by realm roles Search by role name 1-8

<input type="checkbox"/>	Name	Description
<input type="checkbox"/>	admin	\${role_admin}
<input type="checkbox"/>	create-realm	\${role_create-realm}
<input type="checkbox"/>	default-roles-master	\${role_default-roles}
<input checked="" type="checkbox"/>	developer	
<input type="checkbox"/>	fabric-manager	
<input type="checkbox"/>	infrastructure-admin	
<input type="checkbox"/>	offline_access	\${role_offline-access}
<input type="checkbox"/>	uma_authorization	\${role_uma_authorization}

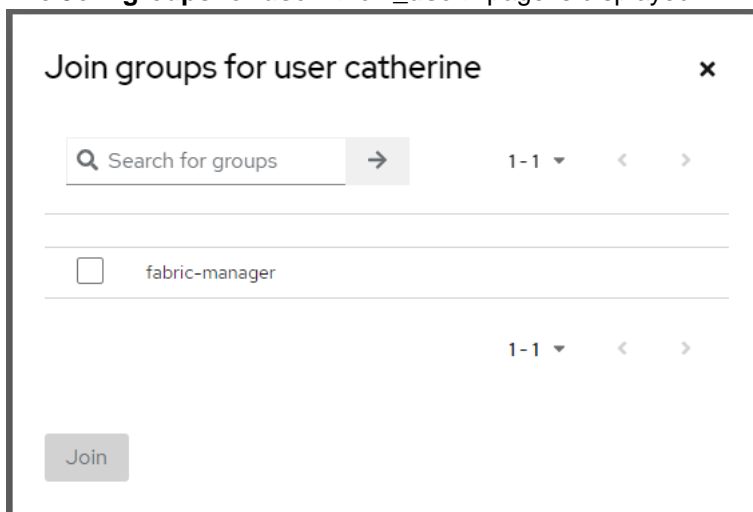
1-8

**Assign** Cancel


20. Click **Assign**. Next, you must assign the user to a group. Every user must belong to at least one group.
21. To add the user to a group, click the **Groups** tab. To add a new group, see [Adding a Group to Keycloak](#) on page 60.
22. Click **Join Group**:



The **Join groups for user <new\_user>** page is displayed:



23. To add the user to a group, click the check box for a group.
24. Click **Join**. The **Groups** page is displayed.
25. In the top right corner of the page, click the down arrow for the admin user, and select **Sign out**:



You can now sign in to the Data Fabric UI using the new user.

## Adding a Group to Keycloak

Describes how to add a Keycloak user group.

By default, the Keycloak software provided with release 7.5.0 and later is preconfigured with only one user group (the `fabric-manager` group). To add a new group:

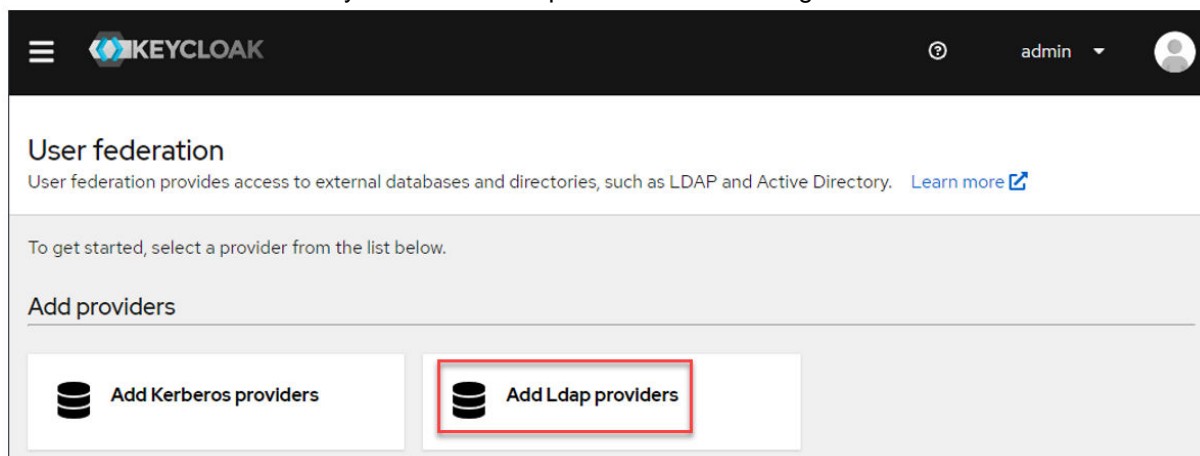
1. Sign in to the Keycloak administration console as described in [Accessing the Keycloak Administration Console](#) on page 50. The master realm information is displayed:
2. In the left navigation pane, click **Groups**.
3. Click **Create a group**.
4. Specify a name for the group, and click **Create**. The **Groups** page is displayed showing the new group. Click the link for the new group.
5. Click the **Attributes** tab.
6. In the **Key** field, type `gidNumber`, then specify a `gidNumber` value, such as `12345`, in the **Value** field.
7. Click **Save**.
8. In the left navigation pane, click **Users**.
9. From the list of users, click a user that you want to add to the new group.
10. Click **Join Group**.
11. Click the name of the group to which you want to add the user, and click **Join**.

## Integrating Your LDAP Directory with Keycloak

Keycloak can interface with an external LDAP directory so that LDAP users can access the Data Fabric UI.

To add an external LDAP provider in Keycloak:

1. Sign in to the Data Fabric UI, and switch to the Fabric manager experience.
2. Click **Security administration**.
3. In the **SSO setup** card, click **Configure LDAP**. The Data Fabric UI opens the Keycloak administration console to the screen where you can start the process of LDAP integration:



4. Click **Add Ldap providers**. The Keycloak **User federation** page is displayed.
5. Fill in the information for your LDAP provider. For field-specific information, click the online help icon (🔗) for the field. For Keycloak documentation, see [this page](#).

## Using LDAP Mappers

Describes how to use mappers to auto-populate Keycloak with the mandatory attributes it needs for users and groups to access the Data Fabric UI.

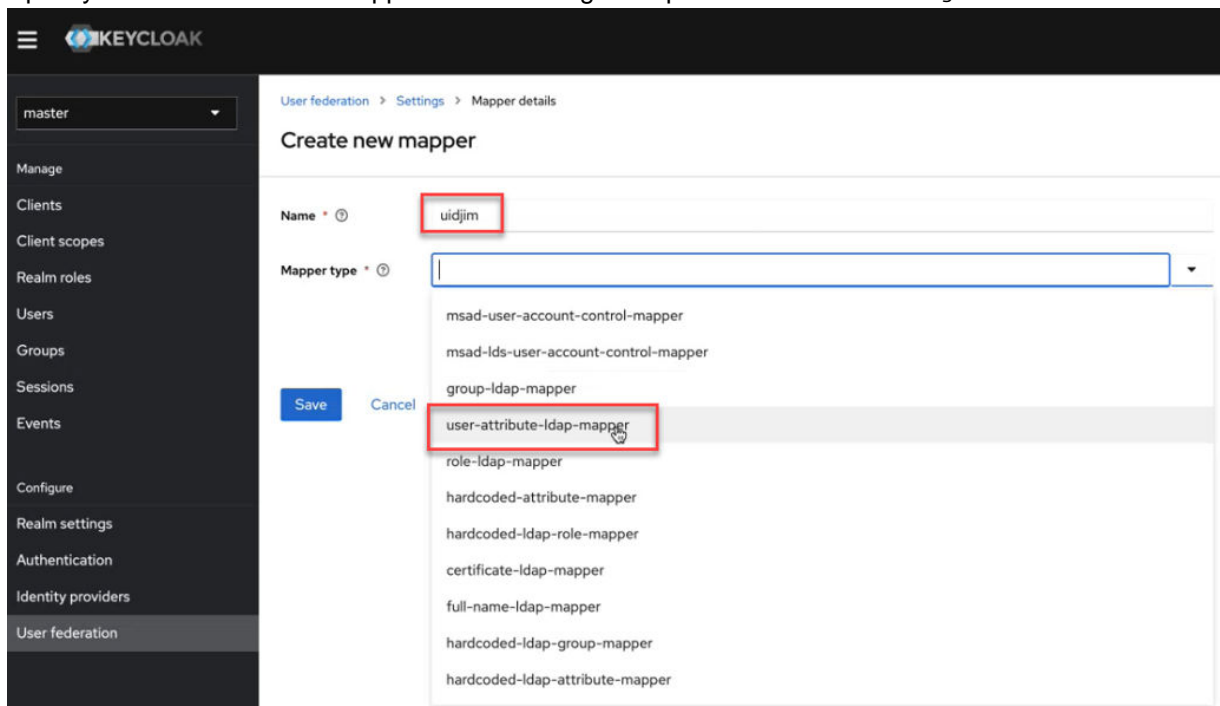
If you integrated your LDAP Directory with Keycloak as described in [Integrating Your LDAP Directory with Keycloak](#) on page 61, you must configure mappers to associate Keycloak user, role, and group attributes with your LDAP users. Three mappers need to be created:

- UID mapper
- GID mapper
- User group mapper

### Creating the UID Mapper

To create the UID mapper:

1. Sign in to the Keycloak administration console as described in [Accessing the Keycloak Administration Console](#) on page 50. The master realm information is displayed:
2. In the left-navigation pane, click **User federation**. The **User federation** screen appears.
3. Click the box for the LDAP provider that you configured in [Integrating Your LDAP Directory with Keycloak](#) on page 61. The **LDAP** screen appears.
4. Click the **Mappers** tab to display the current list of mappers.
5. To add a mapper, click **Add mapper**. The **Create new mapper** screen appears.
6. Specify a name for the UID mapper. The following example uses the name `uidjim`:



7. In the **Mapper type** field, click the down arrow, and select `user-attribute-ldap-mapper`. The **Mapper details** screen appears.
8. Fill out the UID mapper as follows:

The screenshot shows the 'Create new mapper' configuration page in Keycloak. The left sidebar contains navigation options like 'Manage', 'Clients', 'Client scopes', 'Realm roles', 'Users', 'Groups', 'Sessions', 'Events', 'Configure', 'Realm settings', 'Authentication', 'Identity providers', and 'User federation'. The main content area is titled 'Create new mapper' and includes the following fields and options:

- Name:** uidjim
- Mapper type:** user-attribute-ldap-mapper
- User Model Attribute:** uidjim
- LDAP Attribute:** uidNumber
- Read Only:** On
- Always Read Value From LDAP:** On
- Is Mandatory In LDAP:** Off
- Attribute default value:** (empty)
- Force a Default Value:** On
- Is Binary Attribute:** Off

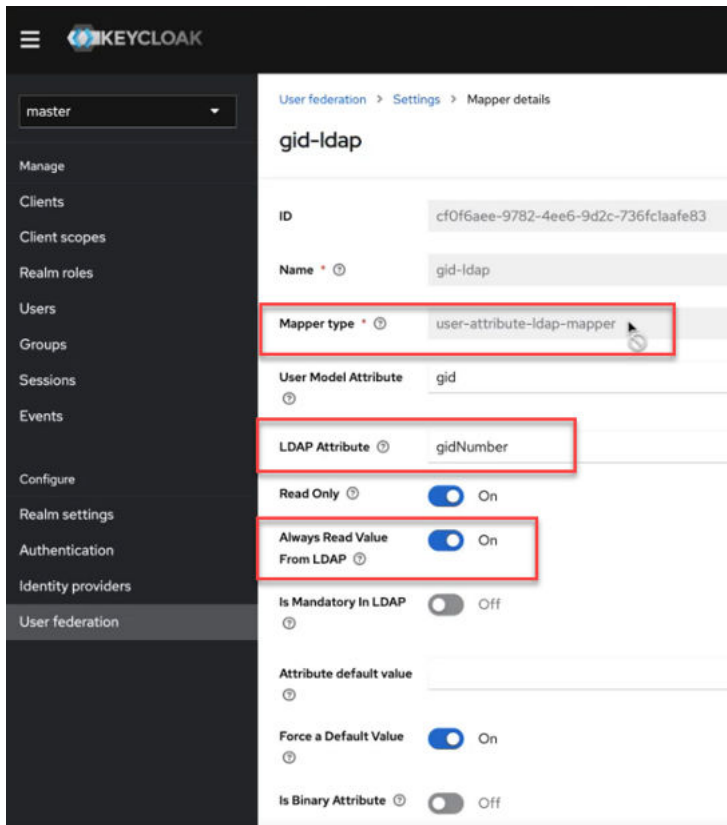
At the bottom, there are 'Save' and 'Cancel' buttons.

9. Click **Save**.

### Creating the GID Mapper

To create the GID mapper:

1. Sign in to the Keycloak administration console as described in [Accessing the Keycloak Administration Console](#) on page 50. The master realm information is displayed:
2. In the left-navigation pane, click **User federation**. The **User federation** screen appears.
3. Click the box for the LDAP provider that you configured in [Integrating Your LDAP Directory with Keycloak](#) on page 61. The **LDAP** screen appears.
4. Click the **Mappers** tab to display the current list of mappers.
5. To add a mapper, click **Add mapper**. The **Create new mapper** screen appears.
6. Specify a name for the GID mapper.
7. In the **Mapper type** field, click the down arrow, and select `user-attribute-ldap-mapper`. The **Mapper details** screen appears.
8. Fill out the UID mapper as follows:



9. Click **Save**.

### Creating the User Group Mapper

To create the User Group mapper:

1. Sign in to the Keycloak administration console as described in [Accessing the Keycloak Administration Console](#) on page 50. The master realm information is displayed:
2. In the left-navigation pane, click **Clients**. The **Clients list** tab appears.
3. Click the **edf-client** entry.
4. In the right pane, click the **Client scopes** tab.
5. Click the **edf-client-dedicated** entry.
6. Click **Add mapper > By configuration**. The **Configure a new mapper** screen appears:



**Configure a new mapper** x

Choose any of the mappings from this table

Name	Description
Allowed Web Origins	Adds all allowed web origins to the 'allowed-origins' claim in the token
Audience	Add specified audience to the audience (aud) field of token
Audience Resolve	Adds all client_ids of "allowed" clients to the audience field of the token. Allowed client means the client for which user has at least one client role
Authentication Context Class Reference (ACR)	Maps the achieved LoA (Level of Authentication) to the 'acr' claim of the token
Claims parameter Token	Claims specified by Claims parameter are put into tokens.
Claims parameter with value ID Token	Claims specified by Claims parameter with value are put into an ID token.
Group Membership	Map user group membership
Hardcoded claim	Hardcode a claim into the token.
Hardcoded Role	Hardcode a role into the access token.
Pairwise subject identifier	Calculates a pairwise subject identifier using a salted sha-256 hash. See OpenID Connect specification for more info about pairwise subject identifiers.
Role Name Mapper	Map an assigned role to a new name or position in the token.
User Address	Maps user address attributes (street, locality, region, postal_code, and country) to the OpenID Connect 'address' claim.
User Attribute	Map a custom user attribute to a token claim.
User Client Role	Map a user client role to a token claim.
User Property	Map a built in user property (email, firstName, lastName) to a token claim.
User Realm Role	Map a user realm role to a token claim.
User Session Note	Map a custom user session note to a token claim.

7. Click the **User Attribute** row. Selecting this row allows you to map a custom attribute to a token claim. The **Add mapper** screen appears.
8. Fill out the form like this, using a name that is appropriate for your installation:

master

Manage

Clients

Client scopes

Realm roles

Users

Groups

Sessions

Events

Configure

Realm settings

Authentication

Identity providers

User federation

Clients > Client details > Dedicated scopes > Mapper details

### Add mapper

If you want more fine-grain control, you can create protocol mapper on this client

Mapper type: User Attribute

Name: userGroupsJim

User Attribute: gidNumber

Token Claim Name: gids

Claim JSON Type: String

Add to ID token:  On

Add to access token:  On

Add to userinfo:  On

Multivalued:  On

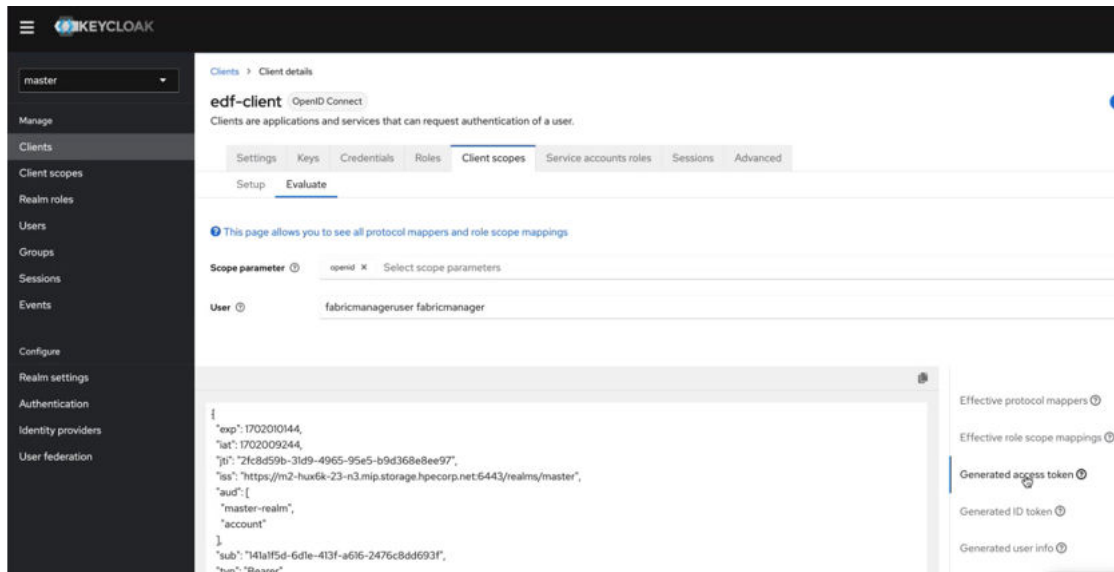
Aggregate attribute values:  On

Save Cancel

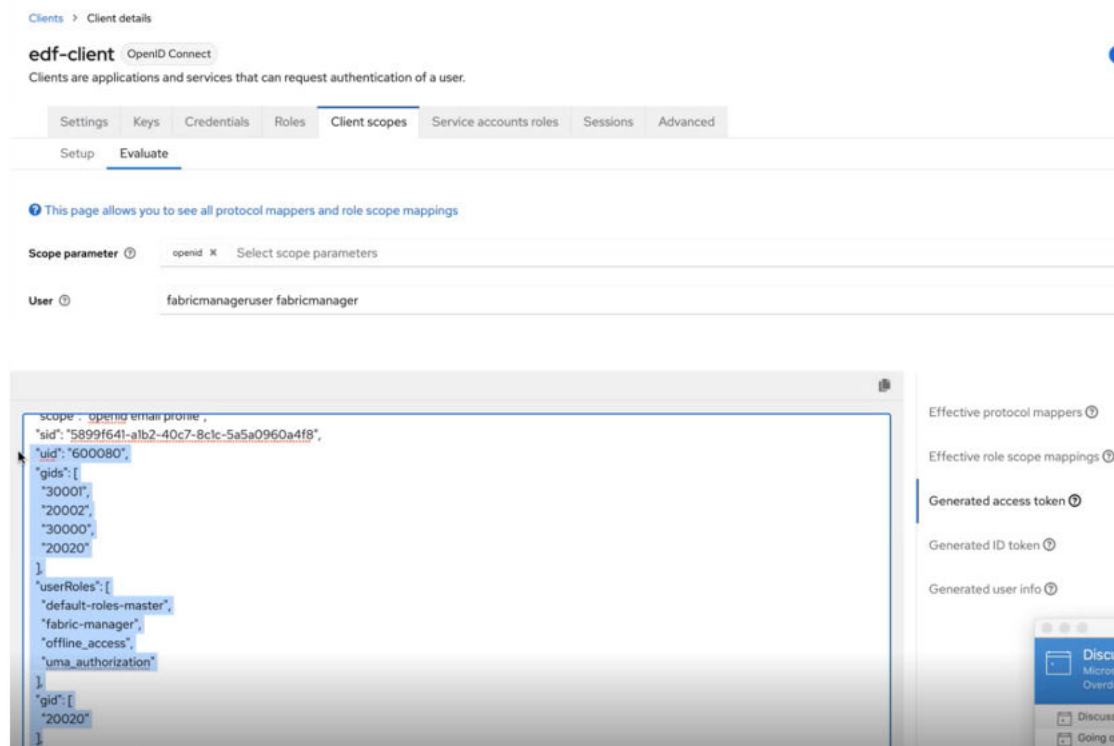
### Confirming that Required Attributes Are Part of the JWT Token for a User

To confirm the required attributes:

1. Sign in to the Keycloak administration console as described in [Accessing the Keycloak Administration Console](#) on page 50. The master realm information is displayed:
2. In the left-navigation pane, click **Clients**. The **Clients list** tab appears.
3. Click the **edf-client** entry.
4. In the right pane, click the **Client scopes** tab.
5. Click **Evaluate**.
6. In the **User** field, type the name of a user. For example, type the name of the fabric manager user:



7. Scroll down to check that the following four items are populated in the JWT token. If any of them are missing, there might be issues with user permissions:



## Completing SSO Setup Using the Data Fabric UI

Describes how to configure the HPE Ezmeral Data Fabric to work with your SSO server.

It is a best practice to complete the SSO setup task soon after your fabric is installed. Non-SSO users have limited capabilities in using the Data Fabric UI.

New installations of release 7.5.0 or later do not need to complete the SSO setup using the Data Fabric UI. But these instructions might be needed if you import a customer-managed cluster.

To complete the SSO setup task:

1. Sign in to the Data Fabric UI, and switch to the **Fabric manager experience**.

2. Click **Security administration**.
3. On the **SSO setup** card, click **Setup SSO**. The **SSO setup** form is displayed:

## Setup SSO Close ×

Provider  
keycloak

Provider URL\*

Client ID\*

Client secret\*

Certificate

Drag and drop
Select File

Setup
Cancel

4. Specify the following parameters:

Parameter	Description	Example
<b>Provider*</b>	Your SSO provider. Currently, Keycloak is the only supported provider.	Keycloak
<b>Provider URL*</b>	The URL of the SSO provider server.	https://myserver.keycloak.com/oauth2/default
<b>Client Secret*</b>	The key that is used to authenticate a client with the Keycloak server.	_Bfj1zbnnQNbNdprf0vnQDSyXcuzziMzyrbm0raB
<b>Client ID*</b>	An identifier that enables communication between the Data Fabric and the SSO provider.	0oa8m2onb7CAohGdW5d8
<b>Certificate</b>	The self-signed certificate from the SSO provider. Drag and drop the certificate into the box in the <b>SSO Setup</b> form. Or click <b>Select File</b> to navigate to the certificate file and select it.	<ssoprovider.crt>

5. Click **Create**. The webserver restarts automatically to ensure that correct authentication is enforced. After submitting, wait at least 15 minutes for the SSO configuration to be propagated. Then sign in again with your SSO credentials.

## Related maprcli Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` commands. These commands are provided for general reference. For more information, see [maprcli Commands in This Guide](#) on page 258.

- `cluster getssoconf`
- `cluster setssoconf`

## Resetting the SSO Configuration

Describes how to update your single sign-on (SSO) configuration information using the Data Fabric UI.

You must have fabric manager permissions to update the SSO configuration.

To view or change your SSO configuration, use these steps:

1. Sign in to the Data Fabric UI, and switch to the **Fabric manager experience**.
2. Click **Administration**.
3. On the **SSO setup** card, view your current SSO configuration details, including:
  - SSO provider
  - Provider URL
  - Client ID
4. To update or correct the information, click **Reset SSO Configuration**. A confirmation dialog box indicates that the operation cannot be undone.
5. Click **Reset**.

## Identifying All CLDB Nodes

Explains how you can identify all the CLDB nodes in an HPE Ezmeral Data Fabric.

Some procedures in this guide require you to find all the *container location database (CLDB)* nodes in your fabric.

### CLDB Nodes in Cloud-Based Fabrics

In AWS, Azure, and GCP deployments of the Data Fabric, ALL nodes are CLDB nodes because fabric creation configures every node as a CLDB node.

### CLDB Nodes in On-Premises Fabrics

In an on-premises deployment, not all nodes are necessarily CLDB nodes. To identify if a node is a CLDB node, ssh into the node as `root` or the Data Fabric `admin` user, and issue the following command:

```
maprcli node list -columns svc,ip
```

A node is a CLDB node if the service list includes `cldb` as one of the services:

```
maprcli node list -columns hostname,csvc
hostname                               ip
configuredservice

mynode67213.mycompany.net  10.163.167.293
keycloak,s3server,cldb,ezotelcol,nfs4,collectd,hoststats,data-access-gateway
,fileserver,mastgateway,opentsdb,gateway,apiserver
mynode67214.mycompany.net  10.163.167.294
```

```
keycloak,s3server,cldb,ezotelcol,nfs4,collectd,hoststats,data-access-gateway,
fileservers,mastgateway,opentsdb,gateway,apiserver
mynode67215.mycompany.net 10.163.167.295
keycloak,s3server,cldb,ezotelcol,nfs4,collectd,hoststats,fileservers,mastgate
way,opentsdb
```

To find the primary CLDB:

```
maprcli node cldbprimary
cldbprimary
ServerID: 5525564767900681920 HostName: mynode167215.mycompany.net
```

## Setting Up Clients

---

Summarizes the steps for enabling client communication with the HPE Ezmeral Data Fabric.

Clients allow hosts to communicate with the HPE Ezmeral Data Fabric. Set up clients and install client libraries to allow applications to access services on the HPE Ezmeral Data Fabric.

To run against the HPE Ezmeral Data Fabric platform, certain application types require a client and the following client libraries:

- HDFS API
- HBase API

## Installing Clients on a Linux Host

Describes how to install the client on a Linux host.

### Basic Steps for Client Installation

To install the Data Fabric client software on a Linux host that you want to communicate with the HPE Ezmeral Data Fabric:

1. Complete the steps in [Preparing to Install a Data Fabric Client](#) on page 70 later on this page.
2. Use one of these procedures to install the client on the Linux host:
  - [Installing the Data Fabric Client on RHEL](#) on page 75
  - [Installing the Data Fabric Client on SLES](#) on page 76
  - [Installing the Data Fabric Client on Ubuntu](#) on page 76
3. Perform the steps in [Installing Client Libraries](#) on page 77 to enable your fabric to communicate with the clients.

### Preparing to Install a Data Fabric Client

Before you install the Data Fabric client, perform the following steps:

1. **Verify that the operating system on the machine where you plan to install the client is supported.** For a list of operating systems that are compatible with the Data Fabric clients, see [Operating System Support Matrix](#) on page 259.
2. **Verify that the machine where you plan to install the client is not a fabric node.** The Data Fabric client is intended for use on a computer that has no other Data Fabric server software installed.
3. **Configure repositories for the client.** The client nodes also need to have the Data Fabric repositories configured in order to pull the client packages. See [Setting up the Data Fabric Repository](#) on page 71.

4. **Install the Data Fabric package key.** The package key must be installed before you can install Data Fabric packages. To install the package key, issue the command appropriate for your Linux distribution:



**IMPORTANT:** To access the Data Fabric internet repository, you must specify the email and token of an HPE Passport account. For more information, see [Accessing the HPE Ezmeral Token-Authenticated Internet Repository](#) on page 74.

- RHEL/Rocky/Oracle Enterprise Linux

```
wget --user=<email> --password=<token> -O /tmp/maprgpg.key -q https://
package.ezmeral.hpe.com/releases/pub/maprgpg.key && rpm --import /tmp/
maprgpg.key
wget --user=<email> --password=<token> -O /tmp/hpeezdf.pub -q https://
package.ezmeral.hpe.com/releases/pub/hpeezdf.pub && rpm --import /tmp/
hpeezdf.pub && gpg --import /tmp/hpeezdf.pub
```

- Ubuntu

```
wget --user=<email> --password=<token> -O /tmp/maprgpg.key -q https://
package.ezmeral.hpe.com/releases/pub/maprgpg.key && sudo apt-key
add /tmp/maprgpg.key
wget --user=<email> --password=<token> -O /tmp/gnugpg.key -q https://
package.ezmeral.hpe.com/releases/pub/gnugpg.key && sudo apt-key
add /tmp/gnugpg.key
```

For SLES only, you do not have to install the package key because `zypper` allows package installation with or without the key.

To install the client, obtain the Data Fabric packages for your operating system at <https://package.ezmeral.hpe.com/> and complete the installation steps described in one of the subsequent topics.

### Setting up the Data Fabric Repository

This section describes how to make packages available through the HPE Ezmeral Data Fabric repository.

The HPE Ezmeral Data Fabric repository on the internet provides all of the packages required to install a Data Fabric cluster using native tools such as:

- `yum` on RHEL
- `zypper` on SLES
- `apt-get` on Ubuntu

Installing from the internet repository is generally the easiest installation method, but requires the greatest amount of bandwidth. With this method, each node is connected to the internet to download the required packages.

Set up repositories by completing the steps for your RHEL, SLES, or Ubuntu distribution.

### Adding the Data Fabric Repository on RHEL

This section describes how to install the Data Fabric repository.

#### Procedure

1. Change to the `root` user or use `sudo`.

2. Create a text file called `maprtech.repo` in the `/etc/yum.repos.d/` directory with the following content, replacing `<version>` with the version of Data Fabric software that you want to install:



**IMPORTANT:** To access the Data Fabric internet repository, you must specify the user name (email) and token of an HPE Passport account. For more information, see [Accessing the HPE Ezmeral Token-Authenticated Internet Repository](#) on page 74.

```
[maprtech]
name=HPE Ezmeral Data Fabric
baseurl=https://package.ezmeral.hpe.com/releases/v<version>/redhat/
username=<email-address>
password=<token>
enabled=1
gpgcheck=1
protect=1

[maprecosystem]
name=HPE Ezmeral Data Fabric
baseurl=https://package.ezmeral.hpe.com/releases/MEP/MEP-<version>/redhat
username=<email-address>
password=<token>
enabled=1
gpgcheck=1
protect=1
```

3. If your connection to the Internet is through a proxy server, you must set the `http_proxy` environment variable before installation: You should also set the value for the `http_proxy` environment variable by adding the following section to the `/etc/yum.conf` file:

```
http_proxy=http://<host>:<port>
export http_proxy
```

```
proxy=http://<host>:<port>
proxy_username=<username>
proxy_password=<password>
```

### Adding the Data Fabric Repository on SLES

This section describes how to install the Data Fabric repository.

#### Procedure

1. Change to the `root` user or use `sudo`.



- Use the following command to add the repository for Data Fabric packages, replacing <version> with the version of Data Fabric software that you want to install:



**IMPORTANT:** For SLES distributions, if your user name is an email address that includes special characters – such as the @ symbol – you must URL encode the special characters so that the correct email address is passed to the authentication protocols in the repository. For most email addresses, changing the @ symbol to %40 is sufficient. For example:

**Unencoded email address:** jane.smith@company.com

**URL encoded email address:** jane.smith%40company.com

To encode other special characters, see "URL Encoded Emails" at [HPE Software Delivery Repository](#).

```
zypper ar https://<email>:<token>@package.ezmeral.hpe.com/releases/
v<version>/suse/ maprtech
```

- Use the following command to add the repository for ecosystem packages:

```
zypper ar https://<email>:<token>@package.ezmeral.hpe.com/releases/MEP/
MEP-<version>/suse/ maprecosystem
```

- If your connection to the Internet is through a proxy server, you must set the `http_proxy` environment variable before installation:

```
http_proxy=http://<host>:<port>
export http_proxy
```

- Update the system package index by running the following command:

```
zypper refresh
```

- Data Fabric packages require a compatibility package in order to install and run on SLES. Execute the following command to install the SLES compatibility package:

```
zypper install mapr-compat-suse
```

### *Installing sshpass*

#### **About this task**

Before installing a cluster on a SLES image, you must run the following command to install `sshpass`:

```
zypper --non-interactive -q --no-gpg-checks -p http://download.opensuse.org/
distribution/leap/42.3/repo/oss/ install sshpass
```

#### **Adding the Data Fabric Repository on Ubuntu**

This section describes how to install the Data Fabric repository.

#### **Procedure**

- Change to the `root` user or use `sudo`.

2. Create the following file:

```
# cat /etc/apt/auth.conf.d/package.ezmeral.hpe.com.conf
machine package.ezmeral.hpe.com
login <HPE-Passport-email>
password <HPE-Passport-token>
```



**IMPORTANT:** To access the Data Fabric internet repository, you must specify the email and token of an HPE Passport account. For more information, see [Accessing the HPE Ezmeral Token-Authenticated Internet Repository](#) on page 74.

3. Add the following lines to `/etc/apt/sources.list`, replacing `<version>` with the version of Data Fabric software that you want to install:

```
deb https://package.ezmeral.hpe.com/releases/v<version>/ubuntu/ binary
bionic
deb https://package.ezmeral.hpe.com/releases/MEP/MEP-<version>/ubuntu/
binary bionic
```

4. Update the package indexes:

```
apt-get update
```

5. If your connection to the Internet is through a proxy server, add the following lines to `/etc/apt/apt.conf`:

```
Acquire
{
  Retries "0";
  HTTP
  {
    Proxy "http://<user>:<password>@<host>:<port>";
  };
};
```

### Accessing the HPE Ezmeral Token-Authenticated Internet Repository

Describes special considerations for using the token-authenticated internet repository for Data Fabric software and the ecosystem components.

#### Accessing the Token-Authenticated Repository

Using a browser to access the new token-authenticated package repository requires you to supply the email address associated with your HPE account and a token. Use these steps:

1. Navigate to the repository at <https://package.ezmeral.hpe.com/>.

The authorization dialog box is displayed:

2. In the **Username** field, paste the email address for your HPE Passport account. To obtain an HPE Passport Account, see [Obtaining an HPE Passport Account](#) on page 75.
3. In the **Password** field, paste a token. To obtain a token, see [Obtaining a Token](#) on page 75.
4. Click **Sign in**.

### Format for Passing an HPE User Name and Token to the Repository

Any files or scripts that point to the new Data Fabric internet repository must include the email address and token associated with a valid HPE account expressed in the following format:

```
https://<email-address>:<token>@package.ezmeral.hpe.com/
```

### Examples for Accessing the Repository

In examples that require you to run Linux commands that point to the repository, this guide shows the format that is needed for including the user name and password. For example, to use a `wget` command with the new repository, you must add the email address and token as follows:

```
wget --user=jane.smith@company.com --password=<token> https://  
package.ezmeral.hpe.com/releases/installer/mapr-setup.sh -P /tmp
```

Depending on the Linux distribution, other formats might be needed.

### Obtaining an HPE Passport Account

An HPE Passport account is required to obtain support for Data Fabric products and gives you access to important HPE services. To obtain an HPE Passport account, visit the [MY HPE SOFTWARE CENTER](#) and click **Sign In** to create a new account.

When you fill in information about your account, be sure to complete ALL of the fields (even fields that are not required). Leaving some fields blank can cause issues when you later try to access HPE repositories.

### Obtaining a Token

A token associated with your HPE Passport account is required to obtain access to the HPE Ezmeral internet repositories. You can create a new token at any time by using the following steps. A token created in this way does not expire. The token remains valid even after you create a new token.

To create a token for your HPE Passport account:

1. Visit the [HPE Support Center User Token page](#).
2. Sign in if needed using your HPE Passport user ID and password.

### Installing the Data Fabric Client on RHEL

This section describes how to install the Data Fabric client on Red Hat Enterprise Linux (RHEL).

These steps assume that you have already installed the package key as described in [Installing Clients on a Linux Host](#) on page 70 and set up a Data Fabric repository as described in [Setting up the Data Fabric Repository](#) on page 71.

1. Remove any previous Data Fabric software. You can use `rpm -qa | grep mapr` to get a list of installed Data Fabric packages, then type the packages separated by spaces after the `rpm -e` command:

```
rpm -qa | grep mapr  
rpm -e mapr-fileserver mapr-core
```

2. Install the client for your target architecture:

```
yum install mapr-edf-clients
```

3. Open the Data Fabric UI to complete the configuration, as described in [Installing Client Libraries](#) on page 77.

### Installing the Data Fabric Client on SLES

This section describes how to install the Data Fabric Client on SLES.

1. Remove any previous Data Fabric software. You can use `rpm -qa | grep mapr` to get a list of installed Data Fabric packages:

```
rpm -qa | grep mapr
```

Then type the package names separated by spaces after the `zypper rm` command. For example:

```
zypper rm mapr-fileserver mapr-core
```

2. Run the following command to install the Data Fabric client:

```
zypper install mapr-edf-clients
```

3. Open the Data Fabric UI to complete the configuration, as described in [Installing Client Libraries](#) on page 77.

### Installing the Data Fabric Client on Ubuntu

This section describes how to install the Data Fabric client on Ubuntu.

1. Remove any previous Data Fabric client software. You can use `dpkg --get-selections | grep mapr` to get a list of installed Data Fabric packages. Then type the packages separated by spaces after the `dpkg -r` command. For example:

```
dpkg -r mapr-core mapr-fileserver
```

2. Update your Ubuntu repositories. For example:

```
apt-get update
```

3. Make sure the client is running JDK 11 or later:

```
$ echo $JAVA_HOME
/Library/Java/JavaVirtualMachines/jdk-11.0.1.jdk/Contents/Home
$ /Library/Java/JavaVirtualMachines/jdk-11.0.1.jdk/Contents/Home/bin/
java -version
openjdk version "11.0.1" 2018-10-16
OpenJDK Runtime Environment 18.9 (build 11.0.1+13)
OpenJDK 64-Bit Server VM 18.9 (build 11.0.1+13, mixed mode)
```

4. Run the following command to install the Data Fabric client:


```
apt-get install mapr-edf-clients
```

5. Open the Data Fabric UI to complete the configuration, as described in [Installing Client Libraries](#) on page 77.

## Installing Client Libraries

Describes how to install the client libraries to enable communication between your Linux hosts and the HPE Ezmeral Data Fabric.

To install client libraries on a fabric:

1. Install the client packages using the steps described in [Installing Clients on a Linux Host](#) on page 70.
2. Sign in to the Data Fabric UI.
3. Select the **Fabric user** view on the **Home** page.
4. On the **Home** page, click the ellipsis (  ) in the **Action** column of the fabric for which you want to install the client libraries.

Alternatively, navigate to the **Fabric Details** page, open the **Actions** dropdown menu, and select **Client library**.

5. Click the **Client library** option. The **Client library** side drawer opens.
6. Download the `config.tar` and the `jwt_tokens.tar.gz` files listed in the **Client library** side drawer. These files include information needed to set up the client libraries for your fabric.
7. Copy the downloaded files to the client machine. On the client machine, perform the following steps.
8. Run the command to extract the setup:

```
tar xf config.tar --directory /opt/mapr
```

9. Run the command to extract the JWT tokens:

```
tar xf jwt_tokens.tar.gz --directory /root
```

10. Run the commands to export the JWT tokens:

```
export MAPR_JWT_TOKEN_LOCATION="/root/jwt_access"
```

```
export MAPR_REFRESH_TOKEN_LOCATION="/root/jwt_refresh"
```

11. Run the command to configure your client libraries:

```
/opt/mapr/server/configure.sh -R
```

12. Run the command to test that your client libraries are set up correctly:

```
hadoop fs -ls /
```

If the client libraries are not set up correctly, the command returns an error message.

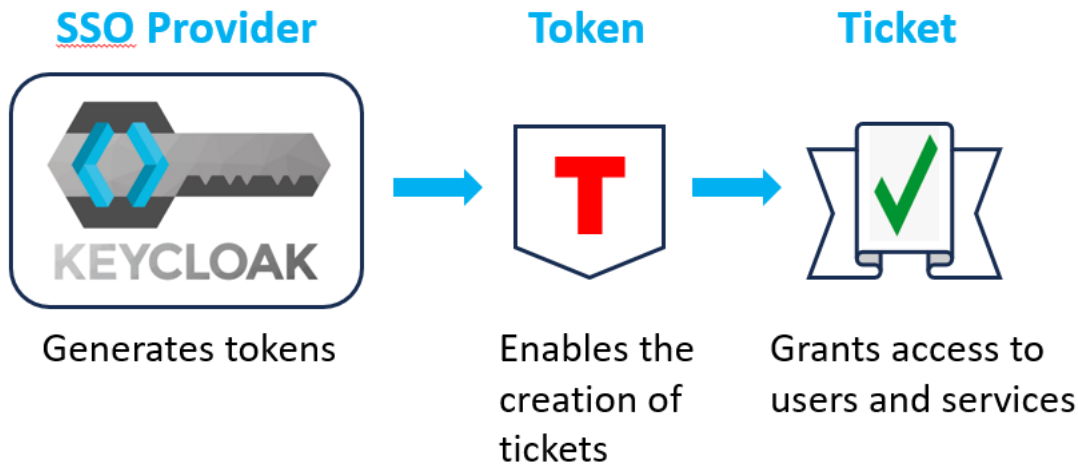
## About Access and Refresh Tokens

Describes how to use the downloadable tokens provided for client access.

### How Clients Use Tokens

In Data Fabric deployments where SSO is configured, you must provide an SSO user name and password for access to a fabric. Clients that aren't aware of SSO user names and passwords can gain access to RPC communications by using JSON web tokens (JWTs). A JSON Web Token (JWT) is a method for securely transmitting information between services in a computing system.

In a Data Fabric deployment, the Keycloak SSO provider can generate tokens when a user supplies an SSO user name and password. The tokens enable the creation of short-lived Data Fabric tickets that facilitate communication with the file system.



Issuing commands from any of the following command line interfaces (CLIs) or clients requires the user or client to have a valid ticket.

- `maprcli`
- `hadoop`
- `mc`
- `fuse (service start)`
- `loopback nfs (service start)`

The ticket allows the command line to connect to the CLDB service. To facilitate the process, you must obtain a token from the SSO provider and place it in a secure location that can be specified in an environment variable. Exporting the environment variable creates a temporary ticket, enabling the CLI to talk to the CLDB server. This method permits the use of any command without a password for the duration of the ticket.

For more information about Data Fabric tickets, see [Managing Tickets](#).

### Downloading the Tokens

In the Data Fabric UI, you use the **Client library** command to download the tokens. The tokens are contained in the `jwt_tokens.tar.gz` file, which are extracted to the client. The `jwt_tokens.tar.gz` file contains the following token files:

Token Type	File	Function
Access	<code>jwt_access</code>	Encapsulates the user's authentication information within the fabric.
Refresh	<code>jwt_refresh</code>	Enables the creation of a new access token when the current access token expires.

For more information about downloading the tokens, see [Installing Client Libraries](#) on page 77.

## Exporting the Tokens

To enable a client to use the tokens, you must export the path to each token. This must be done each time you establish a host session. To export the paths:

Client	To export the tokens . . .
Hadoop	Use these commands: <pre>export MAPR_JWT_TOKEN_LOCATION="/root/jwt_access" export MAPR_REFRESH_TOKEN_LOCATION="/root/jwt_refresh"</pre>
Fuse POSIX	Add the export paths shown for the Hadoop client in the first row of this table to the top of the following file: <pre>/opt/mapr/initscripts/mapr-posix-client-basic</pre>
Loopback NFS	Add the export paths shown for the Hadoop client in the first row of this table to the top of the following file: <pre>/usr/local/mapr-loopbacknfs/initscripts/mapr-loopbacknfs</pre>

Alternatively, you can add the tokens to the `core-site.xml` file. Adding them to `core-site.xml` file causes the fabric to use the designated tokens *every time you log on*. To add the tokens, specify the following property in the `core-site.xml`:

```
<property>
  <name>fs.mapr.sso.tokenpath</name>
  <value>/root/jwt_access</value>
</property>
```

## Token and Ticket Expiration and Renewal

Tokens and tickets expire after a short time. By default, Keycloak-generated tokens expire after two (2) hours. Short-lived tickets expire after 20 minutes.

If an access token expires or becomes invalid, the client application can use a refresh token to obtain a new access token without requiring the user to re-authenticate. The client application sends Keycloak a token-refresh request along with the current refresh token. Keycloak validates the refresh token and issues a new access token. This automatic-refresh mechanism repeats itself to allow client jobs to run for days or weeks as long as the tokens remain valid.

## Changing Token and Ticket Durations

You can change the valid duration of tokens and tickets. Note that a ticket is valid for no more than 20 minutes or the expiry time of its associated access token, *whichever is lower*. Thus, if a ticket expiry time is set for 20 minutes and the associated access token is valid only for 10 minutes, the ticket will be valid for only 10 minutes.



**CAUTION:** Setting long lifetimes for tokens or tickets can introduce a considerable security risk. Hewlett Packard Enterprise recommends finding a balance between security and usability and, whenever possible, erring on the side of security in your use of tokens and tickets.

To check or change the expiry setting for short-lived tickets, see [Checking and Changing the Temporary Ticket Duration](#).

To change the expiration setting for a token, you must be the fabric manager and have access to the Keycloak UI.

### Access Token Expiry

You can configure the access token expiry time at the realm level or at the client level.

1. Log in to the Keycloak admin console. See [Accessing the Keycloak Administration Console](#) on page 50.
2. Select the realm for which you want to configure the access token expiry time.
3. Go to the **Realm Settings > Tokens** tab.
4. In the **Access Token Lifespan** field, specify the desired expiration time for the access tokens in hours, minutes, or days.
5. Save your changes.

### Refresh Token Expiry

You typically configure the refresh token expiry time at the realm level:

1. Log in to the Keycloak admin console. See [Accessing the Keycloak Administration Console](#) on page 50.
2. Select the realm for which you want to configure the access token expiry time.
3. Go to the **Realm Settings > Sessions** tab.
4. In the **SSO Session Max** field, specify the desired maximum lifespan for refresh tokens in minutes, hours, or days.
5. Save your changes.

## Upgrade

---

This section contains information that describes how to upgrade the HPE Ezmeral Data Fabric as-a-service platform.

### Upgrading a Data Fabric

Describes how to upgrade your fabric and what to know before upgrading.

Upgrading to a newer Data Fabric release is supported for both on-premises and cloud-based fabrics.

#### Before Upgrading

Here are some things you should know about the upgrade process:

- Upgrades do not happen automatically. The fabric manager must initiate an upgrade.
- It is a best practice to upgrade to the latest Data Fabric software as soon as you are prompted. Upgrades provide new features and fix defects that can affect your ability to use the Data Fabric.
- You can upgrade only to the latest currently supported Data Fabric version. For example, if your fabric is running release 7.4.0 and an upgrade to release 7.7.0 is available, you will have the option to upgrade to release 7.7.0 but not to release 7.5.0 or 7.6.x. For a list of supported releases, see [Release History](#) on page 255.



- The upgrade proceeds one node at a time as a "rolling" upgrade. This means that the Data Fabric UI is running on a node that eventually will get shut down and upgraded. At that point, the UI will become temporarily unusable.
- The upgrade process upgrades only the packages installed on fabric nodes. Upgrading does not change the Linux OS running on the nodes.
- Upgrades initiated through the Data Fabric UI are supported only for fabrics that have a consumption-based license. A customer-managed Data Fabric with a term-based license cannot be upgraded in the same way.
- If your fabric has fewer than three nodes, you can expect the entire fabric to be unavailable for the client user during the upgrade. If your fabric has three or more nodes, the fabric will still work during the upgrade, and all client operations will continue to work. However, the upgrade of the node hosting the Data Fabric UI will cause the UI to go offline intermittently until the node is upgraded.

### Special Considerations for Upgrading to Release 7.7.0 and Later

In release 7.7.0, the port for accessing the Keycloak Administration Console changed. Newly created cloud or on-premises fabrics running releases 7.7.0 and later use port 443 to access the console. Earlier releases used port 6443. However, this port change does not apply to fabrics that are upgraded from releases 7.5.0 or 7.6.x to releases 7.7.0 or later. Fabrics upgraded from releases 7.5.0 or 7.6.x can continue to use port 6443.

In releases 7.7.0 and later, you do not need to specify port 443 explicitly in the URL. If you specify the webserver node, HTTPS connects by default to port 443. See [Accessing the Keycloak Administration Console](#) on page 50.

### Special Considerations for Upgrading from Release 7.4.0

Note these special considerations for upgrading from release 7.4.0 to a later Data Fabric release:

- You must initiate the upgrade from the primary fabric in the global namespace. If you initiate the upgrade from a fabric that is not the primary, you might not be able to monitor the upgrade status while the upgrade is in progress. To identify the primary fabric, issue the following `maprccli` command:

```
maprccli clustergroup getcgtable -showprimary true -json
```

To access a `maprccli` prompt, see [Displaying a maprccli Prompt](#) on page 49.

- Because of a known issue (EZINDFAAS-581), upgrades from release 7.4.0 to 7.5.0 can fail because the `keypair.pem` file has the wrong permissions. This issue affects upgrades on AWS, Azure, and GCP, but does NOT affect upgrades for on-premises fabrics.

To prevent the issue, you must change the file permissions BEFORE upgrading from release 7.4.0 to 7.5.0. Use the following steps:

1. Find the installer node on your cloud fabric. The installer node has the `keypair.pem` file. When you create a fabric using the [seed node deployment steps](#), the seed node displays the endpoint of the installer node.

You can also identify the installer node because it is the node that contains the `tmp/terraform_output.json` file in the deployment directory.

2. Use the steps in [SSH Access to a Cloud-Based Fabric](#) on page 126 to ssh to the installer node and display a Linux prompt.

3. Change to the directory containing the `keypair.pem` file. For example:

```
cd /opt/mapr/installer/ezdfaas/deployments/  
<name_of_cloud_fabric>/infrastructure/terraform/[aws|azure|gcp]/  
<name_of_new_cloud_fabric>-keypair.pem
```


4. Use the `chmod` command to change the file permissions:

```
chmod -R 0400 <pem-file-name>
```

5. When you are ready, initiate the upgrade from release 7.4.0 to 7.5.0.

### Checking to See if an Upgrade Is Supported for Your Fabric

The Data Fabric UI prompts you when an upgrade is supported for the current version of your software. For example:



If you are not sure that an upgrade is supported, you can check for supported upgrades by using the following steps:

1. Sign in to the Data Fabric UI as a fabric manager.
2. Click the **Fabric manager experience**.
3. Click the **Global namespace** button. The UI displays the fabrics in your global namespace.
4. See the **Version** column in the table.

**IMPORTANT:** Version information for a fabric is displayed only if you have obtained a consumption-based license and [added the activation key](#).

If the **Version** column includes a prompt to upgrade to a newer release, an upgrade is supported for the fabric. For example:

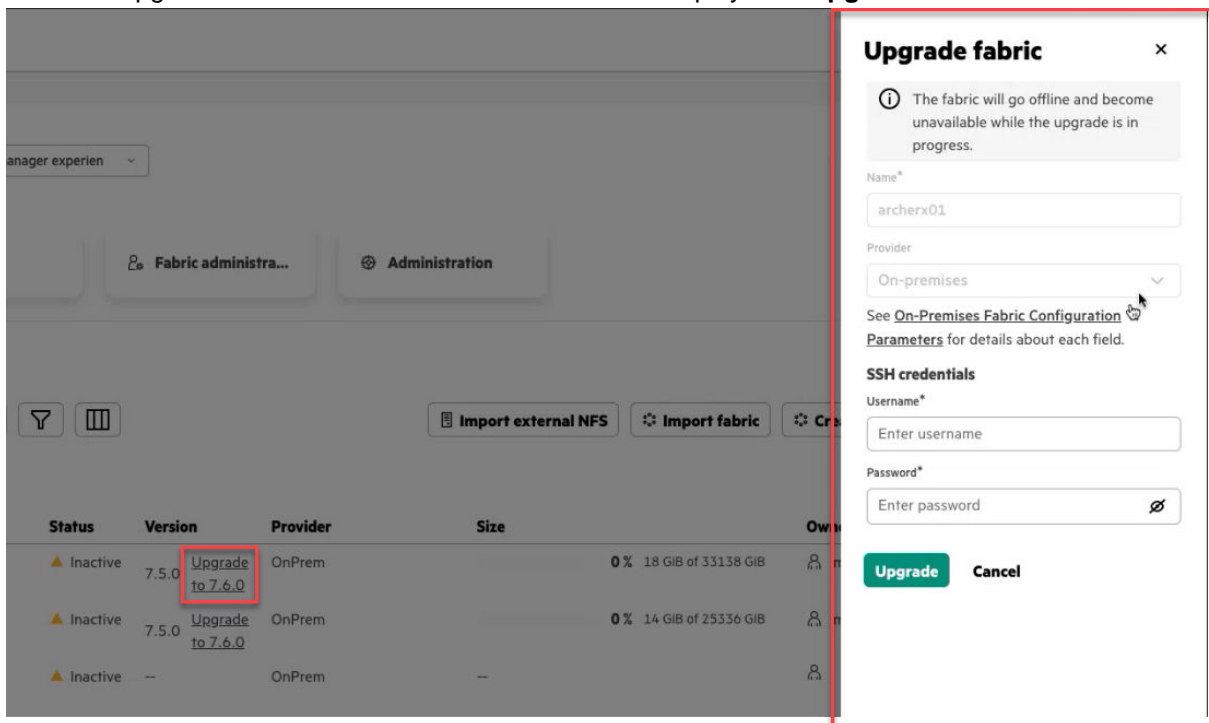
The screenshot shows the HPE Ezmeral Data Fabric UI interface. At the top, there is a notification banner: "An upgrade is available for some of your fabrics. Check upgrade". Below this, the user is logged in as "mapruser82" in the "Fabric manager experience". The main navigation includes "Global namespace", "Monitoring", "Fabric administra...", and "Administration". The "Global namespace" section is active, showing a search bar, filter, and list view icons, along with "Import external NFS" and "Import fabric" buttons. A table displays three fabric resources:

Resource Name	Type	Status	Version	Provider	Size
archerx01	Fabric	Inactive	7.5.0 Upgrade to 7.6.0	OnPrem	0% 18 GiB of 33138 GiB
archerx02	Fabric	Inactive	7.5.0 Upgrade to 7.6.0	OnPrem	0% 14 GiB of 25336 GiB
archerx03	Fabric	Inactive	--	OnPrem	--

## How to Upgrade

To begin the upgrade process:

1. Click the Upgrade to <7.n.n> link. The Data Fabric UI displays the **Upgrade fabric** form:



2. Fill in the information requested by the form. For more information, see [Upgrade Fabric Parameters](#) on page 84.
3. Click **Upgrade**. The **Upgrade fabric** status dialog box is displayed:

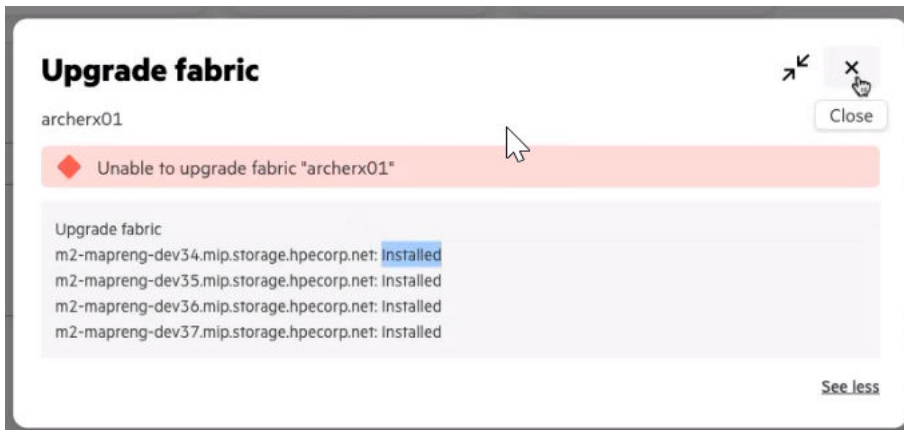


**IMPORTANT:** You can minimize the status dialog box, but do not close the dialog box until the upgrade is complete. If you close the dialog box, you will not be able to reopen it.

When the upgrade is successful, the value in the **Version** column shows the new version (for example, 7.5.0)


## Troubleshooting Upgrade Issues

If an upgrade fails, position your cursor over the failure message in the Data Fabric UI to obtain more information. For example:



Try to determine which node failed and the reason for the failure. In some cases, you can fix the issue manually. For example, if there is a repository or network issue, you might be able to resolve the issue on your own. You can then reinitiate the upgrade.

To reinitiate the upgrade:

1. In the global namespace list of resources, click the ellipsis (  ) for the fabric to be upgraded. The available commands are displayed.
2. Click **Reinitiate**. The **Upgrade fabric** form is displayed.
3. Fill in any empty values in the **Upgrade fabric** form.
4. Click **Upgrade**.

If the failure cannot be resolved manually, contact [HPE Support](#).

### Related concepts

[Viewing the Software Version](#) on page 115

Describes several ways to identify the core software version for a fabric.

## Upgrade Fabric Parameters

This page describes the configuration values that you need to specify to upgrade a data fabric.

Parameters with an asterisk (\*) are required. Before you can initiate the **Upgrade fabric** process, you must specify all required parameters.

By default, the upgrade software knows the fabric name and the provider (cloud or on-premises). For an on-premises upgrade, you must provide the SSH credentials used to create the fabric.

<b>Name*</b>	The name of the fabric. This field is typically grayed out because the Data Fabric UI prefills the <b>Name</b> information based on the fabric you selected for upgrading.
<b>Provider</b>	The fabric hosting information, which can be <b>On-premises</b> or one of several supported cloud services (AWS, Azure, or GCP). This field is typically grayed out because the Data Fabric UI prefills the <b>Provider</b> information based on the fabric you selected for upgrading.
<b>Username*</b>	The SSH user name, which is required for an on-premises upgrade.
<b>Password*</b>	The SSH password, which is required for an on-premises upgrade.

## User Assistance

Describes how to access different resources that can help you learn how to use the HPE Ezmeral Data Fabric.

To make the most of the HPE Ezmeral Data Fabric, be sure to review all of these user-assistance resources:

Resource	Description	To Access
HTML-Based Documentation	This guide, containing release notes, conceptual information, and step-by-step instructions.	Using any browser, navigate to <a href="https://docs.ezmeral.hpe.com/datafabric/home/index.html">https://docs.ezmeral.hpe.com/datafabric/home/index.html</a> .
Guided Tours	Two-minute interactive tours provided in the Data Fabric UI. The following tours are available: <ul style="list-style-type: none"> <li>"Get Started with HPE Ezmeral Data Fabric UI as a fabric manager"</li> <li>"Get Started with HPE Ezmeral Data Fabric UI as a fabric user"</li> </ul>	<ol style="list-style-type: none"> <li>Sign in to the Data Fabric UI.</li> <li>Either: <ul style="list-style-type: none"> <li>Click the <b>Guided Tour</b> button in the lower left corner of the screen.</li> <li>If the <b>Welcome</b> screen appears, click <b>Start Tour</b>.</li> </ul> </li> </ol>
In-Application Online Help	Tool tips and help buttons ( ? ) that describe fields and screens.	Hold your cursor over the button, or click the button.
Videos	Narrated product demonstrations.	The following product videos are available: <ul style="list-style-type: none"> <li><a href="#">Create Standard Volume</a></li> <li><a href="#">Create Mirror Volume</a></li> </ul>

### Related concepts


[Other Resources](#) on page 338

Provides links to additional resources such as on-demand training, videos, blogs, and the HPE Ezmeral Data Fabric community.

## Troubleshooting Online Help for the Data Fabric UI

Describes how to recover from the "Unable to fetch help topic" message.

For some deployments of the HPE Ezmeral Data Fabric, the following message can appear when you click a help button ( ? ) to display user assistance:



If your deployment has no internet access, this message is to be expected. Some online help information requires an internet connection for display purposes. As a workaround, you will need to consult this website for more information about specific features.

If your deployment has internet access, but a proxy is used for internet traffic with the Data Fabric, online help screens can time out or fail to fetch help content.

To resolve this issue, add the following proxy servers to the `/opt/mapr/apiserver/conf/properties.cfg` file:

- `http.proxy=<proxyServer>:<proxyPort>`
- `https.proxy=<proxyServer>:<proxyPort>`

#### Related concepts

[User Assistance](#) on page 85

Describes how to access different resources that can help you learn how to use the HPE Ezmeral Data Fabric.

#### Related tasks

[Configuring a Proxy Server for Data Fabric Access to the Internet](#) on page 108

Describes the procedure to configure an https or http proxy for scenarios where communication between the internet and Data Fabric must happen over a proxy server.

## Platform

---

This section contains conceptual information that can help you to understand and use the HPE Ezmeral Data Fabric.



**IMPORTANT:** To view platform information for the HPE Ezmeral Data Fabric – Customer Managed platform, see [this website](#).

## Data Fabric UI

---

Describes the graphical user interface for the HPE Ezmeral Data Fabric.

### About the Data Fabric UI

The Data Fabric UI is the browser-based, graphical user interface that you use to monitor and manage the HPE Ezmeral Data Fabric.

The Data Fabric UI can give you access to all the fabrics in the [global namespace](#). Depending on your user privileges, you can perform tasks such as the following (this is a partial list):

- Monitor system resources
- Monitor your billing and storage consumption
- Create or import fabrics
- Create volumes
- Create volume mirrors and snapshots
- Create buckets
- Create topics
- Manage users
- Control access to data

The **Home** page provides capacity and system resource information:

**My capacity** ⓘ All Fabrics

Total	Volumes	Buckets	Topics	Tables
<b>2.00 GiB</b>	<b>1.00 GiB</b>	<b>1.00 GiB</b>	<b>0.00 GiB</b>	<b>0.00 GiB</b>

**Resources** ⓘ

Search  Create Volume Bucket Topic Table 80%

24 entities

- GlobalNamespace
  - venkat-cluster214
    - 3 volumes
    - 3 tables
    - 2 topics
    - 3 buckets
  - venkat-cluster215
    - 2 volumes
    - 3 tables
    - 2 topics
    - 4 buckets

Switching to the **Fabric manager** view allows you to monitor and administer fabrics and resources:

**Welcome, admin** Fabric manager

Global namespace | Fabric metrics | **Fabric administration** | Security administration

**Fabric administration**  
 Select fabric: venkat-cluster214  
 Fabric ID: 8916186041906929698  
 Build version: 7.5.0.0.20231025110312.GA

Name	Type	Permissions
mapr	User	Login, Create volumes, Admin, Full control, Create security policy
root	User	Login, Create volumes, Admin, Full control, Create security policy

Quota by defaults:
 

- User quota: --
- Group quota: --
- User hard quota: --
- Group hard quota: --
- Fabric reserve limit: 90%

## Launching the Data Fabric UI

To launch the Data Fabric UI, navigate to the host that is running the WebServer in the fabric. Access to the fabric typically uses HTTPS on port 8443. For example:

```
https://<host-name>:8443/app/dfui
```

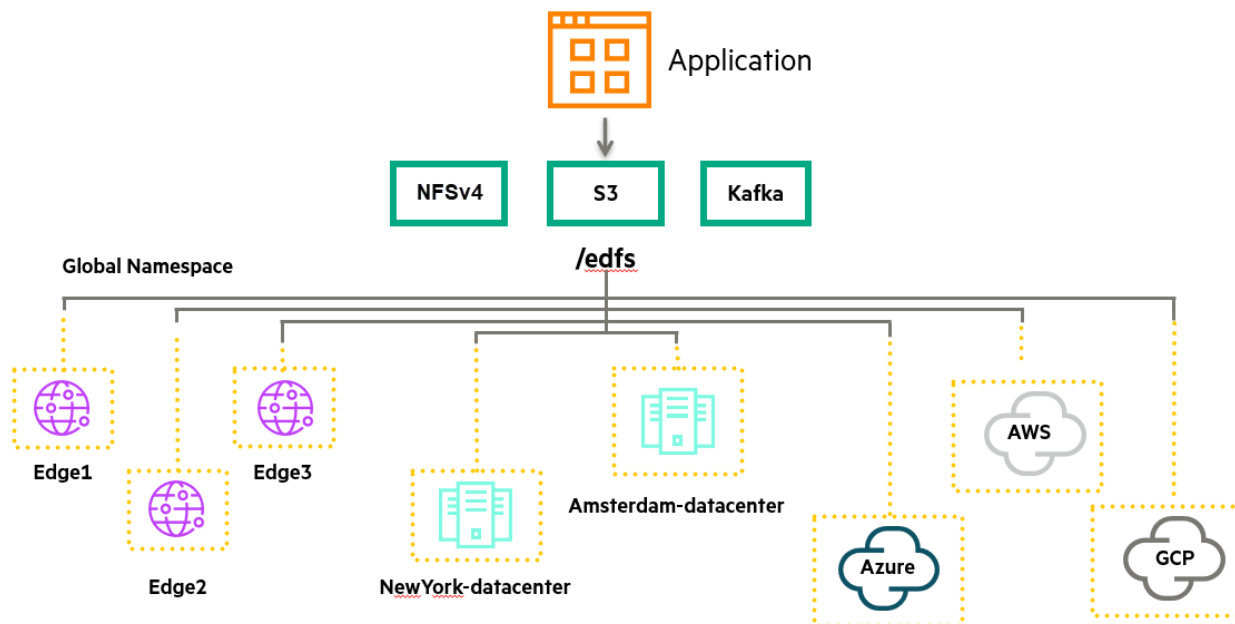
## Global Namespace (GNS)

Describes the data plane that connects all of your HPE Ezmeral Data Fabric deployments.

Because enterprise data is scattered across multiple sources between edge, core, and multi-cloud, a mechanism is needed to enable access to the data seamlessly, irrespective of the data location. The global namespace is a solution that aggregates disparate and remote data sources and provides a namespace that encompasses all of your infrastructure and deployments. The global namespace maintains the native security model of the HPE Ezmeral Data Fabric, so that location details are abstracted from the application.


Global namespace technology lets you manage globally deployed data as a single resource. Because of the global namespace, you can view and run multiple fabrics as a single, logical, and local fabric. The global namespace is designed to span multiple edge nodes, on-prem data centers, and clouds:





A global namespace stitches multiple data sources into a single accessible entity and provides multiple data access points, where each access point shows the same hierarchical namespace. The entire group of data sources could be accessed and operated by using any one of the access points, without any need to know the source's physical data location.

In the global namespace, all fabrics can view all other fabrics. The Data Fabric UI shows the global namespace on the **Home** page in the **Graph view**. For example:



The Data Fabric also supports adding external NFSv4 and external S3 object stores to a global namespace.

In support of S3 object stores, S3 clients can access multiple fabrics by using a single endpoint after specifying a single pair of access key and secret key.

### S3 Federation in Global Namespace

Provides an overview of S3 federation in the global namespace

Data Fabric provides a native implementation of the S3 object store for object management. You can access data from multiple fabrics with the native S3 object store implementation.

Additionally, Data Fabric provides a mechanism to access data that is stored on S3-compliant object stores from vendors other than HPE such as Amazon Web Services (AWS), Google Cloud Platform (GCP), , Scality, WEKA, VAST, to name a few . A third-party object store that is managed by vendors other than HPE is referred to as an external S3 object store (external S3 server).

S3 federation is a federation of one or more Data Fabric native S3 servers and/or one or more external S3 servers in the global namespace.

You can access remote S3 servers imported into your global namespace from any fabric in the same global namespace, by obtaining the access points to the object stores.

All buckets and objects from your native and external S3 servers can be accessed on the global namespace via the Data Fabric UI.

Data Fabric acts as the intermediary between a S3 client and an external S3 server that is imported into the global namespace. Data Fabric forwards requests coming from S3 clients to the external S3 server. The responses from the external S3 server are transported back to the requesting S3 clients via Data Fabric.

By default, Data Fabric works in forwarding mode with external S3 servers in their communication with S3 clients via Data Fabric. Data Fabric forwards the S3 client request to the external S3 server by using the secret key-access key pair that has been used to import the S3 server into the global namespace. The forwarding mechanism that happens during the S3 client-external S3 server communication consumes CPU and memory resources on Data Fabric. This can impact the Data Fabric performance, depending on the use case for frequency of communication between the S3 client and external S3 server.



**NOTE:** Currently, there is no performance benchmark available on Data Fabric performance for S3 client-external S3 server communication.

Data Fabric can work in redirection mode for S3 client-native S3 server communication with the `clustergroup s3gns` command. Data Fabric redirects an incoming S3 client request made to a native S3 server on another fabric in the same global namespace. Once the redirection is successful, the communication between the S3 client and the native S3 server on the other fabric takes place directly. In case of redirection, there is no impact on Data Fabric performance.

### Prerequisites for S3 Federation

Following are the prerequisites for S3 federation on a global namespace.

- External S3 servers and fabrics in the global namespace must be able to communicate with one another over the network.
- An external S3 server and an individual fabric in the global namespace must have a pair of secret key and access key that is authorized on the external S3 server, and can be used by the fabric to forward S3 client request to the external S3 server.

See [Working with an External S3 Object Store](#) on page 238 for details on managing external S3 servers via the Data Fabric UI.

## S3 Global Namespace

Describes what is meant by S3 global namespace.

### Introduction

S3 global namespace is a subset of the [global namespace](#) that comprises the native S3 storage (S3 storage on fabrics) and external S3 storage that are part of the global namespace.

A S3 client can authenticate all the S3 storage across fabrics in the S3 global namespace, by using a single set of access key/secret key.

A S3 client uses keys to connect to a S3 server. When the S3 GNS is explicitly enabled, the S3 client can choose any S3 server out of all the available native s3 servers to connect, and the same pair of keys should be valid on all of the native S3 servers. This is achieved by way of replicating the required tables across connected native fabrics. Enabling of s3 global namespace can be done from the command line.

The S3 global namespace also provides a unified control plane for all native S3 storage (S3 storage on fabrics) accessed from DFUI.

For the incoming s3 client requests intended for a remote/external s3 server, there are two ways the request can be handled:

- by redirecting the request to the correct s3 server
- by returning http forward error and letting the client to connect to the target s3 server directly.

Data Fabric works in redirection mode for S3 client - native/fabric S3 storage communication. Data Fabric redirects an incoming S3 client request made to a native S3 server on another fabric in the same global namespace. Once the redirection is successful, the communication between the S3 client and the native S3 server on the other fabric takes place directly.

Data Fabric works in forwarding mode for S3 client - external S3 storage communication.



**NOTE:** Data Fabric serves only as a unified data plane for external S3 servers that are imported into the global namespace. The control plane for the external S3 servers is accessible by using the respective interface or utility provided by the individual external S3 server. For example, bucket management can be performed on AWS S3 by using the AWS Management Console or AWS CLI.

### Enable S3 global namespace

By default, S3 global namespace is disabled on a global namespace.

When you enable S3 global namespace, the access keys, secret key and IAM policies from the primary fabric are replicated across all fabric storage or native S3 storage. The keys and IAM policies that exist on the non-primary fabrics in the S3 global namespace are overwritten on enabling S3 global namespace.



**CAUTION:** Enabling of S3 global namespace is a disruptive operation and must be performed only when it is feasible to have a single set of keys and policies across all the native S3 object stores/servers in your global namespace.

Use the `clustergroup s3gns enables3gns` command to enable S3 global namespace on the S3 storage across fabrics in your global namespace.

## Network File System in Global Namespace

Describes the federation of network file system in the global namespace.

A network file server(NFS) is hosted on a remote network, typically in a different physical location. The remote network file servers provide file-sharing services using NFS, allowing clients from different networks or locations to access and share files over the network.

HPE Ezmeral Data Fabric offers a native implementation of NFSv4. Data stored in the native NFSv4 system can be accessed from the global namespace.

The Data Fabric global namespace can be extended by importing one or more NFSv4 systems offered by vendors other than HPE. The NFSv4 system that is offered by vendors other than HPE is referred to as an external NFS with respect to Data Fabric.

An external NFS that is imported into Data Fabric appears as a part of the global namespace on the Data Fabric UI.

Data Fabric aggregates native NFSv4 systems and/or external NFSv4 systems in the global namespace, such that data from all the relevant NFSv4 systems in the global namespace is accessible via Data Fabric.

The addition of an external NFS to the Data Fabric global namespace facilitates the federation of all your NFS data, irrespective of the NFS vendor and the location of the data.

Data Fabric refers requests coming from an NFSv4 client to an external NFSv4 server that has been imported into Data Fabric. Once the connection between the NFSv4 client and the external NFSv4 server is established, the communication and data exchange between the NFSv4 client and external NFSv4 server takes place directly. Data Fabric is not a part of this communication.

See [Working with an External NFS Server](#) on page 236 for details on importing an external file system into the global namespace via the Data Fabric UI.

## Single Sign-On (SSO) Support

---

Describes how the HPE Ezmeral Data Fabric supports single sign-on (SSO).

### Keycloak IAM Support

The HPE Ezmeral Data Fabric supports SSO when configured with the Keycloak identity and access management (IAM) solution. Other IAM solutions are not currently supported.

### Keycloak Is Preinstalled and Preconfigured

Starting with release 7.5.0, Keycloak is preinstalled and preconfigured whenever you create a new fabric. You can create new users and roles easily and quickly by using the Keycloak administration console. For more information, see [SSO Using Keycloak](#) on page 50.

### Limitation for Non-SSO Users

SSO users with sufficient credentials can view and manage resources on all fabrics. Non-SSO users can view and manage resources only on the fabric to which they are signed in. Non-SSO users cannot view or manage resources on other fabrics. The Data Fabric UI does not display these resources to non-SSO users because the UI cannot connect to other fabrics without the same login information.

## Iceberg Support

---

Describes support for Iceberg in HPE Ezmeral Data Fabric 7.6.x.

### Apache Iceberg

Apache Iceberg is an open-source table format that helps to simplify the data processing of huge data sets on a file system or object store. Iceberg brings the simplicity of SQL tables to huge data sets.

Iceberg has the following capabilities:

- Iceberg tables are fast, safe, scalable, and can easily integrate with analytics engines like Spark, PrestoDB, Hive, and so on.
- Iceberg supports Atomicity, Consistency, Isolation, and Durability (ACID) transactions.
- You can use analytics engines like Spark, PrestoDB, Hive, and Impala to safely perform ACID transactions on the same table at the same time.
- Iceberg supports schema evolution, hidden partitioning, partition layout evolution, and time travel, which minimize unpleasant surprises.

See the [Apache Iceberg](#) documentation for details.

### Data Fabric and Iceberg

Starting from Data Fabric 7.6.x, you can perform the following operations in the HPE Ezmeral Data Fabric Object Store:

- Create a schema for Avro, ORC, or Parquet data types, and modify the schema if needed.
- Create Iceberg tables using a specific schema and perform ACID transactions.

- Create a snapshot of a table to check time travel.
- Grant access permissions for an Iceberg table to different users.
- Perform data migration of data files into an Iceberg table, as well as migrate the metadata.
- Query an Iceberg table through Apache Spark.
- Create an Iceberg table in an external S3 bucket and query it through the HPE Ezmeral Data Fabric Object Store.

With these features, you can build a reliable and scalable Data-Lakehouse architecture.

### Related concepts

[Getting Started with Iceberg](#) on page 252

Summarizes what you need to know to begin using Iceberg with HPE Ezmeral Data Fabric release 7.6.x.

## Fabric Resources

---

Describes fabric resources.

### What is a fabric resource?

Fabric resources are the entities or resources that are associated with a fabric.

Fabric resources are used to store organizational data. Organizational data is available in both structured and unstructured formats, in static and streaming format.

Different fabric resources can be used to store the organizational data, depending on the data format.

### Fabric Resource Types

Following is the list of fabric resources supported by HPE Ezmeral Data Fabric.

- **Volumes:** Volumes are used to store static data or structured data.
- **Buckets:** Buckets are used to store large objects or data that is disparate in nature such as audio files, video files, and images.
- **Topics:** Topics are used to store streaming or real-time data.

## Volumes

Brief conceptual information about volume.

A volume is a logical unit that allows you to apply policies to a set of files, directories, and sub-volumes. You can use volumes to enforce disk usage limits, set replication levels, establish ownership and accountability, and measure the cost generated by different projects or departments.

The volume structure defines how data is distributed across the nodes in your cluster.

You can create a volume for each user, department, or project.

You can mount volumes under other volumes to build a structure that reflects the needs of your organization.

Sub-volumes are created by mounting a volume in a sub-directory of an already mounted volume. This establishes a parent-child relationship between the volumes whereas the parent volume is mounted in top-level directory and the child volume is mounted in the sub-directory. Create multiple small volumes with shallow paths at the top of a fabric volume hierarchy to spread the load of access requests across the nodes.

A well-structured volume hierarchy is an essential aspect of fabric performance. As the data in a fabric grows, having an efficient volume hierarchy maximizes data availability as the data in a fabric grows. Fabric performance is negatively affected when a volume structure is not in place.

## Buckets

Describes buckets and the objects that they store.

Buckets are containers that store objects. Objects comprise disparate types of data, such as audio files, video files, and images.

Object-based storage is the preferred method of storing and efficiently managing gigantic amount of data.

Underlying each Data Fabric bucket is a volume. Every bucket created in a Data Fabric user account is automatically associated with a volume.

## Topics

Describes topics that are relevant to streaming data.

Topics are used to store streaming data.

A topic can be thought of as a persistent message queue. The lifetime of a topic or the time for which a topic must persist is configurable.

One or more client applications called producers publish streaming data in the form of messages to a topic. One or more client applications called consumers subscribe to the topics of their choice to consume the messages that are published to topics by subscribers.

Multiple producers can publish messages to the same topic and multiple subscribers can subscribe to and consume the message from such a topic.

Messages thus published to a topic are arranged or queued in the sequence of the publishing time.

## Data Storage Management

---

Summarizes options that the HPE Ezmeral Data Fabric provides to give you access to your data.

HPE Ezmeral Data Fabric facilitates efficient storage of data that is based on the frequency at which data is accessed.

Data Fabric provides rule-based automated tiering functionality that allows you to integrate seamlessly with:

- Low-cost storage as an additional storage tier in the fabric for storing file data that is less frequently accessed ("warm" data) in an erasure-coded volume.
- Third-party cloud object storage as an additional storage tier in the fabric to store file data that is rarely accessed or archived ("cold" data).

In this way, valuable on-premise storage resources can be used for more active or hot file data and applications, while warm and/or cold file data can be retained at minimum cost for compliance, historical, or other business reasons. Data Fabric provides consistent and simplified access to and management of the data.

## Data Tiering

Describes data tiering for efficient data access and data storage.

Data that is active and frequently accessed is considered as hot data. Data that is rarely accessed is considered cold data.

Hot data, and cold data is identified based on the rules and policies set by the administrator.

The mechanism used to store hot data is referred to as the hot-tier (or the data fabric cluster), and the mechanism to store cold data is referred to as the cold tier (or low-cost storage alternative on the cloud).

Data starts off as hot when it is first written to local storage (on the data fabric cluster). It becomes cold based on the rules and policies the administrator configures.

Data can be set up to be automatically offloaded using the data fabric automated storage tiering (MAST) Gateway service to the low-cost storage alternative on the third party cloud object store (cold tier) like S3.

The mechanism used to store hot data is referred to as the hot-tier that is nothing but the fabric storage or volumes. The mechanism to store cold data is referred to as the cold tier (or low-cost storage alternative on the cloud).

### **Cold Tier**

Describes a cold tier.

On the data fabric cluster, every cold tier (referred to as remote target on the Data Fabric UI) has a bucket on a third-party cloud store where volume data is offloaded based on the policy configured by the administrator.

Volume data in 64KB data chunks is packed into 8MB sized objects and offloaded to the bucket on the tier and the corresponding volume metadata is stored in a visible tier-volume as HPE Ezmeral Data Fabric Database tables on the data fabric cluster.

During writes and reads, volume data is recalled to the data fabric cluster if necessary. Data written to the volume is periodically moved to the remote target, releasing the disk space on the filesystem. See [Data Reads, Writes, and Recalls](#) for more information.

Data stored on the data fabric cluster requires thrice the amount of disk space of the regular volume on premium hardware due to replication (default being 3). After offloading to the cloud, the space used by data (including data in the namespace container) in the volume on the data fabric cluster is freed and only the metadata of the volume in the namespace container is 3-way replicated on the data fabric cluster.

There is also a visible tier-volume on the data fabric cluster for storing the metadata associated with the volume. When you create a cold tier, the tier volume named `mapr.internal.tier.<tierName>` is by default created in the `/var/mapr/tier` path. A directory/folder for the volumes associated with the tier, identifiable by `volumeid`, is created under the path after the first offload of data from the volume to the tier.

You can create one tier per volume or create and associate multiple volumes with the same tier using the Data Fabric UI.

### **Data Read, Write, and Recall**

Once offloaded to the storage tier, data is considered to be cold on the storage tier, but the data can still be accessed (read, written, and recalled).

### **Read of Tiered Data**

When the standard volume data is outside of the fabric storage, and in the cloud (cold tiering), Data Fabric processes the request to read standard volume data and mirror volume data.

### **Data Reads on Tiering-enabled Standard Volume**

When a client attempts to read, the read request is first sent to the volume on the fabric and if the data exists in the volume on the fabric, the data is returned from the volume. On the other hand, if the data was offloaded to a storage tier, Data Fabric recalls the data from the cold-tier to process the read request.

### **Data Reads on Tiering-enabled Mirror Volume**

When a client attempts to read, the read request is first sent to the volume on the Data Fabric cluster and if the data exists in the volume on the cluster, the data is returned from the volume. On the other hand, if

the data was offloaded, Data Fabric recalls or fetches a copy of the data (from the tier) into an associated cache-volume, from where data is returned to the client.

### Write on Tiered Data

When writes happen, if the write is:

- An append, new data is offloaded when the data meets the criteria in the rule (associated with the volume) for offload.
- A change to existing data (overwrite), the data is recalled to the Data Fabric file system to allow the write to succeed and then offloaded when the data meets the criteria in the rule (associated with the volume) for offload.



**NOTE:** If cold data is accessed (read/written) frequently, I/O to that file may suffer large latencies. In such scenarios, recall the whole volume or the corresponding files.

### Recall of Tiered Data

Offloaded data is automatically recalled when a client performs a read or overwrite on the data in the cold-tier, or when a client performs an overwrite on the data in the warm-tier. Data Fabric fetches a copy of the data to allow the operations to succeed.

## Data Storage Policy

Describes the use of data storage policy rules.

The storage policy simplifies the lifecycle management of data in the volume including automated migration of files to low-cost storage alternatives. The policy can contain rules for files that have a well-defined lifecycle or for files you want to switch to different storage tiers during their lifecycle.

You can specify the rules, at the volume level, to selectively identify files to offload (such as file size, file owner, and file modification time), the schedule for offloading the data (for example, 2 months after file modification), and the settings for storing (such as the location and credentials for the tier) and recalling the offloaded data. You can configure one rule per volume using the CLI or REST API. You can also associate a schedule to automatically offload data at scheduled intervals based on the associated rules.

Data offload is driven by rules, which are configured per volume. Data offload rule can be based on size of file, owner of the file, and/or file modification timestamp. You can apply one rule per volume.

When a rule is associated with a volume, the rule is first applied on the files in the tiering-enabled volume. When applied on the files in the tiering-enabled volume, the offload is triggered for all files in the snapshot chain as well when the criteria in the rule is met. If the file does not exist in the tiering-enabled volume, rule is applied on the latest state of the file in the snapshot chain. If the file exists in the tiering-enabled volume but has no latest state or if the file was deleted in the tiering-enabled volume, offload does not happen.

## AWS Architecture Notes

---

Describes architectural considerations for the HPE Ezmeral Data Fabric software-as-a-service (SaaS) platform when deployed on Amazon AWS.

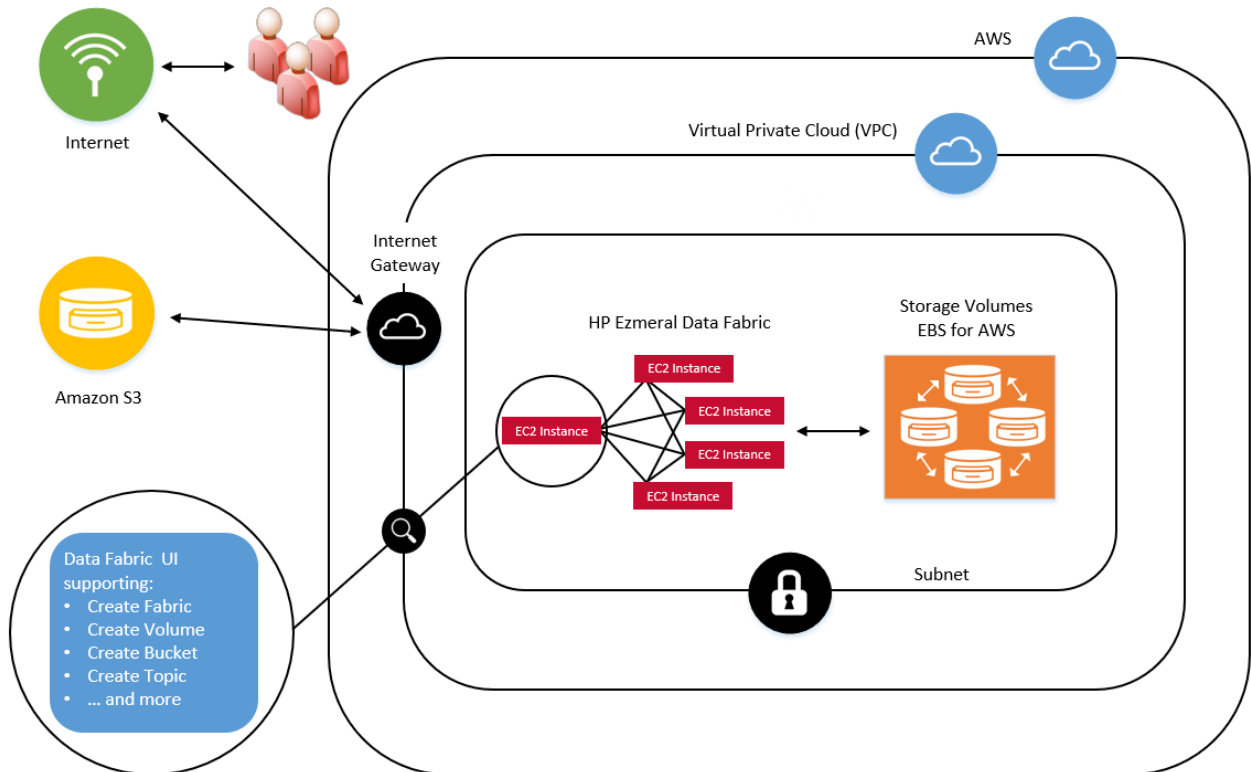
### Deployment Topology

To take advantage of the benefits of cloud computing, you can provision an HPE Ezmeral Data Fabric in Amazon AWS and in other public clouds. A single instance of the Data Fabric is referred to as a *fabric*. The fabric provides a high-performance file system for files, objects, tables, and streaming files and can be deployed quickly and easily. The HPE Ezmeral Data Fabric is designed so that many fabrics deployed in



different public clouds or on premises can communicate with each other seamlessly in a *global namespace (GNS)*.

The following diagram shows the high-level architecture for a single cloud-based fabric on AWS:



### Deployment Prerequisites

At a minimum, the user who deploys the Data Fabric on AWS must have [AmazonEBSCSIDriverPolicy](#) and [AmazonEC2FullAccess](#) permissions and must provide information such as the:

- Fabric name
- Access key
- Secret key
- Region
- Virtual private cloud (VPC) ID
- Public subnet ID

For more information, see [AWS Fabric Configuration Parameters](#) on page 29.

### Public and Private Subnets

To enable a global namespace consisting of many fabrics accessible over the internet, the user must provide a public subnet. The global namespace cannot currently be implemented with private subnets. The Data Fabric architecture does not prevent the use of private subnets, but some code changes are required before private subnets can be supported. Note that air-gapped, on-premises installations are fully supported.

## Regions and Availability Zones

The Data Fabric can be deployed into the following AWS regions:

- US East (Ohio)
- US East (N. Virginia)
- US West (N. California)
- US West (Oregon)
- Asia Pacific (Mumbai)
- Asia Pacific (Hyderabad)
- Asia Pacific (Singapore)
- Asia Pacific (Sydney)
- Asia Pacific (Melbourne)

In the current architecture, all fabric instances reside in a specific subnet, which is contained within a single availability zone (the default availability zone).

## Amazon Machine Images (AMIs)

Users of the HPE Ezmeral Data Fabric do not need to create or manage the AMIs needed to support the Data Fabric on AWS. HPE provides a set of publicly available AMIs that facilitate installation and upgrade of the fabric without the need for user interaction.

## Security Groups

During fabric creation, a security group is created for each fabric. The security group is configured with predefined in-bound and out-bound rules to support the list of [ports](#) required for fabric-to-fabric communication.

## STS Support for AWS S3 Object Stores

With release 7.7.0 and later, Data Fabric provides a new option for gaining access to AWS S3 object stores. You can import an external AWS S3 server by using the `maprcli clustergroup addexternal` command and specifying an Amazon Resource Name (ARN) to enable STS authentication. For more information, see [Integrating the AWS Security Token Service \(STS\) with Data Fabric](#) on page 243.

Using STS simplifies the process of accessing AWS services by using STS tokens for authentication. With STS tokens, the Data Fabric user can assume an AWS role and get temporary credentials to perform S3 actions. Once the external S3 object store is imported into the global namespace, **all S3 operations** automatically use STS.

## Instance, Disk, and Memory Information

See [AWS Cloud Instance Specifications](#) on page 256.

## Upgrades

When a new software version is available, the user is notified. At the user's discretion, the platform can perform a non-disruptive, rolling upgrade from one major software version to another. However, upgrade capability is currently limited to on-premises deployments. See [Upgrading a Data Fabric](#) on page 80.

## Scaling

Adding nodes to a fabric can be done using a rolling upgrade process. Note that adding nodes is currently supported only for on-premises deployments. See [Adding Nodes \(On-premises Deployment\)](#) on page 198.

## Administrative Interface

The [Data Fabric UI](#) on page 86 provides a browser-based graphical user interface for monitoring and managing all fabrics in a global namespace.

## SSO and Predefined Roles

The Data Fabric leverages the Keycloak identity and access management (IAM) solution to ensure that all the fabrics in a global namespace have access to the same user information. Keycloak can be used as a passthrough with other popular IAM solutions.

SSO-configured fabrics support the following predefined roles:

- Infrastructure Admin
- Fabric Manager
- Fabric User

For more information about the permissions granted to each role, see [Pre-defined Roles and Associated Permissions](#) on page 129.

## Azure Architecture Notes

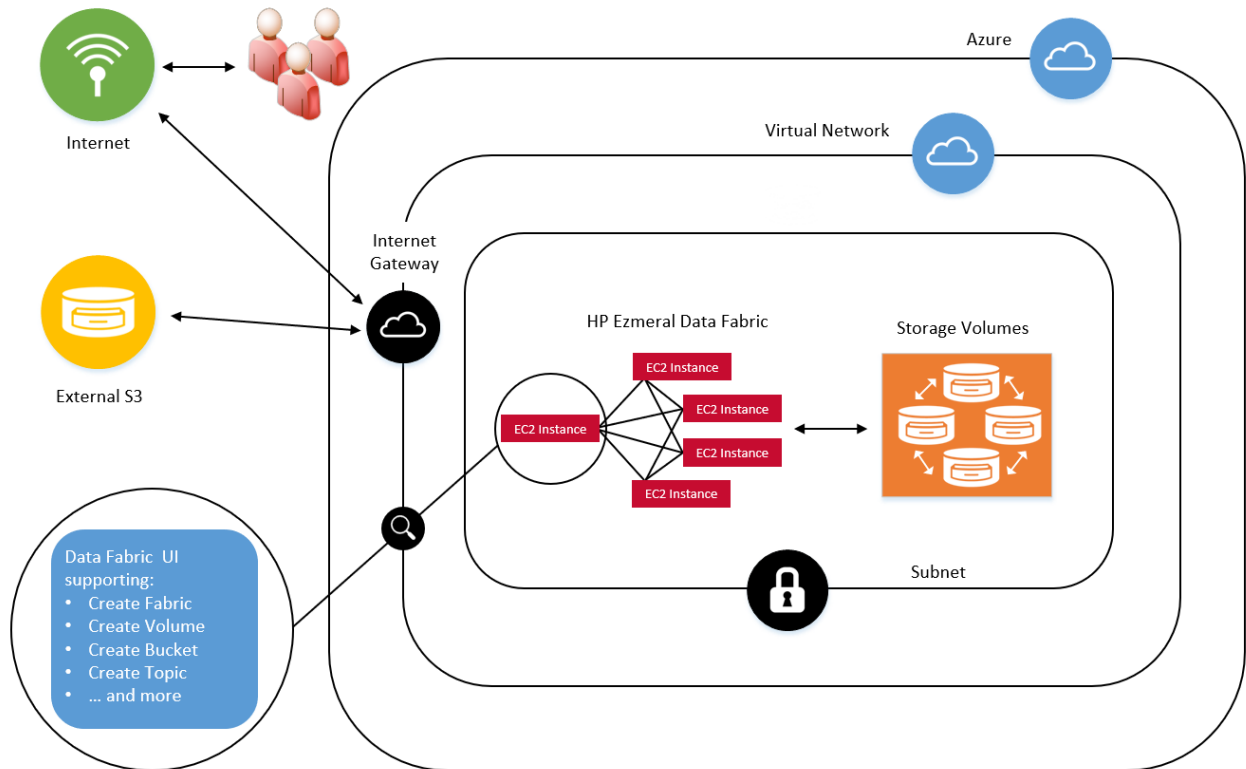
---

Describes architectural considerations for the HPE Ezmeral Data Fabric software-as-a-service (SaaS) platform when deployed on Microsoft Azure.

### Deployment Topology

To take advantage of the benefits of cloud computing, you can provision an HPE Ezmeral Data Fabric in Microsoft Azure and in other public clouds. A single instance of the Data Fabric is referred to as a *fabric*. The fabric provides a high-performance file system for files, objects, tables, and streaming files and can be deployed quickly and easily. The HPE Ezmeral Data Fabric is designed so that many fabrics deployed in different public clouds or on premises can communicate with each other seamlessly in a *global namespace (GNS)*.

The following diagram shows the high-level architecture for a single cloud-based fabric on Azure:



### Deployment Prerequisites

At a minimum, the user who deploys the Data Fabric on Azure must have the [Contributor](#) role and must provide information such as the:

- Azure tenant ID
- Subscription ID
- Client ID
- Client secret
- Region
- Resource group name
- Virtual network
- Subnetwork

For more information, see [Azure Fabric Configuration Parameters](#) on page 31.

### Public and Private Subnets

To enable a global namespace consisting of many fabrics accessible over the internet, the user must provide a public subnet. The global namespace cannot currently be implemented with private subnets. The Data Fabric architecture does not prevent the use of private subnets, but some code changes are required before private subnets can be supported. Note that air-gapped, on-premises installations are fully supported.

## Regions and Availability Zones

The Data Fabric can be deployed into the following Azure regions:

- East US
- East US 2
- West US
- West US 2

In the current architecture, all fabric instances reside in a specific subnet, which is contained within a single availability zone (the default availability zone).

## Network Security Groups

During fabric creation, a network security group is created for each fabric. The network security group is configured with predefined in-bound and out-bound rules to support the list of [ports](#) required for fabric-to-fabric communication.

## Instance, Disk, and Memory Information

See [Azure Cloud Instance Specifications](#) on page 256.

## Upgrades

When a new software version is available, the user is notified. At the user's discretion, the platform can perform a non-disruptive, rolling upgrade from one major software version to another. However, upgrade capability is currently limited to on-premises deployments. See [Upgrading a Data Fabric](#) on page 80.

## Scaling

Adding nodes to a fabric can be done using a rolling upgrade process. Note that adding nodes is currently supported only for on-premises deployments. See [Adding Nodes \(On-premises Deployment\)](#) on page 198.

## Administrative Interface

The [Data Fabric UI](#) on page 86 provides a browser-based graphical user interface for monitoring and managing all fabrics in a global namespace.

## SSO and Predefined Roles

The Data Fabric leverages the Keycloak identity and access management (IAM) solution to ensure that all the fabrics in a global namespace have access to the same user information. Keycloak can be used as a passthrough with other popular IAM solutions.

SSO-configured fabrics support the following predefined roles:

- Infrastructure Admin
- Fabric Manager
- Fabric User

For more information about the permissions granted to each role, see [Pre-defined Roles and Associated Permissions](#) on page 129.

## HPE Ezmeral Unified Analytics

Describes the HPE Unified Analytics Software and provides a link to more information.

Hewlett Packard Enterprise recommends the HPE Ezmeral Data Fabric as the [hybrid data lakehouse](#) for HPE Ezmeral Unified Analytics Software.

HPE Ezmeral Unified Analytics Software is a usage-based Software-as-a-Service (SaaS) model that operationalizes hybrid and multi-cloud analytical workloads through a simple user interface.

Available for use in connected or air-gapped environments, HPE Ezmeral Unified Analytics Software separates compute and storage for flexible, cost-efficient scalability. With HPE Ezmeral Unified Analytics, you can securely access data stored in multiple data platforms, and run traditional and advanced analytics workloads using open-source tools.

For more information, see the [Unified Analytics Software Documentation](#) home page.

## GCP Architecture Notes

Describes considerations related to Google Cloud Platform(GCP) for the HPE Ezmeral Data Fabric software platform for deployment to GCP.

### Deployment Overview

You can provision an HPE Ezmeral Data Fabric in Google Cloud Platform (GCP) and in other public clouds to take advantage of the benefits of cloud computing. A single instance of the Data Fabric is referred to as a *fabric*. The fabric provides a high-performance file system for files, objects, tables, and streaming files and can be deployed quickly and easily. The fabric makes use of the storage provided by the Google Cloud Platform to store the fabric data.

### GCP Credentials

You would require the following GCP cloud credentials to access GCP for fabric creation on GCP.

- service account key file

### GCP Services

The following table details the GCP services are that relevant to provisioning Data Fabric on GCP.

GCP Service	Description/Purpose
VPC	High-level GCP service for cloud-based fabric deployment
Google cloud image (AMI equivalent)- OS image	Required to create an operating system image on GCP
Google compute instance/ VM instance- custom image (base OS images+DF packages)	Required to host base operating system image instances plus Data Fabric packages. Can be used to generate a VM instance from the custom image on which Data Fabric would be installed.
Google compute disk	Required for storage of fabric data.
Subnet	Required for network communication between GCP and Data Fabric. You must create the required number of public subnets. Private subnet is not supported for Data Fabric. You must configure firewall rules to manage subnet access for Data Fabric. For information on required ports on which various Data Fabric components communicate, see <a href="#">port information</a>

### Permissions required to access GCP resources

GCP implements IAM roles for the purpose of managing access to GCP cloud resources.

In general, you must be assigned one of the following roles to be able to provision and manage fabrics on GCP.

- Compute Instance Admin(v1)
- Compute Security Admin

### Regions and Zones

A fabric can be provisioned on any the GCP zones that are available to you while creating the fabric via the Data Fabric UI.

Following are the GCP zones on which you should be able to create a fabric.

- Council Bluffs, Iowa, North America -us-central1-a
- Council Bluffs, Iowa, North America -us-central1-b
- Council Bluffs, Iowa, North America -us-central1-c
- Council Bluffs, Iowa, North America -us-central1-f
- Moncks Corner, South Carolina, North America -us-east1-b
- Moncks Corner, South Carolina, North America -us-east1-c
- Moncks Corner, South Carolina, North America -us-east1-d
- Ashburn, Virginia, North America -us-east4-a
- Ashburn, Virginia, North America -us-east4-b
- Ashburn, Virginia, North America -us-east4-c
- Columbus, Ohio, North America -us-east5-a
- Columbus, Ohio, North America -us-east5-b
- Columbus, Ohio, North America -us-east5-c
- Dallas, Texas, North America -us-south1-a
- Dallas, Texas, North America -us-south1-b
- Dallas, Texas, North America -us-south1-c
- The Dalles, Oregon, North America -us-west1-a
- The Dalles, Oregon, North America -us-west1-b
- The Dalles, Oregon, North America -us-west1-c
- Los Angeles, California, North America -us-west2-a
- Los Angeles, California, North America -us-west2-b
- Los Angeles, California, North America -us-west2-c
- Salt Lake City, Utah, North America -us-west3-a

- Salt Lake City, Utah, North America -us-west3-b
- Salt Lake City, Utah, North America -us-west3-c
- Las Vegas, Nevada, North America -us-west4-a
- Las Vegas, Nevada, North America -us-west4-b
- Las Vegas, Nevada, North America -us-west4-c

### Instance, Disk, and Memory Information

See [GCP Cloud Instance Specifications](#) on page 257.


### Upgrades

When a new software version is available, the user is notified. At the user's discretion, the platform can perform a non-disruptive, rolling upgrade from one major software version to another. See [Upgrading a Data Fabric](#) on page 80.

## Administration

---

This section describes how to administer fabric resources in the global namespace of your HPE Ezmeral Data Fabric as-a-service platform.

 **IMPORTANT:** To administer the HPE Ezmeral Data Fabric – Customer Managed platform, see [this website](#).

## IPv6 Support in Data Fabric

---


Describes the IPv6 support feature for Data Fabric.

Data Fabric can be installed on hosts with IPv6 addresses. In other words, external endpoints for Data Fabric can have IPv6 addresses. Data Fabric can communicate with clients over IPv6 addressing. Inter-fabric traffic and intra-fabric traffic over IPv6 connections is supported with IPv4 compatibility.

Data Fabric deployment over IPv6 addresses is possible when both the hardware hosting Data Fabric and the Data Fabric software are able to detect and support IPv6 addresses.

The underlying hardware that hosts Data Fabric must have a network interface card (NIC) that supports IPv6 addressing.

An application that wishes to communicate with Data Fabric over IPv6 can do so, when Data Fabric is installed on IPv6-compatible hardware and IPv6 support is enabled on Data Fabric.

 **NOTE:** When you have a network interface card (NIC) that supports both IPv4 and IPv6 addressing, each IP address must be identifiable by a distinct hostname. In other words, a single hostname must NOT map to both the IPv4 and IPv6 addresses.

The following table describes the terminology related to IPv6 client/server nodes.

Term	Description
IPv6-aware	The term denotes readiness of the underlying hardware. It indicates that the NIC associated with a node that hosts Data Fabric is IPv6 compatible, and can communicate with other nodes with IPv6 and IPv4 addresses.



Term	Description
IPv6-unaware	The term denotes readiness of the underlying hardware. It indicates that the NIC associated with a node that hosts Data Fabric is incompatible to handle IPv6 traffic, and can handle IPv4 traffic only.
IPv6-enabled	The term denotes that IPv6 is enabled on Data Fabric software. The Data Fabric node on which IPv6 is enabled is able to communicate with IPv6 addresses. The node is able to communicate with IPv4 addresses.
IPv6-only	The term denotes that IPv6 is enabled on Data Fabric software. The Data Fabric node on which IPv6 is enabled is able to communicate exclusively with IPv6 addresses only. Communication with IPv4 addresses is not supported on this node.

The following matrix explains in detail the communication between a client node and a Data Fabric node for various IP address type combinations.

Type of application (Type of client node)	IPv6-unaware server (IPv4-only server node)	IPv6-unaware server (IPv6-enabled server node)	IPv6-aware server (IPv6-only server node)	IPv6-aware server (IPv6-enabled server node)
IPv6-unaware client (IPv4-only node)	client-server communication takes place over IPv4	client-server communication takes place over IPv4	no communication	client-server communication takes place over IPv4
IPv6-unaware client (IPv6-enabled node)	client-server communication takes place over IPv4	client-server communication takes place over IPv4	no communication	client-server communication takes place over IPv4
IPv6-aware client (IPv6-only node)	no communication	no communication	client-server communication takes place over IPv6	client-server communication takes place over IPv6
IPv6-aware client (IPv6-enabled node)	client-server communication takes place over IPv4	client-server communication takes place over IPv4	client-server communication takes place over IPv6	client-server communication takes place over IPv6

## Enabling IPv6

Describes the procedure to enable IPv6 communication on an on-premises fabric.

Using the Data Fabric UI to enable IPv6 on a cloud fabric or a customer-managed installation is not currently supported.

### Prerequisites

Note the following prerequisites:

- The hardware on which Data Fabric is installed must be IPv6 compatible; that is, the hardware must have an IPv6-compatible NIC.
- To enable intra-fabric communication over IPv6, IPv6 must be enabled over all communicating nodes of the fabric.
- IPv6 must be enabled for the operating system (OS) used by your installation. To ensure that IPv6 is enabled for your OS, check the vendor documentation.

### About this task

Enabling IPv6 is a one-time operation and the operation is not reversible.

You can enable IPv6 for an on-premises fabric:

- When you initiate a new installation.
- After installation by using the Data Fabric UI.

- After installation by using `maprcli` commands.

Follow the steps given below to enable IPv6 on an on-premises fabric.

### Enabling IPv6 for a New Installation

For on-premises deployments, you can enable IPv6 as part of the seed-node installation procedure. See [On-Premises Fabric Configuration Parameters](#) on page 33 and [Fabric Deployment Using a Seed Node](#) on page 18.

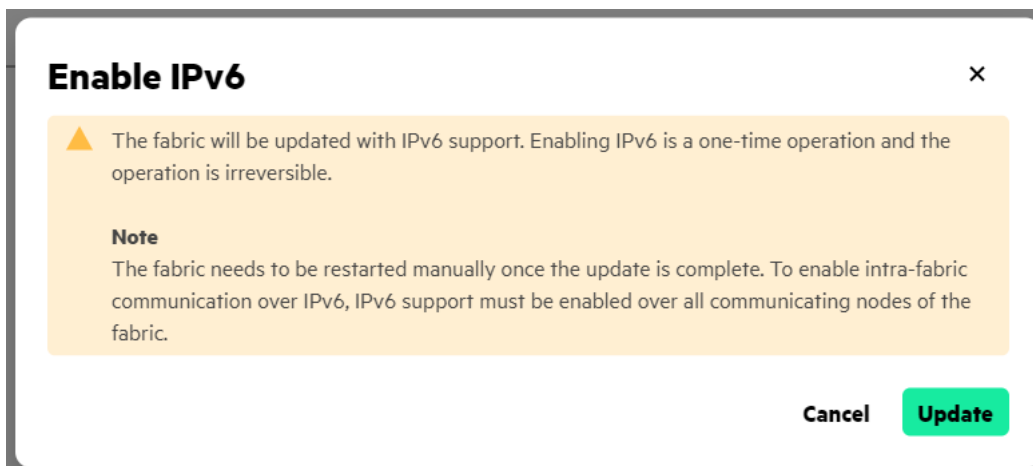
### Enabling IPv6 by Using the Data Fabric UI

Use the following steps:

1. Log on to the Data Fabric UI as a fabric manager.
2. If you are a fabric user, click the Table View icon on the **Resources** card. If you are a fabric manager, select the **Fabric manager** option, and click **Global namespace** to display the table view.
3. Under the **Resource Name** column, click the link for the fabric.
4. Click the **Overview** tab if it is not already selected.
5. Next to the **Internet protocol version** entry, click **Enable IPv6**:

The screenshot shows the HPE Ezmeral Data Fabric UI interface. The top navigation bar includes a home icon, the text 'HPE Ezmeral Data Fabric', a notification bell with '5+', a help icon, and the user 'admin'. The breadcrumb trail is 'Dashboard > cluster166'. The main content area is titled 'cluster166 | Fabric' with an 'Actions' dropdown. Below this is a tabbed interface with 'Overview' selected. The 'Fabric details' section lists various attributes: Fabric ID, Global namespace (GlobalNameSpace), Provider (OnPrem), Internet protocol version (IPv4), Nodes (1 node), Build version (7.8.0.0.20240709133642.GA), NFS server root, Object server storage endpoint, and Data at rest encryption (Enabled). The 'Alarms' section shows 0 Errors, 2 Warnings, and 1 Info. The 'Internet protocol version' entry is highlighted with a red box, and the 'Enable IPv6' button is visible next to it.

A confirmation dialog box appears:



6. Click **Update**.
7. Restart the fabric manually by using the steps in [Enabling IPv6 by Using maprccli Commands](#) on page 107.

### Enabling IPv6 by Using maprccli Commands

Use the following steps:

1. Log on to the command line of the primary node of the fabric on which you wish to enable IPv6.
2. Run the following command to enable the IPv6 feature:

```
maprccli cluster feature enable -name cldb.ipv6.support
```

3. To confirm that the feature has been enabled, run the following command:

```
maprccli config load -json |grep cldb.ipv6.support
```

4. Restart Warden on the current primary CLDB to initiate a CLDB node failover:

```
systemctl restart mapr-warden
```

This command causes all of the cluster nodes (MFS, NFS, FUSE) to re-register with the new primary CLDB and enables the primary CLDB to use the IPv6 addresses. For more information about the CLDB, see [Identifying All CLDB Nodes](#) on page 69.

## Administering Fabrics

This section describes fabric operations that you can perform using the Data Fabric UI.

A data fabric is a collection of nodes that work together under a unified architecture, along with the services or technologies running on that architecture. Fabrics help you manage your data, making it possible to access, integrate, model, analyze, and provision your data seamlessly.

Using the [Data Fabric UI](#) on page 86, you can create fabrics hosted by the following providers:

- AWS
- Azure

- GCP
- On-premises

## Configuring a Proxy Server for Data Fabric Access to the Internet

Describes the procedure to configure an https or http proxy for scenarios where communication between the internet and Data Fabric must happen over a proxy server.

### Prerequisites

You must be a fabric manager to be able to perform this task.

### About this task

Depending on the network configuration, Data Fabric might require to connect to the internet via a proxy server for download of Data Fabric packages, or to communicate with the SSO server that might be outside of the network in which Data Fabric is installed. A https or http proxy server can serve as a common routing point for internet access from Data Fabric.

It is recommended to use an https proxy for secure communication between Data Fabric and any server to be accessed over the internet.

You may, optionally, specify a http proxy server for any communication over http.

The format to specify the proxy server is <proxy-server-url>:<portnumber>.

Some Data Fabric components are auto-restarted when you save the proxy settings. You are required to refresh the Data Fabric UI after about 15 to 20 minutes after saving the proxy settings.



**NOTE:** If you have signed in to the Data Fabric UI by using your single sign-on (SSO) credentials, your session is active, and it should suffice to refresh the page in the web browser after the restart of Data Fabric components. If you are a non-SSO user, you must log off and log back in to the Data Fabric UI.

Follow the steps given below to configure a proxy server.

### Procedure

1. Log on to the Data Fabric UI and switch to the **Fabric manager experience**.
2. Click **Fabric Administration**.
3. On the **Fabric Settings** card, click the Edit icon adjacent to **Proxy configuration**.
4. Enter the https proxy URL and port number in **HTTPS proxy** in the <proxy-server-url>:<portnumber> format.
5. Optionally, enter the http proxy URL and port number in **HTTP proxy** in the <proxy-server-url>:<portnumber> format.
6. Click **Save**.

### Results

The proxy settings are saved. Some Data Fabric components restart and this might result in an error after about 15 minutes of saving the proxy settings. Refresh the Data Fabric UI page that you are working on at such time.

## Creating a Fabric

Fabrics make it possible for you to create volumes, buckets, and topics in a cloud or on-premises deployment. If your organization has multiple departments or multiple use cases to support, you can create multiple fabrics.

This page describes the basic steps to create a new *fabric* for any of the supported fabric providers (AWS, Azure, GCP, and on-premises).

### Before Creating a Fabric

Note these considerations:

- If you are creating your first fabric, see [Fabric Deployment Using a Seed Node](#) on page 18. You must use the seed node steps to create your first fabric. For all subsequent fabrics you can use the steps on this page.
- To create a fabric, you must have fabric manager [credentials](#). The **Create Fabric** button is not displayed for developer and infrastructure admin credentials.
- Currently, only SSO users can create fabrics.
- You must obtain a license for and register each new fabric that you create.
- Always create fabrics one at a time. You cannot create multiple fabrics at the same time.
- Creating an on-premises fabric requires that you provide host nodes *before* starting fabric creation. These nodes must meet certain requirements. Before creating an on-premises fabric, review [Prerequisites for On-Premises Installation](#) on page 27.

### Steps for Creating a Fabric


Use the following steps to create a new fabric.

1. Log on to the Data Fabric UI with Fabric Manager [credentials](#).
2. Click **Create fabric**. The **Create fabric** side drawer appears.
3. Fill in the configuration parameters for the type of fabric you want to create:
  - [AWS Fabric Configuration Parameters](#) on page 29
  - [Azure Fabric Configuration Parameters](#) on page 31
  - [GCP Fabric Configuration Parameters](#) on page 32
  - [On-Premises Fabric Configuration Parameters](#) on page 33
4. Click **Create**.
5. To monitor the progress of fabric creation, check the status bar in the **Fabric details** dialog box, or click **See details**. Fabric creation can take 20 minutes or more.

If fabric creation fails, you can retry the operation. Click the ellipsis in the **Action** column, and select **Reinitiate**.

It is possible for fabric creation to fail because of a host-name resolution issue. See [Troubleshooting Seed Node Installation](#) on page 34.

If fabric creation continues to fail, and the failure cannot be resolved manually, contact [HPE Support](#).

6. When the installation status shows **Installed**, click the ellipsis (  ) in the **Action** column, and select **View endpoints**. The URL for the new fabric is displayed, and you can copy the URL to the clipboard.
7. Add your fabric activation key. See [Adding an Activation Key](#) on page 42.
8. Register the fabric. See [Registering a Fabric](#) on page 43.
9. Set the billing model. See [Setting the Billing Model](#) on page 43.

## Importing a Fabric

This section provides the steps to import an as-a-service fabric into the global namespace.

### Importing an as-a-Service Fabric

Describes the steps to import an as-a-service fabric into the global namespace.

### Considerations for Importing an as-a-Service Fabric

An as-a-service fabric is a fabric that exists as part of a global namespace and was created using the **Create fabric** functionality of the Data Fabric UI. From any as-a-service fabric, you can import another as-a-service fabric into the global namespace by using the **Import fabric** command.

A fabric can belong to only one global namespace at a time. Thus, the act of importing an as-a-service fabric necessarily removes the fabric from the global namespace to which it currently belongs.

To view the current list of fabrics in your global namespace, display the **Table view** or **Graph view** on the **Resources** card of the **Home** page.

Note these considerations:

- Only fabrics configured for SSO can be imported.
- To import a fabric, you must be an SSO user and have Fabric Manager or Fabric User [credentials](#).
- You can only import one fabric at a time.
- You must have a consumption license for each new fabric that you import.

### Preparing to Import an as-a-Service Fabric

1. On the fabric that you plan to import, stop Keycloak:
  - a. Use the following command to identify the host running the `mapr-keycloak` service:

```
maprcli node list -columns svc
```

In the command output, look for a host that shows `keycloak` in the `service` column. If no host shows the `mapr-keycloak` service in the `service` column, go to step 2.

- b. Stop the `mapr-keycloak` service:

```
maprcli node services -name keycloak -action stop -nodes -json
```

2. On the fabric that you plan to import, reset the SSO information:
  - a. Reset the SSO configuration:

```
maprcli cluster resetssoconf -json
```

- b. Restart the `mapr-apiserver` services on the fabric hosts:

```
maprcli node services -name apiserver -action restart -nodes
host1,host2 -json
```

3. Use the following command to disable the `pbs.master` role for the fabric to be imported:

```
maprcli config save -values {cldb.pbs.global.master:0} -json
```

If any security policies have been created on the fabric to be imported, they must be manually re-created on the importing fabric after the import operation is completed. To re-create the policies, refer to [Administering Security Policies](#) on page 206.

4. On the cluster to be imported, create a tar ball of the fabric directory:

```
/opt/mapr/installer/ezdfaas/deployments/<cluster-name>
```

5. Copy the contents of the tar ball to the importing cluster's `/opt/mapr/installer/ezdfaas/deployments` directory. Extract the contents, and be sure to delete the `.tar` file:

6. Obtain the SSO configuration from the importing fabric, and configure it on the fabric to be imported:

- a. Use the following command to fetch the SSO parameters from the importing fabric:

```
maprcli cluster getssoconf -json
```

For example:

```
maprcli cluster getssoconf -json
{
  "timestamp":1699432649586,
  "timeofday":"2023-11-08 12:37:29.586 GMT-0800 AM",
  "status":"OK",
  "total":1,
  "data":[
    {
      "issuerendpoint":"https://<hostname>:443/realms/master",
      "providername":"keycloak",
      "clientid":"edf-client",
      "clientsecret":"<secret>"
    }
  ]
}
```

- b. Obtain the SSO certificate from the importing fabric's `/opt/mapr/keycloak/conf/.cert`, and use it to set the SSO configuration information for the fabric to be imported. Use the following command:

```
maprcli cluster setssoconf -issuerendpoint "https://:8443/realms/
TestReallm" -providername keycloak -clientid edf-client -clientsecret
<secret> -certfile -json
```

- c. Restart the `mapr-apiserver` services.

```
maprcli node services -name apiserver -action restart -nodes
host1,host2 -json
```

- d. Wait for a minute to ensure that the SSO configuration is active, then try signing in to the UI:

```
https://<apiserver>:8443/app/dfui
```

You should be redirected to the Keycloak sign-in screen.

7. Use the Data Fabric UI to complete the **Import** operation as described in the next section.

### Completing the Import Operation by Using the Data Fabric UI

Use the following steps to complete the import operation:

1. Log on to the Data Fabric UI as a Fabric Manager.
2. Click **Import fabric**. The **Import fabric** menu appears.
3. Specify the current **Name** of the fabric to be imported. Do not change the name of the fabric to be imported.
4. Specify the **Public IP address** of the APIserver of the fabric to be imported.
5. Specify the port of the APIserver for the fabric to be imported.
6. Click **Import**.
7. Once the **Import** operation is finished, ensure that the imported fabric is part of a cluster group and the fabric is listed as part of the global namespace.
8. Complete the following tasks for the new fabric:
  - [Registering a Fabric](#) on page 43
  - [Adding an Activation Key](#) on page 42
  - [Setting the Billing Model](#) on page 43

## Viewing the Fabric Status

Describes how to use the **Global namespace** card.

### About the Global namespace Card

The **Global namespace** card shows the current status and software version for each fabric and allows you to perform certain fabric-level actions (subject to your role). These actions can include:

- Viewing fabric status and error information
- Viewing fabric access points
- Importing a fabric
- Creating a fabric
- Registering a fabric



- Adding an activation key
- Setting the billing model
- Importing an external S3 server
- Importing an external NFS server
- Upgrading fabric software (if a new software version is available)
- Deleting a fabric
- Reinitiating (retrying) an upgrade operation

### Viewing the Fabric Status

To view the fabric status:

1. Sign in to the Data Fabric UI, and switch to the **Fabric manager** view.
2. Click **Global namespace**.
3. Click the **Table view** icon to display the resource table with status values.

### Fabric Status Information

Fabric status values include:

Status	Description
Active	An activation key has been added as described in <a href="#">Adding an Activation Key</a> on page 42.
Deleting	Fabric removal, as described in <a href="#">Deleting a Fabric</a> on page 127, is currently in progress.
Expired	The activation key for the fabric is no longer valid.
Inactive	An activation key has not been added for the fabric. See <a href="#">Adding an Activation Key</a> on page 42.
Install Failed	There was a problem during installation, and the fabric was not successfully installed.
Installed	Fabric installation completed successfully.
Installing	Fabric installation is currently in progress.
Upgrade Failed	There was a problem during the software upgrade, and the fabric was not successfully upgraded.
Upgrading	A software upgrade is currently in progress.

### Viewing Fabric Settings

Describes how to view and change the fabric settings, which include fabric auditing, data auditing, and gateway information.

#### About the Fabric Settings

Fabric settings currently include:

Setting	Default Value	Description
Fabric auditing	Off	Auditing of fabric-management operations and fabric administration.
Data auditing	Off	Auditing of data-access operations.

Setting	Default Value	Description
Gateway	N/A	The Gateway parameter is not currently supported. For table replication, generate a DNS record that specifies the location of the gateways in the fabric, which can copy and paste into the zone file for your domain. Source fabrics can look up gateways during table replication using this record. Before generating the record, ensure that you have configured gateways in your cluster.
Proxy configuration	No proxy configured	Enables the configuration of an https or http proxy for environments that have a proxy server.

### Viewing the Fabric Settings

To view the **Fabric Settings** card:

1. Sign in to the Data Fabric UI, and switch to the **Fabric manager**.
2. Click **Fabric administration**. The **Fabric Settings** card appears under the **Quota by** card.

### Changing the Fabric Settings

To change fabric settings:

To	Do this
Turn on fabric auditing	Click and drag the slider to the right.
Turn on data auditing	Click the <b>Edit</b> icon (✎), select <b>On</b> , and click Update.
Copy the gateway information to the clipboard	Click the the <b>Copy</b> icon (📄).
Configure the proxy	See <a href="#">Configuring a Proxy Server for Data Fabric Access to the Internet</a> on page 108.

#### Related reference

[Auditing Fabric and Fabric Data](#) on page 205  
Auditing in Data Fabric

### Viewing the Fabric Endpoint

Describes how to view the endpoint for a fabric on the Data Fabric UI.

#### Prerequisites

You must be a fabric manager to perform this operation.

#### About this task

You can view the endpoint for a fabric from the Data Fabric UI.

#### Procedure

1. Log on to the Data Fabric UI.
2. Select the **Fabric manager** from the dropdown next to the welcome message on the Home page.
3. Click **Global namespace**.
4. On the table view, click the ellipsis under **Action** for the fabric whose endpoint you wish to view.
5. Click the **View endpoint** option.

## Results

The fabric endpoint is displayed. The endpoint can be used to access the fabric and fabric resources. You can download the endpoint.

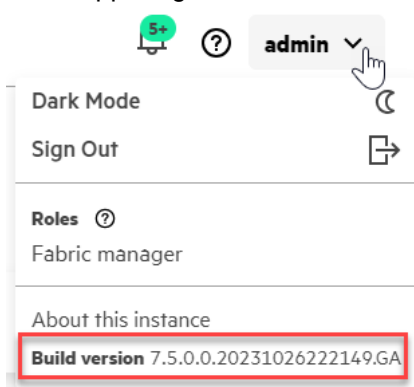
## Viewing the Software Version

Describes several ways to identify the core software version for a fabric.

### View the Software Version with User Information

To view the core software version:

1. Sign in to the Data Fabric UI.
2. In the upper right corner of the home screen, click the down arrow next to the user name. For example:



### Viewing the Software Version with Fabric Details

To view the core software version:

1. Sign in to the Data Fabric UI, and switch to the **Fabric manager** view.
2. Click **Global namespace**.
3. Click the **Table view** icon to display the resource table with status values.
4. Click the fabric name to display the fabric details page. The core software version is displayed as the **Build Version**:

The screenshot shows the HPE Ezmeral Data Fabric administration interface. The top navigation bar includes the HPE logo, the title 'HPE Ezmeral Data Fabric', and a user profile 'admin'. The breadcrumb trail is 'Dashboard > tejatest'. The main content area is titled 'tejatest | Fabric' and has an 'Actions' dropdown. Below this is a navigation menu with tabs: Overview (selected), Services, Security policies, Storage policies, Remote targets, and Settings. The 'Fabric details' section lists various attributes: Fabric ID, Global namespace (GlobalNameSpace), Provider (OnPrem), Node IP addresses (with a 'View all' link), Build version (7.5.0.0.20231026222149.GA, highlighted with a red box), NFS server root (with a 'View all' link), Object server storage endpoint (with a 'View all' link), and Data at rest encryption (Enabled). The 'Alarms' section shows a summary: Error (0), Warning (2), and Info (0).

### Related concepts

[Release History](#) on page 255

Describes the currently released versions of the HPE Ezmeral Data Fabric as-a-service platform.

## Generating S3 Access Keys for the Global Namespace


Describes how to obtain an S3 user access key and secret key that can be used to perform operations on S3 resources anywhere in the global namespace.

For the fabric manager you can generate and download a user access key and secret key. These keys facilitate command-line and programmatic (API) access to S3 resources for all fabrics in the namespace. For example, once you have generated the keys for the fabric manager, you can use MinIO client (`mc`) commands to set an alias for the fabric manager. Then you can use the same alias to perform operations on all fabrics in the namespace. For information about the supported `mc` commands, see [MinIO Client \(mc\) Commands](#).

You can generate the keys only twice for the same global namespace. More attempts to generate keys result in an error message.

Use these steps:

1. Sign in to your local fabric as a fabric manager. The **Global namespace** screen appears.
2. In the Global namespace card, click the **Graph view**.
3. Click the icon for the global namespace.
4. Click **View access points**:



The **Access points** screen is displayed.

- Click the **S3 servers** tab. The screen displays the Access keys and S3 server details:

 A screenshot of the 'Access points' screen in the GlobalNameSpace. The title 'Access points' is at the top right with a close button (X). Below the title, 'GlobalNameSpace' is displayed. There are two tabs: 'NFS servers' and 'S3 servers', with 'S3 servers' being the active tab. Below the tabs, the section 'Access keys' is shown. It includes an information icon and the text 'You can generate up to 2 keys.' Below this, it states 'No S3 access keys exist. Once a key is generated, it will be displayed here.' A green 'Generate key' button is present. The 'S3 server details' section follows, showing 'Default:' and 'External:' URLs. A 'Download' button is at the bottom.
 

**Access points** ×

GlobalNameSpace

**NFS servers** **S3 servers**

---

**Access keys**

ⓘ You can generate up to 2 keys.

No S3 access keys exist.  
Once a key is generated, it will be displayed here.

**Generate key**

**S3 server details**

**Default:** <https://10.0.1.204:9000> ⓘ

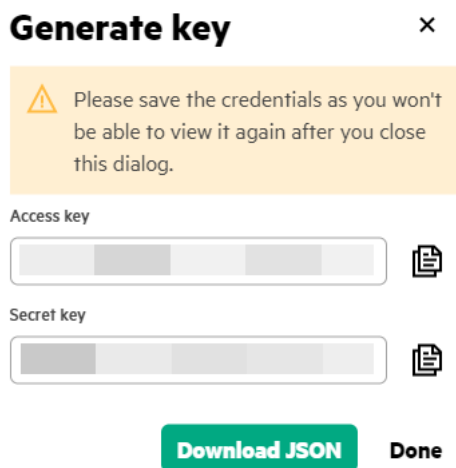
**External:** <https://3.101.79.104:9000> ⓘ

<https://3.101.79.104:9000> ⓘ

---

**Download**

- Click **Generate key**. A confirmation dialog box asks if you want to generate a new S3 key.
- Click **Generate key**. A dialog box displays the access key and secret key information. For example:



8. Click **Download JSON** if you want to download the keys for the global namespace as a JSON file.
9. Click **Download** if you want to download the S3 end points information as a JSON file.

## Setting a Quota for a User

Set quota for an individual user.

### About this task

You can set quotas for individual users via the Data Fabric UI.

The following quotas can be set.

- Advisory quota, which raises an alarm when the threshold is reached, but does not prevent further writes.
- Hard quota, which raises an alarm when the limit is reached and prevents further writes. which raises an alarm when the limit is reached and prevents further writes.

Advisory quota and hard quota can be expressed in megabytes (MB), gigabytes (GB), or terabytes (TB). GB is the default unit. The advisory quota must be less than the hard quota.

Follow the steps given below to set quota for a user.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric manager** option from the dropdown on the Home page.
3. Click **Fabric Administration**.
4. Select the **users** option in the dropdown for **Quota by**.
5. Click **Edit fabric quotas** on the Quotas card.
6. On the **Edit Fabric Quotas** dialog box, enter the values for user advisory quota, user hard quota, group advisory quota and group hard quota. Change the unit, as required.
7. Specify the **Fabric reserve limit**.
8. Click **Update**.

Alternatively, you can edit the user quota for an individual user from the **Settings** tab for a fabric . To get to the **Settings** tab,

- a) Select the **Fabric user** on the Home page.
- b) Click the table view icon on the Resources card and click the fabric name on the table view of resources.

### Results

The specified quota is saved for the individual user.



**NOTE:** The default quota specified for the fabric applies to other groups, unless a group-specific quota is set for any other group on the fabric.

## Setting a Quota for a Group

Set quota for an individual group.

### About this task

You can set quotas for individual groups via the Data Fabric UI. This quota overrides the default group quota set for the fabric.

The following quotas can be set.

- Advisory quota, which raises an alarm when the threshold is reached, but does not prevent further writes.
- Hard quota, which raises an alarm when the limit is reached and prevents further writes. which raises an alarm when the limit is reached and prevents further writes.

Advisory quota and hard quota can be expressed in megabytes (MB), gigabytes (GB), or terabytes (TB). GB is the default unit. The advisory quota must be less than the hard quota.

Follow the steps given below to set quota for a group.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric manager** option from the dropdown on the Home page.
3. Click **Fabric Administration**.
4. Select the **groups** option in the dropdown for **Quota by**.
5. Click **Edit fabric quotas** on the Quota by groups card.
6. On the **Edit Fabric Quotas** dialog box, enter the values for advisory quota, hard quota. Change the unit, as required.
7. Specify the **Fabric reserve limit**.
8. Click **Update**.

Alternatively, you can edit the group quota for an individual group from the **Settings** tab for a fabric. To get to the **Settings** tab,

- a) Select the **Fabric user** on the Home page.
- b) Click the table view icon on the Resources card and click the fabric name on the table view of resources.

## Results

The specified quota is saved for the individual group.



**NOTE:** The default quota specified for the fabric applies to other groups, unless a group-specific quota is set for any other group on the fabric.

## Viewing Fabric-Related Metrics

Explains the various fabric-related metrics visible on the [Data Fabric UI](#) on page 86.

### About this task

You can view various useful fabric-related metrics on the [Data Fabric UI](#) on page 86.

The metrics include data related to billing and storage consumption by fabric, storage usage by various users of the fabrics that are being monitored via the [Data Fabric UI](#) on page 86, and the top fabrics nearing the total storage capacity.

Click any of the following links to view details about the fabric-related metrics that are visible on the Home Page of the [Data Fabric UI](#) on page 86.

### View Storage Consumption by User

View fabric storage consumption by users.

### Prerequisites

- You must be a fabric manager or an infrastructure admin to perform this operation.
- Users and/or groups must have been created.

### About this task

View how much storage is used by individual fabric users via the Data Fabric UI.

You can determine the fabric storage consumption trends by users on the fabrics that are being monitored via the Data Fabric UI.

Viewing fabric storage consumption trends by groups is available if groups have been defined on the Data Fabric UI.

The storage consumption by a user/group is aggregated storage size of volumes and topics owned by the user/group.



**NOTE:** Bucket storage consumption is not included as there is no concept of bucket owner in Data Fabric.

You can import a fabric to monitor the fabric usage via the Data Fabric UI. See [Importing a Fabric](#) on page 110 for information importing fabrics.

### Procedure

1. Log on to the Data Fabric UI.
2. Select the **Fabric manager** from the dropdown next to the welcome message on the Home page.
3. Click **Fabric metrics** on the Home page.
4. Check the **Storage use by Users** card.

## Results

You are able to view a list of all fabric users and/or groups in the order of the storage utilization.



### View System Resource Utilization by Fabric

View system resource utilization for fabric.

#### Prerequisites

You must be a fabric manager or an infrastructure admin or a fabric user to perform this operation.

#### About this task

View a graphical representation of CPU utilization and memory utilization by fabric for the selected time duration.

#### Procedure

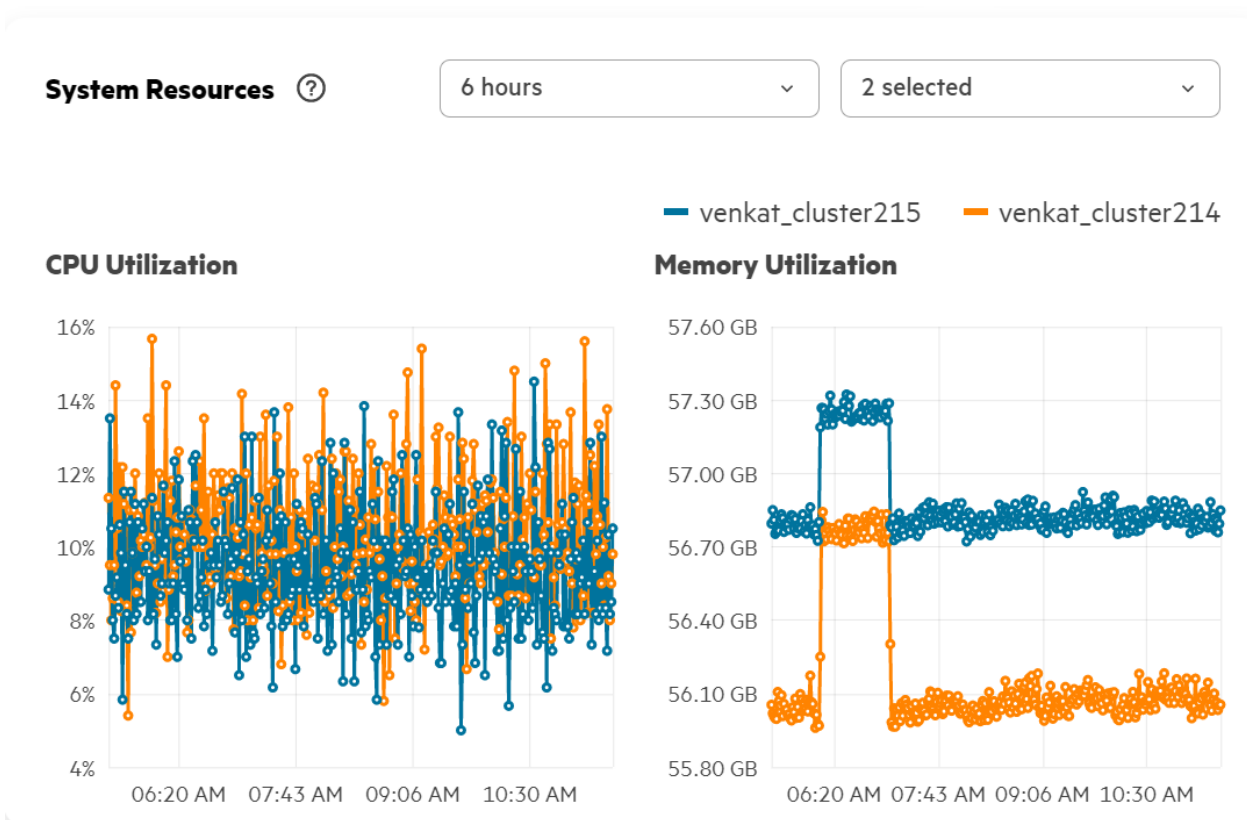
1. Log on to the Data Fabric UI.
2. If you are a fabric manager or an infrastructure admin, click and check the **Fabric utilization** card to view the CPU and memory utilization by fabric.
3. If you are a fabric user, scroll down the Home page to view the System Resources - CPU and memory utilization card.
  - a) Select the fabric and the time duration for which you wish to view the CPU and memory utilization of the fabric.

#### Results

For an infrastructure admin or a fabric manager, the CPU and memory utilization is seen as shown in the following image.

Name	CPU(Cores)	Memory(GB)
venkat_cluster214	1 of 32 (3%)	51 GB of 129 GB (38%)
venkat_cluster215	4 of 32 (13%)	52 GB of 129 GB (38%)
cluster-onprem-199	3 of 40 (8%)	47 GB of 129 GB (35%)

For a fabric user, you can view the system utilization by the fabric during the selected time duration in a graphical format. The following image shows CPU and memory utilization by the two selected fabrics for the last 6 hours.



### View Top Fabrics by Storage Capacity

View up to top five fabrics by consumption of the available storage capacity.

#### Prerequisites

- You must be a fabric manager or an infrastructure admin to perform this operation.
- One or more fabrics must have been created on or imported into Data Fabric.

#### About this task

View the top fabrics that have consumed maximum of the total storage available to the fabric.

You can use this data to understand what fabrics are nearing the total storage capacity available to the fabric.

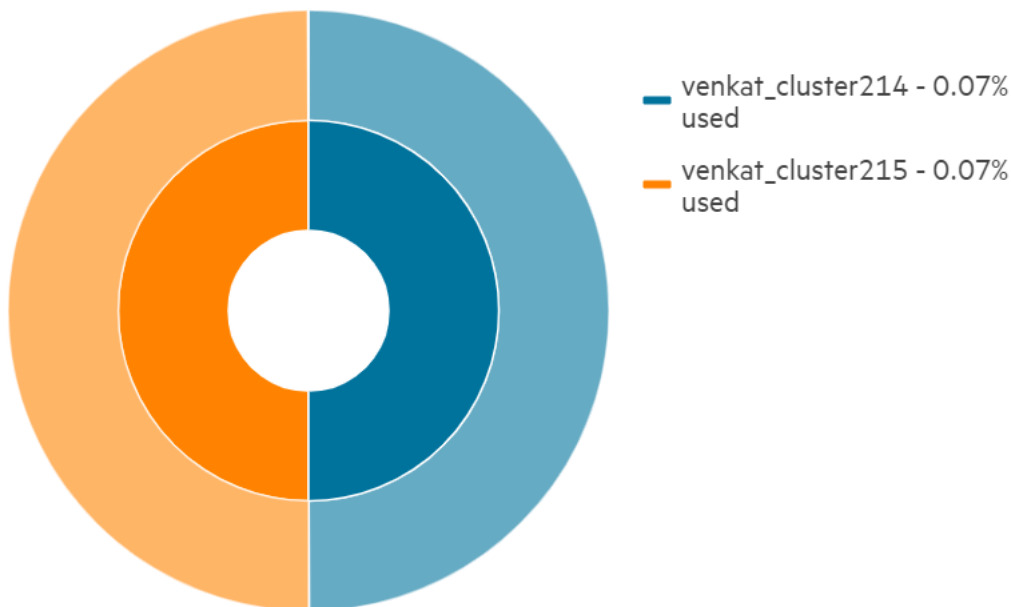
#### Procedure

1. Log on to the Data Fabric UI .
2. Click **Fabric metrics** on the Home page.
3. Scroll down to see the **Fabric Storage** card.

#### Results

You are able to see up to five fabrics that have consumed maximum storage capacity out of the available storage capacity available to the individual fabrics. If you have clicked the fabric area on the **Fabric Storage** card, you are navigated to the Resource page to be able to view the fabric resources.

### Fabric Storage ?



[View all](#)

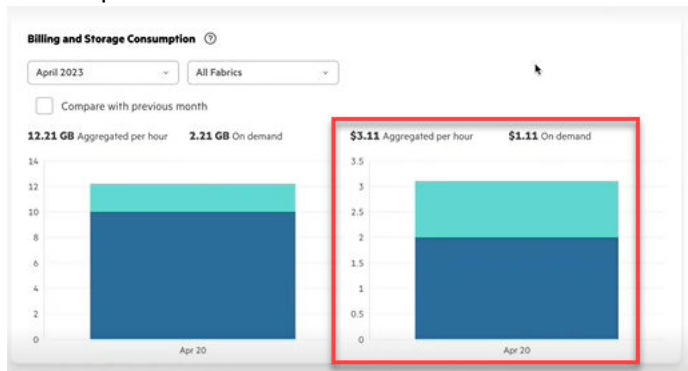
#### View Billing Data by Fabric

View estimated billing charges alongside storage consumption.

#### About this task

If you [set the billing model](#) as part of installing the HPE Ezmeral Data Fabric, the **Fabric Metrics** page displays a **Billing and Storage Consumption** card. The **Billing and Storage Consumption** card shows estimates of your storage consumption charges alongside your aggregated and on-demand storage consumption. The estimates are approximate and are not guaranteed to reflect your actual charges.

In the bar chart, dark blue represents consumption that is *below* the commit amount. Teal represents consumption *above* the commit amount.



Only storage consumption information is displayed if you are logged on as a Developer role. If you log on as an Infrastructure Admin or Fabric Manager, both billing and storage consumption information are displayed.

### Procedure

1. Sign in to the Data Fabric UI with Fabric Manager [credentials](#).
2. Click **Fabric metrics**.
3. Scroll down to see the **Billing and Storage Consumption** card.

### Results

You see the billing charges for the specified fabrics. If the charges are not displayed, it is likely that you have not yet [set the billing model](#).

## Setting Default Quotas for Users/Groups

Set default values for user and group quotas on a fabric via the Data Fabric UI.

### Prerequisites

You must be a fabric manager to edit user quotas and group quotas for a fabric.

### About this task

Quotas limit the disk space used by a volume or an entity such as a user or a group. A volume quota limits the space used by a volume. A user/group quota limits the space used by all volumes owned by a user or group. These quotas work on tenant volumes as well.

You can set hard quota and advisory quota defaults for users and groups. When a user or group is created, the default quota and advisory quota apply unless overridden by specific quotas. Quotas can be specified in mega bytes (MB), gigabytes (GB), terabytes (TB), petabytes (PB), exabytes (EB), and zettabytes (ZB).

- User quota is the total space allocated to user on fabric.
- User hard quota is the total space allocated to user on fabric.
- Group quota is the total space allocated to group on fabric.
- Group hard quota is the total space allocated to group on fabric.
- Fabric reserve limit is the percentage of the total cluster capacity to allocate for the volumes on the cluster.

The size of a disk space quota is expressed in terms of the actual data stored from the user's point of view. Only post-compression data blocks are counted, and snapshot and replica space do not count against quotas. For example, a 10G file that is compressed to 8G and has a replication factor of 3 consumes 24G (3\*8G), but charges only 8G to the user or volume's quota.

Follow the steps given below to set default quotas for users and/or groups.

### Procedure

1. Log on to the Data Fabric UI.
2. Select the **Fabric manager** from the dropdown next to the welcome message on the Home page.
3. Click
4. In the tabular list of fabrics on the **Global namespace** card, click the fabric name for which you wish to set the default user quota and/or group quota.
5. Click Settings. Click the pencil/edit icon next to the **Default quota**.

6. On the **Edit Fabric Quotas** dialog box, enter the values for user quota, user hard quota, group quota and group hard quota. Change the unit, as required.
7. Specify the **Fabric reserve limit**.
8. Click **Update**.

### Results

The user quota, group quota, and fabric reserve limit specified for the fabric are saved. You can view the fabric default user quota and group quota on the **Settings** tab for the fabric.

You may choose to specify quotas for individual users and/or groups. Such values override the default user quota and group quota for the respective users and/or groups. See [Setting a Quota for a User](#) on page 118 for details.

## Viewing the Fabric Service Status

View status of various services running on a fabric.

### Prerequisites

- You must have permissions to access and view the fabric details.

### About this task

A fabric functions as a set of core services and monitoring services that run in the background. You can view the status of the services on the Data Fabric UI to determine if the services are running or they have stopped. This can be useful in preliminary troubleshooting of the fabric operation.

Use the following steps to view fabric service status.

### Procedure

1. Log on to the Data Fabric UI
2. If you are a fabric user, click the Table View icon on the **Resources** card. If you are a fabric manager, select the **Fabric manager** option, and click **Global namespace** and check the table view.
3. Click the link for the fabric under the **Resource Name** column.
4. Navigate to the **Services** tab for the fabric.

### Results

The details about the various fabric-related services along with the status of each service is visible on the **Services** tab.

## View Capacity Usage by User on Fabric

Describes how to check the capacity used by various internal volumes, buckets, topics, and binary tables created by the user that is logged in to the Data Fabric UI.

### About this task

When you are logged in to the Data Fabric UI, you can view the total capacity for a fabric that is used up by the data stored on the volumes, buckets, topics, and binary tables created by you.



**NOTE:** The capacity displayed on the **My Capacity** card is related to the storage resources on the selected fabric only. The capacity displayed on the Data Fabric UI is exclusive of any resources such as volumes, buckets, topics, binary tables that you might have created on an external NFS server or external S3 server that has been added to the global namespace.

**Procedure**

1. Log on to the Data Fabric UI.
2. Select **Fabric user** on the Home page.
3. Check the **My Capacity** card.
4. Select the fabric for which you wish to view the storage capacity that is used up by the volumes, buckets, topics, and binary tables that you have created. Alternatively, select all fabrics if you wish to view the capacity usage on your fabric resources for all fabrics.

**Results**

You are able to view the total capacity usage under **Total**, along with a categorized break-up of the storage capacity used up by the volumes, buckets, topics, and binary tables created by you on the selected fabric, or collectively for all fabrics if you have chosen to view the data for all fabrics.

**SSH Access to a Cloud-Based Fabric**


Describes how to obtain a fabric-specific `.pem` file that enables SSH access to a cloud-based fabric.

Command-line access to cloud-based fabrics (AWS, Azure, or GCP) requires you to download a fabric-specific `.pem` file. The Data Fabric UI makes it easy to download the file.

Note these considerations for downloading the `.pem` file:

- Only a user with the Fabric user or Fabric Manager role (or `fc` access) can use the **Download SSH keys** command.
- To connect to the fabric, you must provide the public IP address or public DNS name of any cloud fabric node. The public IP address or public DNS name are contained in the URL that you use to access the Data Fabric UI. The URL was provided when you performed the seed node installation to create your first fabric.
- SSH access to fabrics should only be used for troubleshooting operations under the supervision of HPE support personnel. SSH access should not be used for daily operations.

Use these steps to download the `.pem` file:

1. Sign in to the Data Fabric UI as a Fabric manager or Fabric user. If you are a Fabric manager, switch to the **Fabric user** experience.
2. In the **Resources** card, click **Table view**.
3. Under the **Action** column, click the ellipsis (  ).
4. Click **Download SSH key**. The Data Fabric UI downloads the `.pem` file as `<fabric_name>_key.pem`.
5. If necessary, copy the file to the workstation that you will use to ssh to the fabric. Suppose you copy the file to `/root/myfabric-keypair.pem`. Remember the path to that location.
6. Reset the permissions on the downloaded `.pem` file to `0400`:

```
chmod 0400 <pem-file-name>
```

- Use one of the following commands to connect to the fabric:

#### AWS or GCP

```
ssh -i "<pem-file-name>" rocky@<public-IP-addr-or-public-DNS-name>
```

#### Azure

```
ssh -i "<pem-file-name>" mapr@<public-IP-addr-or-public-DNS-name>
```

## Deleting a Fabric

Delete a remote fabric from the global namespace.

### Prerequisites

You must be a fabric manager to delete a fabric.

You must delete fabrics one at a time. You cannot delete multiple fabrics at the same time.

To delete a remote fabric, you must have logged in with your single sign-on credentials.

### About this task

A local fabric is the fabric by which you are logged onto the Data Fabric UI. Any fabric, other than the local fabric, that has been added to the global namespace, is a remote fabric.

You can delete a remote fabric via the Data Fabric UI. You cannot delete a local fabric. If no remote fabrics are present, the **Delete** command is not available.



**CAUTION:** When you delete a fabric, the fabric is uninstalled, and any running instances of the fabric are destroyed. The data on a fabric becomes inaccessible when you delete the fabric. Ensure that you do not need the data on the fabric or the required data is backed up, before you delete the fabric.

Follow these steps to delete a fabric:

### Procedure

- Log on to the Data Fabric UI.
- Select **Fabric manager** from the dropdown on the Home page.
- Click **Global namespace**.
- Click the ellipsis under **Actions** for the fabric to delete from the global namespace.
- Click **Delete**. Confirm the fabric deletion.

### Results

The fabric is deleted, removed, and uninstalled from the global namespace.


## Administering Identities

---

This section describes the operations you can perform as an SSO user on identities, that is, IAM policies, SSO users, SSO groups, and roles for the HPE Ezmeral Data Fabric.

Only SSO users can access and use the identity management feature by using the Data Fabric UI.

You can assign SSO users, SSO groups, and identity access management (IAM) policies to a role.

 **NOTE:** The creation and management of SSO user and SSO group is done from the SSO provider console. See your SSO provider documentation for details.

See the individual section about roles, users, groups, IAM policies for details about managing these identities.

## About Roles

Describes roles in Data Fabric

A role comprises SSO users and/or SSO groups, along with the associated permissions on various Data Fabric resources.

The types of roles in Data Fabric are as follows:

- Pre-defined roles

 **NOTE:** Only pre-defined SSO provider roles are recognized and used by the Data Fabric UI

- User-defined roles

An IAM policy can be assigned to or tagged to a user-defined role.

You can create and manage user-defined roles using the Data Fabric UI.

### Pre-defined Roles

The pre-defined roles are as follows:

- Infrastructure administrator
- Fabric manager
- Fabric user or developer

Following are the characteristics of pre-defined roles.

- Predefined roles are made available as a part of the fabric deployment process, and stored in SSO service provider.
- Pre-defined roles have a fixed set of permissions.
- Pre-defined role are available on successfully deploying the first fabric.
- You cannot modify or delete pre-defined roles.

See [Pre-defined Roles and Associated Permissions](#) on page 129 for the permissions assigned to each of the pre-defined roles.

### User-defined Roles

A user-defined role is an identity that is associated with [IAM policies](#) and one or more SSO users and/or SSO groups in Data Fabric.

Following are the characteristics of user-defined roles:

- You can create and manage user-defined roles in Data Fabric.
- You can create user-defined roles after fabric deployment.
- You can create, modify, and delete user-defined roles, as required.
- You can define permissions in user-defined roles through IAM policies.



## Pre-defined Roles and Associated Permissions

This page describes the roles supported by the HPE Ezmeral Data Fabric as-a-service platform.

### Roles and Permissions When SSO Is Configured

SSO-configured fabrics support the following pre-defined roles:

Role	Permissions	Corresponding ACL Permission Code <sup>1</sup>
Infrastructure Admin	Permission to log in and start or stop services	login, ss
Fabric Manager	Full control of the fabric, create volume permission, and login permission <sup>2</sup>	login, cv, cp, fc
Fabric User	Login permission <sup>2</sup> and create volume permission	login, cv, cp

<sup>1</sup>Shows the equivalent access control list (ACL) permission code for the HPE Ezmeral Data Fabric – Customer Managed cluster. For more information, see [Security Policy Permissions](#) on page 218 and [Creating Cluster-Level ACLs](#).

<sup>2</sup>The login user can log in to the Data Fabric UI and issue commands. Includes read access for existing objects.

### Resource Actions Supported for the Roles

The following table shows the create, delete, and modify actions that each role can perform on various resources:

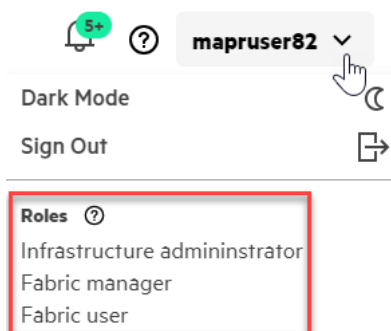
Role	Resource	Create	Delete	Modify
Fabric Manager	Fabric	Allow	Allow	Allow
	Volumes	Allow	Allow	Allow
	Buckets	Allow	Allow	Allow
	Directories	Allow	Allow	Allow
	User	Allow	Allow	Allow
	Accounts	Allow	Allow	Allow
	Groups	Allow	Allow	Allow
	S3 Keys	Allow	Allow	Allow
	Objects	Allow	Allow	Allow
	Security Policies	Allow	Allow	Allow
	Storage Policies	Allow	Allow	Allow
	Storage Tiers / Remote Targets	Allow	Allow	Allow
	SMTP Configuration	Allow	Allow	Allow
	IAM Policies	Allow	Allow	Allow
User-defined role	Allow	Allow	Allow	

Role	Resource	Create	Delete	Modify
Infrastructure Admin	<b>Resource</b>	<b>Create</b>	<b>Delete</b>	<b>Modify</b>
	Fabric	Deny	Deny	Allow
	Volumes	Deny	Deny	Deny
	Buckets	Deny	Deny	Deny
	Directories	Deny	Deny	Deny
	User	Deny	Deny	Deny
	Accounts	Deny	Deny	Deny
	Groups	Deny	Deny	Deny
	S3 Keys	Allow	Allow	Allow
	Objects	Deny	Deny	Deny
	Security Policies	Deny	Deny	Deny
	Storage Policies	Deny	Deny	Deny
	Storage Tiers / Remote Targets	Deny	Deny	Deny
	SMTP Configuration	Deny	Deny	Deny
Fabric User	<b>Resource</b>	<b>Create</b>	<b>Delete</b>	<b>Modify</b>
	Fabric	Deny	Deny	Deny
	Volumes	Allow	Allow	Allow
	Buckets	Allow	Allow	Allow
	Directories	Allow	Allow	Allow
	S3 Keys	Allow	Allow	Allow
	Objects	Allow	Allow	Allow
	Security Policies	Allow	Allow	Allow
	Storage Policies	Allow	Allow	Allow
	Storage Tiers / Remote Targets	Allow	Allow	Allow
	SMTP Configuration	Deny	Deny	Deny

### Displaying Role Information

To display role information for the currently signed-in user:

1. Sign in to the Data Fabric UI.
2. In the upper right corner of the home screen, click the down arrow next to the user name. For example:



### Limitation for Non-SSO Users

SSO users with sufficient credentials can view and manage resources on all fabrics. Non-SSO users can view and manage resources only on the fabric to which they are signed in. Non-SSO users cannot view or manage resources on other fabrics. The Data Fabric UI does not display these resources to non-SSO users because the UI cannot connect to other fabrics without the same login information.

## Creating a Role

Describes how to create a user-defined role.

### Prerequisites

The following prerequisites must be satisfied before you can create a user-defined role.

- You must be a fabric manager to be able to create a user-defined role.
- Single sign-on must be enabled to Data Fabric.

### About this task

By default, Data Fabric has pre-defined roles that can be assigned to SSO users and SSO groups in Data Fabric.

If you need a user-defined set of permissions to apply to SSO users or SSO groups, you can configure user-defined roles and associate such roles with the required set of users.

A user-defined role can be attached to one or more SSO users and/or groups.



**NOTE:** Although you can create a role without associating IAM policies and users/groups with the role, it is a best practice to create related users/groups and IAM policies, before you create a role.

A user-defined role must have a unique name.

User-defined roles are shared across all fabrics in a global namespace.



**NOTE:** User-defined roles that are created using the Data Fabric UI are not visible on the Keycloak console.

Follow the steps given below to create a new role.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric Manager** for the fabric manager view.
3. Click the **Security Administration** tab.

4. On the **Roles** card, click Create New Role. Alternatively, click **View All** on the **Roles** card and then click **Create new role**.
5. Enter the **Name** and **Description** for the role.
6. Click **Add user and group** and then click **Add+** seen above **User** or **Group**.
7. Search and select one or more users and then click **Add** to add all selected users. Repeat in a similar way for groups, if you are adding groups.
8. Click **Assign Policy**. Search and select one or more IAM policies that the role is to be tagged with or associated with.
9. Click **Apply**.

### Results

The new role is created. The newly created role is visible on the **Roles** card..

## Viewing Roles

Describes how to view all roles in Data Fabric.

### Prerequisites

You must be a fabric manager to perform this operation.

### About this task

You can view all roles that have been configured in Data Fabric.

- The pre-defined and user-defined roles are displayed on the Data Fabric UI.
- Any roles created in the SSO provider, other than the predefined roles, are not visible on the Data Fabric UI.
- User defined roles that are created by using the Data Fabric UI or maprccli command line are listed on the Data Fabric UI.

Follow the steps given below to view a list of all existing roles.

### Procedure

1. Log on to the Data Fabric UI.
2. Click the **Security Administration** tab.
3. On the **Roles** card, click **View All**.

### Results

A list of all existing roles, consisting of both pre-defined roles and user-defined roles, is displayed.

## Editing a Role

Describes how to edit a user-defined role.

### Prerequisites

You must be a fabric manager to perform this operation.

### About this task

user-defined roles cannot be renamed.

You can modify the role description, the users and/or groups, and the IAM policies that are associated with a user-defined role.

Follow the steps given below to edit a user-defined role.

### Procedure

1. Log on to the Data Fabric UI.
2. Click **Security Administration** on the Home page.
3. Scroll down to the **Roles** card.
4. On the list of roles, click the ellipsis under **Actions** for the user-defined role to edit.
5. Click **Manage role** to make changes to the user-defined role.
6. Make the required changes.
7. Click **Save**.

### Results

The changes are saved and applied to the user-defined role.

## Deleting a Role

Describes how to delete a role.

### Prerequisites

You must be a fabric manager to perform this operation.

### About this task

You can delete a user-defined role only when there is no user or group attached to the role.



**NOTE:** You cannot delete pre-defined roles.

Follow the steps given below to delete a user-defined role.

### Procedure

1. Log on to the Data Fabric UI.
2. Click **Security Administration** on the Home page.
3. Scroll down to the **Roles** card.
4. On the list of roles, click the ellipsis under **Actions** for the user-defined role to delete.
5. Click **Delete** and type DELETE in the provided text area to confirm deletion.
6. Click **Save**.

### Results

The selected user-defined role is deleted.

## Assigning a Role to a User

Describes how to assign a role to a Data Fabric SSO user.

### Prerequisites

You must be a fabric manager to perform this task.

### About this task

You can assign one or more roles to an SSO user by using the Data Fabric UI.

Note the following with respect to role assignment from the Data Fabric UI.

- The Data Fabric UI supports the assignment of user-defined roles to SSO users/groups only. Assignment of SSO roles to SSO user/groups must to be done by using the SSO provider console or UI.
- The Data Fabric UI recognizes the assignment of roles to users/groups assignment made through SSO pre-defined roles only.
- Assignment of SSO users/groups to user-defined roles is facilitated but is not enforced, by using the Data Fabric UI.

To assign a role to a Data Fabric user, perform the following steps.

1. Sign in to the Data Fabric UI, and switch to the **Fabric manager** experience.
2. Click **Security administration**.
3. Scroll down to the **Users Roles** card.
4. Click the pencil icon seen under **Actions** for the role that you wish to assign to the user.
5. Select **Manage Role**.
6. Click **Add user and group**.
7. Click **Add+** for Users and select the required user or users.
8. Click **Assign Policy** and then click **Apply**.

### Results

The user is assigned the selected roles.

## Assigning a Role to a Group

Describes how to assign a role to a SSO group via the Data Fabric UI.

### Prerequisites

You must be a fabric manager to perform this task.

### About this task

To assign a role to a SSO group, perform the following steps.

1. Sign in to the Data Fabric UI, and switch to the **Fabric manager** experience.
2. Click **Security administration**.
3. Scroll down to the **Roles** card.
4. Click the pencil icon seen under **Actions** for the role that you wish to assign to a group.

5. Select **Manage Role**.
6. Click **Add user and group**.
7. Click **Add+** for Group and add the required group or groups.
8. Click **Assign Policy** and then click **Apply**.

### Results

The selected role is assigned to the group.

## Viewing a List of Users

Describes how to display a searchable list of SSO users that includes the names of the users and their roles.

The **Users** card can display only a small number of SSO users. Clicking the **View all** button displays a searchable, configurable, full-page listing of SSO users who have access to the Data Fabric UI. If you are a fabric manager, you can also edit the roles for a user.

Note the following with respect to the Data Fabric UI:

- The **Users** card and the **User Identity** tab display user information available through the SSO provider. (Keycloak is currently the only supported SSO provider).
- Non-SSO users are not listed here.
- Only pre-defined roles that are assigned to SSO users are displayed on the **Users** card or the **User Identity** tab. A user-defined role assigned to an SSO user is not visible on the **Users** card or the **User Identity** tab.

To display the full list of SSO users:

1. Sign in to the Data Fabric UI, and switch to the **Fabric manager** or **Infrastructure admin** experience.
2. Click **Security administration**.
3. Scroll down to the **Users** card.
4. Click **View all**. The list of SSO users is displayed.

## Viewing a List of Groups

Describes how to display a searchable list of all SSO groups that includes the names of the groups and their roles.

### Prerequisites

You must be a fabric manager to perform this task.

### About this task

The **Groups** card can display up to five SSO groups on the Groups card. Clicking the **View all** button displays a searchable, configurable, full-page listing of all the existing SSO groups.

Note the following with respect to the Data Fabric UI:

- The **Groups** card and **Group Identity** tab display SSO user and SSO group information only.
- Only pre-defined roles are displayed on the **Groups** card and **Group Identity** tab. Any user-defined roles are not visible here, even if they are assigned to SSO users/SSO groups.

To display the full list of SSO groups, perform the following steps:

1. Sign in to the Data Fabric UI, and switch to the **Fabric manager** experience.
2. Click **Security administration**.
3. Scroll down to the **Groups** card.
4. Click **View all**.


### Results

The list of SSO groups is displayed.

## Administering IAM Policies

Provides an overview of IAM policies in Data Fabric.

### Identity Access Management Policy

 **IMPORTANT:** An IAM policy can be assigned to a user-defined role only. It is mandatory to have single sign-on enabled on the Data Fabric to be able to use the IAM policy feature.

An identity access management (IAM) policy is a security mechanism that states actions that can or cannot be performed by identities such as SSO users or groups on one or more resources that belong to one or more fabrics in a global namespace. An IAM policy defines allowable and disallowable actions on volumes, objects, and tables on more than one fabrics in a global namespace.

A fabric manager can create, modify, view, and delete IAM policies. A fabric manager can assign IAM policies to roles and vice-versa. A fabric manager can manage IAM policies from the Data Fabric UI, regardless of whether the fabric manager is logged on to the Data Fabric UI from the primary fabric or a non-primary fabric.

Use of IAM policies is recommended if you wish to associate a common set of allowable/disallowable actions on multiple disparate fabric resources at one go, that is, volumes, objects, tables in a global namespace for one or more SSO users and/or SSO groups.

#### TIP:

- Security policy can be configured exclusively for volumes belonging to a single fabric.
- Bucket policies can be applied exclusively to buckets on an object store from a fabric.
- A security policy or a bucket policy is associated with resources, while identity access management policy is associated with identities such as SSO users/groups and roles. As they apply to separate entities, security policies, bucket policies, and IAM policies can co-exist simultaneously.

### About IAM Policy

Introduction to identity access management (IAM) policy.

An identity is an entity such as a user, a group or a role. A policy that defines the permissions on Data Fabric resources such as volumes, buckets, and database tables can be tagged to an identity.

An identity policy could apply to:

- one or more resources from multiple clusters
- multi modal resources (volumes/objects/tables)

All resources in a fabric must be represented with consistent naming scheme that is known as Unique Resource Name (URN) to uniquely locate/identify a specific resource.



As a fabric manager, you can create IAM policies by using the Data Fabric UI.

### Identity Access Management Policy Life Cycle

Describes the life cycle of identity management policy.

An identity access management (IAM) policy goes through the following stages during its life cycle.

It is important to remember the following with respect to enforcement of an IAM policy:

- An IAM policy must be active and assigned to be enforceable on the SSO users and/or SSO groups, and roles. Similarly, an IAM policy must be inactive to be able to disarm the IAM policy.
- The time taken to enforce an IAM policy is anywhere between 2 and 30 minutes in a multi-cluster environment, depending on the number of identities involved in the operation.

Policy State	Policy State Description	Supported Transitions to other Policy States
Unassigned	Policy is yet to be assigned to an identity	Policy can be assigned, modified, or deleted
Assigned	Policy is assigned to an identity but not yet enforced	Policy can be unassigned from the current fabric resources to which it is assigned, modified, and deleted
Enforcing	Policy enforcement is in progress	The IAM policy cannot be assigned, unassigned, modified or deleted during the transient phase.
Enforced	Policy enforcement is complete	The IAM policy can be assigned, unassigned, modified, or deleted to other fabric resources.
Disarming	Policy deactivation is in progress	Policy cannot be assigned, unassigned, modified, or deleted.
Disarmed	Policy is deactivated	The IAM policy can be assigned, unassigned, modified, or deleted.

### Resource-level Permissions in an IAM Policy

Describes various resource-level permissions that can be allowed or denied in an IAM policy.

You can configure the resource-level permissions at the following levels:

- Fabric level
- Bucket level
- Volume level

The following sections describe permissions at each of the aforementioned levels.

### Fabric-level Permissions

You can configure the following permissions in an IAM policy for fabrics, external S3 servers, or external NFS servers.

Permission	Description
ViewClusterConfig	Permission to view cluster configuration.
ManageClusterServices	Permission to manage various fabric-level services.
ManageClusterOperations	Permission to manage cluster operations.
ManageClusterSettings	Permission to manage various cluster settings.

Permission	Description
ManagemRoleOperations	Permission to manage operations related to user-defined roles.
ManageClusterACE	Permission to manage fabric <a href="#">ACE</a> .
ManageStartStopService	Permission to start and stop fabric services
ManageClusterVolume	Permission to manage volumes on a fabric.

### Volume-level Permissions

The following permissions can be granted in an IAM policy for volumes.

Permission	Description
ReadVolume	Permission to read a volume.
WriteVolume	Permission to write a volume.
DeleteVolume	Permission to delete a volume.
MountVolume	Permission to mount a volume.
MirrorVolume	Permission to mirror a volume.
ManageVolumeConfig	Permission to manage volume configuration.
ManageVolumeACE	Permission to manage volume <a href="#">ACE</a> .
VolumeFullControl	Permission to perform all allowable operations on a volume.

### Bucket-level Permissions

The following permissions can be granted in an IAM policy for Data Fabric S3 buckets.

Permission	Description
AbortMultiPartUpload	Permission to abort multi-part upload of an object to S3 bucket
DeleteBucket	Permission to delete an S3 bucket
ForceDeleteBucket	Permission to force an S3 bucket deletion
DeleteBucketPolicy	Permission to delete an S3 bucket policy
DeleteObject	Permission to delete object from an S3 bucket
GetBucketLocation	Permission to get location or region for an S3 bucket
GetBucketNotification	Permission to get notification for an S3 bucket
GetBucketPolicy	Permission to get the policy of an S3 bucket.
GetObject	Permission to retrieve an object from S3 server/bucket.
HeadBucket	Permission to access S3 bucket to check for its existence and contents
ListAllMyBuckets	Permission to retrieve a list of S3 buckets owned by the sender of the request
ListBucket	Permission to retrieve a list of S3 buckets
ListBucketVersions	Permission to retrieve a list of S3 bucket versions
ListBucketMultiPartUploads	Permission to retrieve a list of the multi part uploads for an S3 bucket.
ListMultiPartUploadParts	Permission to retrieve list of parts in a multi-part upload into an S3 bucket

Permission	Description
PutBucketLifeCycle	Permission to create a new lifecycle configuration for S3 bucket or replaces an existing lifecycle configuration.
GetBucketLifeCycle	Permission to retrieve lifecycle configuration for S3 bucket
PutBucketNotification	Permission to enable notifications for specified events related to S3 bucket.
PutBucketEncryption	Permission to configure encryption and keys on an S3 bucket
DeleteObjectTagging	Permission to delete object tagging
PutBucketPolicy	Permission to apply bucket policy to an S3 bucket
PutObject	Permission to add object to S3 bucket
PutObjectRetention	Permission to configure object retention settings on an object
GetObjectRetention	Permission to retrieve object retention configuration on an object
GetObjectLegalHold	Permission to retrieve the legal hold status for an object
PutObjectLegalHold	Permission to configure legal hold for an object.
GetBucketObjectLockConfiguration	Permission to retrieve the object lock configuration for an S3 bucket
PutBucketObjectLockConfiguration	Permission to configure the object lock settings for an S3 bucket
GetBucketTagging	Permission to retrieve the tags associated with an S3 bucket
PutBucketTagging	Permission to set tags for an S3 bucket
GetObjectVersion	Permission to access a specific version of an object
GetObjectVersionTagging	Permission to retrieve tag of an object version
DeleteObjectVersion	Permission to delete an object version
DeleteObjectVersionTagging	Permission to delete an object version tagging
PutObjectVersionTagging	Permission to set a tag for an object version
GettObjectTagging	Permission to retrieve the set of tags for an object
PutObjectTagging	Permission to set the tags for an object
GetBucketEncryption	Permission to retrieve the encryption settings for an S3 bucket
PutBucketVersioning	Permission to set the versioning state for an S3 bucket
GetBucketVersioning	Permission to retrieve the versioning state for an S3 bucket
GetReplicationConfiguration	Permission to retrieve replication configuration
PutReplicationConfiguration	Permission to set replication configuration

### Creating an IAM Policy

Describes the procedure to create an IAM policy.

### Prerequisites

You must be a fabric manager to perform this operation.

### About this task

You can create an IAM policy that comprises one or more statements for different resource types. Resources are of the following types:

- Fabric

- Volumes
- Buckets

For instance, you can allow a set of fabric-level permissions for multiple fabrics, including external S3 servers and/or external NFS servers by selecting the resource type as fabric for your statement. Another statement can be added for denial of various actions or operations related to specified buckets for the selected fabrics in the same policy.

An IAM policy is enforced when it is active and is attached or assigned to an identity like an SSO user, SSO group, and/or a user-defined role. The selected actions are allowed on the selected resources to the SSO users and/or SSO groups, when the policy is tagged/assigned to a user-defined role.

Follow the steps given below to create a IAM policy.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric Manager** for the fabric manager view.
3. Click the **Administration** tab.
4. On the **IAM Policies** card, click **Create Policy**.
5. Enter the **Name** and **Description** for the IAM policy.
6. Turn the **Active** toggle off, if you wish to make the policy inactive. The **Active** toggle is on, by default.
7. Click **Manage Resources** and select the resource and the respective operations to allow or deny permissions on.
8. Click **Add** for **Statements** to add a statement to the policy.
9. Select the **Resource type** . Select fabric to add fabric-level permissions, volume to add permissions specific to volumes, or bucket to add permissions specific to buckets.<add link>.
10. If you have chosen volume or bucket as the resource, select one or more Fabrics from the **Fabric** dropdown to which the volumes or buckets belong, and then click **Apply**.
11. Click **Add** for **Selected Resources**, select one or more fabrics or volumes or buckets, depending the resource type selected.
12. Select the resource-type specific actions that are to be allowed or denied on the selected resources.
13. Select the **Effect**. Select **Deny** to deny the selected actions on the selected resources. **Allow** is the default value.
14. Repeat steps 7 through 12 to add more statements, if required.
15. Click **Save**.

### Results

An IAM policy is created. The newly created IAM policy is visible under the list of policies on the IAM policies card. The policy can be assigned to one or more identities such as SSO users, SSO groups, and/or user-defined roles.

### Editing an IAM Policy

Describes the procedure to edit an identity access management policy

**Prerequisites**

You must be a fabric manager to perform this operation

**About this task**

You can edit an IAM policy to add, change, or remove the assignment of the IAM policy with users/groups/roles.

You can also deactivate or activate an IAM policy.

Follow the steps given below to edit a IAM policy.

**Procedure**

1. Log on to the Data Fabric UI.
2. Select **Fabric Manager** for the fabric manager view.
3. Click the **Administration** tab.
4. On the list of IAM policies seen on the **IAM policies** card, click the ellipsis under **Actions** for the IAM policy to edit.
5. Click **Manage IAM policy** to make changes to the IAM policy.
6. Make the required changes to the various statements in the IAM policy.
7. Click **Save**.

**Results**

The changes are saved to the IAM policy and are reflected across the resources to which the IAM policy has been assigned.

**Deleting an IAM Policy**

Describes how to delete an IAM policy.

**Prerequisites**

You must be a fabric manager to perform this operation.

**About this task**

You can delete IAM policies that are not assigned to any other identity type such as users/groups/roles.

Follow the steps given below to delete a IAM policy.

**Procedure**

1. Log on to the Data Fabric UI.
2. Select **Fabric Manager** for the fabric manager view.
3. Click the **Administration** tab.
4. On the list of IAM policies seen on the **IAM policies** card, click the ellipsis under **Actions** for the IAM policy to delete.
5. Click **Delete** and type DELETE in the provided text area to confirm deletion.

**Results**

The IAM policy is deleted.

### Viewing All IAM Policies

Describes how to view a list of the existing IAM policies.

#### Prerequisites

You must be a fabric manager to perform this operation.

#### About this task

Follow the steps given below to view all existing IAM policies.

#### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric Manager** for a fabric manager view.
3. Click the **Administration** tab.
4. On the **IAM policies** card, click **View all** to view a list of all IAM policies that have been defined.

#### Results

All existing IAM policies are displayed.

### Assigning an IAM Policy

Describes how to assign an IAM policy to an identity by using the Data Fabric UI.

#### Prerequisites

You must be a fabric manager to perform this operation.

#### About this task

When you assign an IAM policy to one or more roles, SSO users, and/or SSO groups, the statements in the IAM policy are applied to the role, SSO users, and/or SSO groups.

An IAM policy is enforced when it is active and is attached or assigned to an identity like a user, group, and/or role. If an IAM policy is inactive, it will not be disarmed, even if it is assigned to an identity.



**NOTE:** After assigning an IAM policy, you can navigate out of the page using the breadcrumb on the top left side of the page.

Follow the steps given below to assign an IAM policy.

#### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric Manager** for the fabric manager view.
3. Click the **Administration** tab.
4. On the list of IAM policies seen on the **IAM policies** card, click the ellipsis under **Actions** for the IAM policy to edit.
5. Click **Assign Policy to** to assign the IAM policy to another identity such as users/groups/roles.
6. Click **Add+** on the **Users** card, enter the search criteria and select the users that you wish to assign the IAM policy to. This is an optional step.

7. Click **Add+** on the **Groups** card, enter the search criteria and select the groups that you wish to assign the IAM policy to. This is an optional step.
8. Click **Add+** on the **Roles** card, enter the search criteria and select the roles that you wish to assign the IAM policy to. This is an optional step.

### Results

The policy is assigned to the selected SSO users, SSO groups, and/or roles. If the policy is active, it is enforced for the selected SSO users, SSO groups, and/or roles.

## Configuring Email Notifications


Describes how to configure the Simple Mail Transfer Protocol (SMTP) to send email notifications from the Data Fabric UI to specified email accounts.

The Data Fabric UI can notify you by email when alarms are generated on a fabric. To configure email notifications, you must set up SMTP:

### Setting Up SMTP

To set up SMTP:

1. Sign in to the Data Fabric UI, and switch to the **Fabric manager experience**.
2. Click **Fabric administration**.
3. On the **SMTP** card, click **Edit SMTP settings**. The **Edit SMTP settings** form is displayed.
4. Specify the following parameters:

Parameter	Description	Example
<b>Provider*</b>	Select <b>Office 365</b> , <b>SMTP</b> , or <b>Other</b> from the drop-down menu. If you select <b>Office 365</b> , the SMTP server and port information is pre-filled for you.   <b>NOTE:</b> Gmail is provided as an option, but is not currently supported because Gmail does not support unsecure emails from third-party applications. For more information, see <a href="#">this page</a> .	Office 365
<b>SMTP server*</b>	The name of the mail server for the SMTP provider that you specified.	smtp.office365.com
<b>This server requires an encrypted connection (SSL)</b>	Check this box if the connection to the SMTP server must be encrypted.	N/A
<b>SMTP port*</b>	The SMTP port to use for sending mail.	587
<b>Sender's full name*</b>	The name that the HPE Ezmeral Data Fabric should use when sending email.	East Lab Data Fabric
<b>Sender's email address*</b>	The email address that the HPE Ezmeral Data Fabric should use when sending email.	jennifer-huang87@outlook.com
<b>Sender's username</b>	(Optional) The user name that the HPE Ezmeral Data Fabric should use when logging on to the SMTP server.	jennifer46
<b>Sender's SMTP password</b>	(Optional) The password that the HPE Ezmeral Data Fabric should use when logging on to the SMTP server.	mySMTP!pw


5. Click **Save**. A message indicates if the configuration was successful.
6. See [Setting Up Alarm Notifications](#) on page 144.

### Editing SMTP Settings

After SMTP has been configured, you can edit the settings by clicking **Edit SMTP settings**, changing the parameter values as needed, and clicking **Save**.

### Setting Up Alarm Notifications

Setting up SMTP does not by itself enable alarm notifications. You must also identify the alarms for which you want to be notified. Currently, setting up alarms must be done using the `maprcli` command line. You must run the `config save` command for each alarm where you want to generate an email.

 **WARNING:** You must have `fc` (full control) or `a` (admin) permissions to run this command.

The format of the command is:

```
maprcli alarm config save -cluster <fabric_name> -values
"<alarm>,<enableEmail>,<email>"
```

Assign values as follows:

Value	Description	Example
alarm	Name of the alarm. Specify the alarm name in uppercase with underscores. For a list of Data Fabric alarms, see <a href="#">Alarms Reference</a> .	DISK_FAILURE_ALARM
enableEmail	Specifies whether individual alarm notifications are sent to any email address (including the default email address) for the alarm type: <ul style="list-style-type: none"> <li>• 0 – Do not send notifications to any email address for the alarm type.</li> <li>• 1 – Send notifications to all email addresses for the alarm type.</li> </ul>	1
email	One or more email addresses other than the default email address. If specified, alarm notifications are sent to these addresses as well, if <code>enableEmail</code> is set to 1. Multiple email addresses must be separated by spaces only. You cannot use commas or other delimiters. For example, <code>user1@mycorp.com user2@mycorp.com</code> is valid.	jennifer-huang87@outlook.com

### Example

The following example command configures an email to be sent to `test@example.com` whenever the **Node Alarm Core Present** alarm is generated:

```
maprcli alarm config save -values
"NODE_ALARM_CORE_PRESENT,1,test@example.com"
```

### Related tasks

- [Viewing Alarms](#) on page 250
- View alarms on the [Data Fabric UI](#) on page 86.

## Viewing and Editing Access Control Information

Describes how to find and use the Access Control card that shows the access privileges for users and groups.

Note the following prerequisites for viewing and changing access control information:



- You must have fabric manager permissions to view or change access control settings.
- The user or group for which you want to assign access must already be configured for the fabric. To add users or groups, see [Adding New Users to Keycloak](#) on page 54 or [Adding a Group to Keycloak](#) on page 60.

### Viewing the Access Control Card


To view the **Access control** card:

1. Sign in to the Data Fabric UI, and switch to the **Fabric manager** view.
2. Click **Fabric administration**. The **Access Control** card appears under the Fabric administration details.

### Changing Access Control Settings

To add or change the access control information for a user or group:

1. On the **Access control** card, click **Edit access**. The Data Fabric UI displays the current settings.
2. Refer to the following table to change access settings for a user or group:

To	Do this
Add a new user or group	Click <b>+Add</b> , specify the <b>Type</b> (User or Group), and select the desired access options.
Change the access for an existing user or group	Select or deselect the the desired access options.
Remove a user or group	Click the garbage can icon (  ).

3. Click **Save** to save your changes, or click **Close** to exit without saving changes.

### Access Control Expression Syntax

This topic explains access control expression.

An [access control expression \(ACE\)](#) on page 339 is defined by a combination of user, group, or *role* definitions. You can combine these definitions using the following syntax:

Operator	Description
u	Username or user ID, as they appear in <code>/etc/passwd</code> , of a specific user. Usage: <code>u:&lt;username or user ID&gt;</code>
g	Group name or group ID, as they appear in <code>/etc/group</code> , of a specific group. Usage: <code>g:&lt;group name or group ID&gt;</code>
r	Name of a specific role. Usage: <code>r:&lt;role name&gt;</code> .
p	Public. Specifies that this operation is available to the public without restriction. Cannot be combined with any other operator. API request or CLI command to save such settings will return an error.
!	Negation operator. Usage: <code>!&lt;operator&gt;</code> .
&	AND operation.
	OR operation
()	Delimiters for subexpressions.
""	The empty string indicates that no user has the specified permission.

An example definition is `u:1001 | r:engineering`, which restricts access to the user with ID 1001 or to any user with the role `engineering`.

In this next example, members of the group `admin` are given access, and so are members of the group `qa`:

```
g:admin | g:qa
```

Another example is to have a list of groups to which you want to give read permissions:

To grant the read permission, you construct the following boolean expression:

```
u:cfkane | (g:admin & !g:c13) | (g:qa & (g:app2 | g:app3)) | (g:ba & g:dept_7a) | g:ds
```

This expression is made up of five subexpressions which are separated by OR operators:

- The first subexpression `u:cfkane` grants the read permission to the username `cfkane`.
- The subexpression `(g:admin & !g:c13)` grants the read permission to the admins for all clusters except cluster `c13`. The operator `g` is the group operator, the value `admin` is the name of the group of all admins. The `&` operator limits the number of administrators who have read permission because only those administrators who meet the additional condition will have it.

The condition `!g:c13` is a limiting condition. The operator `!` is the NOT operator. Combined with the group operator, this operator means that this group is excluded and does not receive the read permission.



**WARNING:** Be careful when using the NOT operator. You might exclude fewer people than you intended. For example, suppose that you do not want anyone in the group `group_a` to have access. You therefore define this ACE: `!g:group_a`. You might think that the data is now protected because members of `group_a` do not have access to it. However, you have not restricted access for anyone else except the members of `group_a`. The rest of the world can access the data. You should not define ACEs through exclusion by using the NOT operator. You should define them by inclusion and use the NOT operator to limit further the access of the groups or roles that you have included. In the subexpression `(g:admin & !g:c13)`, the NOT operator limits the number of members within the admin group who have access. The `admin` group is included, and all users who are also part of the `c13` group are excluded.

- The subexpression `(g:qa & (g:app2 | g:app3))` demonstrates use of a subexpression within a subexpression. The larger subexpression means that only members of group `qa` who are also members of group `app2` or `app3` have read access to the data. The smaller subexpression limits the number of people who have this permission in the `qa` group.

## Administering Buckets

---

Describes the operations you can perform related to buckets for the HPE Ezmeral Data Fabric.

Buckets are storage resources that store objects, which consist of data and its descriptive metadata.

Object-based storage is the preferred method of storing and efficiently managing gigantic volumes of data. Data are stored efficiently in a flat address space called as a storage pool and not as a tiered file structure. The address space is referenced by the metadata that holds the required information to retrieve the data. The metadata facilitates deep analysis of the usage and function of the data that is stored in the storage pool. The access protocol used in object storage architecture is TCP/IP and the communication medium is usually through REST APIs.

Objects can comprise disparate types of unstructured data such as audio files, video files, and images.

A user can store objects in the user's own account. Objects are stored inside containers called buckets. Every user can create buckets and set access policies or bucket policies to govern who can access the resources created by the user.

See [Administering Bucket Policies](#) on page 228 for information on managing bucket policies.

## Creating a Bucket

Create a bucket on a fabric.

### Prerequisites

- A fabric must be available for you to create a bucket.
- You must be a fabric user to create a bucket on the fabric.

### About this task

Buckets can be created on a fabric that exists on a public cloud provider.

When you create a bucket, you must ensure that the bucket name is

- globally unique for your fabric
- starts and ends with a lowercase letter or a number
- between 3 and 63 characters long
- contains characters in lowercase only

You can enable locking of objects that are stored on the bucket. Object locking is useful when you wish to prevent overwriting of objects for a specific time duration. When you lock an object, multiple versions of the object can be stored. Each object version is unalterable. Once object locking is enabled on an object, you cannot disable it for the object. You can specify the retention mode and retention period.

If you wish to create a versioned bucket to store multiple object versions with an object lock, you must select the **Enable Object Lock** check box. The **Object versioning** check box is auto-selected when you enable object locking.

You can create a bucket to store multiple object versions without enabling an object lock. In this case, you must select the **Object versioning** check box and leave the **Enable Object Lock** check box deselected.



**NOTE:** Currently, a bucket can be created by a non-SSO user only.

Follow the steps given below to create a bucket.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** on the Home page.
3. Click Create **Bucket** on the **Resources** card.
4. Enter the **Name**.
5. Select the fabric on which you wish to create the bucket.
6. Enter the **Account name** that owns the bucket.
7. Select the **Enable Object Lock** check box to enable object versioning and prevent deletion of objects. This is an optional step.
8. Select the **Object versioning** check box if you wish to enable object versioning, but do not wish to disable deletion of objects. This is an optional step.

## 9. Click **Create**.

### Results

The bucket is created on the fabric.

You can now upload objects to the bucket. You can create one or more folders on the bucket to store objects.

### Related `maprcli` Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `mc mb`

## Creating a Folder on a Bucket

Create a folder on a bucket to store objects.

### About this task

You can create folders on a bucket to segregate the objects to store on the bucket.

You can create subfolders inside folders.

Follow the steps given below to create a folder on a bucket.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** from the dropdown on the Home page.
3. Click the Table view icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the bucket to which you wish to upload objects (files).
4. Click the bucket name seen under **Resource Name**.
5. Click **Create Folder** on the **Objects** tab.



**NOTE:** If you are creating a sub-folder, you must navigate into the folder where you wish to create the sub-folder, and then click **Create Folder**.

6. Enter the **Folder name**.
7. Click **Create**.

### Results

The folder is created in the bucket. You can store objects in the folder.

## Uploading Objects to a Bucket

Upload one or more objects to a bucket.

### Prerequisites

- Your account must have the permission to upload object on the bucket.
- The bucket to which you wish to upload object must have enough space to store the object.

### About this task

You can upload one or more objects of size upto 1 GB to a bucket.

If object versioning is enabled on the bucket, you can store multiple files with the same name on the bucket.

Follow the steps given below to upload objects to a bucket.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** from the dropdown on the Home page.
3. Click the Table view icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the bucket to which you wish to upload objects (files).
4. Click the bucket name seen under **Resource Name**.
5. Click **Upload Objects** on the **Objects** tab.
6. Click **browse** to select one or more files to upload. Alternatively, drag and drop one or more files to **Upload files** area.
7. If you wish to store the files to upload as versions of the same object, enter the **Destination file name**. The files are stored as multiple versions of the same object in this case.
8. Click **Upload**.

### Results

The selected files are successfully uploaded to the object store and appended to the list of objects for the fabric seen on the Data Fabric UI.

If you have uploaded multiple files with the same name, that is, if you have uploaded multiple versions of an object, the multiple files are seen having the same file name with a version number in parentheses.

## Downloading an Object from a Bucket

Download an object from a bucket.

### Prerequisites

- You must have the permission to download an object from a bucket.

### About this task

You can download one or more objects that are stored on a bucket. You can download only one object at a time.

Follow the steps given below to download an object from the bucket.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric User** on the Home page.
3. Click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the bucket from which you wish to download the object(file).

4. Click the bucket name seen under **Resource Name**.
5. Click the ellipsis seen under **Actions** for the object row to download on the **Objects** tab.
6. Click the **Download** option.

### Results

The object is downloaded to the default download folder on your machine or to the folder that you select. The destination folder for the object depends on your web browser settings.

## Deleting an Object from a Bucket

Delete an object from a bucket.

### Prerequisites

- You must have the permission to delete an object from a bucket.

### About this task

You can delete versioned and unversioned objects from a bucket.

When you delete a versioned object, the object is merely marked for deletion and not actually deleted from the bucket.

When you delete an unversioned object, the object is permanently removed from the bucket.

Follow the steps given below to delete objects from a bucket.

### Procedure

1. Log on to the Data Fabric UI.
2. Click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the bucket from which you wish to delete objects (files).
3. Click the bucket name seen under **Resource Name**.
4. Click the ellipsis seen under **Actions** for the object row to delete on **Objects** tab.
5. Click the **Delete** option.
6. Click **Delete** on the message box that appears.

### Results

If the object is unversioned, the object is permanently removed from the bucket. If the object is versioned, it is marked for deletion.

### Related `maprcli` Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `mc rm`



**NOTE:** You might not be able to delete an object from GCP by using the minio client (`mc`). You would need to use the AWS CLI to do so. This is a known minio client issue. For more information about the issue, see <https://issuetracker.google.com/issues/162653700>.

## Deleting a Folder from a Bucket

Delete a folder from bucket.

### Prerequisites

- The folder to delete must not contain objects.
- You must have the permission to delete a folder from a bucket.

### About this task

You can delete a folder from a bucket when the folder does not contain any objects.

Follow the steps given below to delete a folder from a bucket.

### Procedure

1. Log on to the Data Fabric UI.
2. Click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the bucket from which you wish to delete a folder
3. Click the bucket name seen under **Resource Name**.
4. Click the ellipsis seen under **Actions** for the folder row to delete on **Objects** tab.
5. Click the **Delete** option.
6. Click **Delete** on the message box that appears.

### Results

The folder is deleted from the bucket.

## Deleting a Bucket

Delete a bucket from fabric.

### Prerequisites

- The bucket to delete must be empty.
- You must have the permission to delete a bucket.

### About this task

You can delete a bucket from a fabric.

Follow the steps given below to delete objects from a bucket.

### Procedure

1. Log on to the Data Fabric UI.
2. Click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the bucket that you wish to delete.
3. Locate the row for the bucket to delete in the Resources seen under the fabric.
4. Click the ellipsis seen under the **Actions** column for the bucket row.

5. Click the **Delete** option.
6. Click **Delete** on the message box that appears.

### Results

The bucket is deleted from the fabric.

### Related `maprcli` Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `mc rb`

## Administering Tables

---

Describes the operations you can perform related to tables for HPE Ezmeral Data Fabric.

Tables on HPE Ezmeral Data Fabric provide a native implementation of the HBase Table for improved scalability, stability, and speed on the Data Fabric platform.

With HPE Ezmeral Data Fabric tables, you can:

- Create and access tables in the Data Fabric UI.
- Access tables with a global path.
- List and filter tables.
- Create and manage access controls through the Data Fabric UI.
- Enable read optimization to improve the speed of large read workloads.

## Managing Tables

The topics in this section describe managing tables.

### Creating a Table

This topic describes creating a table.

### Prerequisites

- A fabric must be available for you to create a table.
- You must have the permissions to create a table on the fabric.

### About this task

Follow the steps given below to create a table.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** on the **Home** page.
3. Click **Table** on the **Resources** card.  
The **Create Table** side drawer opens.
4. Enter the table **Name**.



5. The table **Path** field is automatically filled. However, you can edit the **Path** if desired.  
The table is automatically assigned a fabric-specific full path, which you can use to access the table from the global namespace.
6. Select the name of the **Fabric** on which you want to create the table from the dropdown menu.
7. Enter a **Column family name** for your table. During the table creation process, you can create only one column family.
8. Click **Create**.

### Results

The table is created on the fabric.

After successful table creation, you can add additional column families to the table, as described in [Creating a Column Family](#) on page 153.

### Deleting a Table

This topic describes deleting a table.


### Prerequisites

- You must have the permissions to delete tables on the fabric.

### About this task

Follow the steps given below to delete a table.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** on the **Home** page.
3. Click the **Table View** icon on the **Resources** card.
4. Click the ellipsis (  ) in the **Action** column of the table you want to delete.
5. Click **Delete**.

### Results

The table is deleted from the fabric.

## Managing Column Families and Columns

The topics in this section describe managing column families and columns.

### Creating a Column Family

This topic describes creating a column family.

### Prerequisites

- A table must be available for you to create a column family.
- You must have the permissions to create column families on the table.

### About this task

Follow the steps given below to create a column family on a table.

## Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** on the **Home** page.
3. Click the **Table View** icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the table on which you want to create a column family.
4. Click the table name seen under **Resource Name**.
5. Click **Add column family** on the **Column families** tab.  
The **Add column family** wizard opens on the **Column family details** screen.
6. Enter the **Column family name**.
7. Enter the **Max version** and **Min version** for your column family. This is the the maximum and minimum number of versions that are retained upon revision to the column family.
8. Click the toggle to enable **Read optimization**. Read optimization enables in-memory caching for read workloads, greatly improving the speed of read operation for large workloads.
9. **(Optional)** Create user access controls for your column family.  
The wizard automatically creates a default column family **Access control** entry. You can edit the default entry, or delete it with the action button (trash can).
  - a) Select the user **Type**.
  - b) Enter the user **Name**.
  - c) Select the permissions you want the user to have:
    - **Read**: The user has read permissions on the column family.
    - **Write**: The user has write permissions on the column family.
    - **Append**: The user has append permissions on the column family.
    - **Version**: The user can roll the column family back to a previous version.
  - d) **(Optional)** Click **Add** to create another user access control for the column family.
10. Click **Manage column permissions**.  
The **Add column family** wizard opens the **Manage column permissions** window. From this window, you can create columns and column-specific access controls.
11. **(Optional)** To create a column in the column family, enter the column name.
12. **(Optional)** Create user access controls for your column.  
The wizard automatically creates a default column **Access control** entry. You can edit the default entry, or delete it with the action button (trash can).
  - a) Select the user **Type**.
  - b) Enter the user **Name**.
  - c) Select the permissions you want the user to have:
    - **Read**: The user can read the entries in the column.
    - **Write**: The user can alter the existing column entries.

- **Append:** The user can add new records to the column.

d) **(Optional)** Click **Add** to create another user column in the column family.

13. Click **Apply**.

## Results

The column family is created on the table with the defined access controls and columns.

## Configuring Column Family Permissions

This topic describes configuring column family permissions for a table.



### Prerequisites

- A table with a column family must be available for you to configure column family permissions.
- You must have the permissions to configure column families on the table.

### About this task

Follow the steps given below to create or edit column family permissions on a table.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** on the **Home** page.
3. Click the **Table View** icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the table on which you want to create or edit column family permissions.
4. Click the table name seen under **Resource Name**.  
The **Table details** view opens.
5. Select the **Column families** tab. On this tab, you can view all column families created on the table.
6. Perform one of the following:
  - Click the ellipsis (  ) in the **Action** column of the column family for which you want to create or edit column family permissions.
  - Click the name of the column family for which you want to create or edit column family permissions.  
The **Column family details** view opens. Then, click the ellipsis (  ).
7. Click **Edit**.  
The column family wizard opens.
8. Click **Manage column permissions**.  
The column family wizard opens the **Manage column permissions** window. From this window, you can create columns and column-specific user access controls.
9. Enter the column name.
10. Define user access controls for your column.  
The wizard automatically creates a default column **Access control** entry. You can edit the default entry, or delete it with the action button (trash can).
  - a) Select the user **Type**.

- b) Enter the user **Name**.
- c) Select the permissions you want the user to have:
  - **Read**: The user has read permission on the column.
  - **Write**: The user has write permissions on the column.
  - **Append**: The user has append permissions on the column.
- d) **(Optional)** Click **Add** to create another user access control for the column.

11. **(Optional)** Click **Add column permissions** to define user access controls for another column.

12. Click **Apply**.

### Results

The permissions are applied to the column family with the defined user access controls on the columns.

### Deleting a Column Family

This topic describes deleting a column family.


### Prerequisites

- You must have the permissions to delete column families from the table.

### About this task

Follow the steps given below to delete a column family from a table.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** on the **Home** page.
3. Click the **Table View** icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the table on which you want to delete a column family.
4. Click the table name seen under **Resource Name**.
5. Select the **Column families** tab.
6. Click the ellipsis (  ) in the **Action** column of the column family you want to delete.
7. Click **Delete**.

### Results

The column family is deleted from the table.

## Viewing Table Information

The topics in this section describe viewing table information.

### Viewing the List of Tables

View the list of tables created on a fabric.

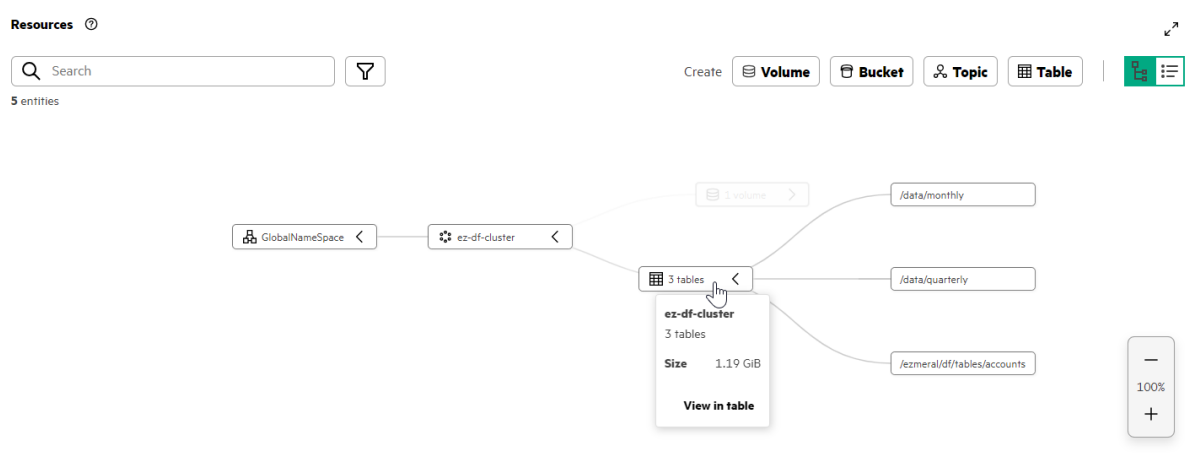
### About this task

Use the following steps to view tables created on a fabric.

### Procedure

1. Log on to the Data Fabric UI.
2. Under the default **Fabric user**, click the **Graph View** icon on the **Resources** card.

The **Graph View** opens:



3. To open the tabular list of tables, hover over the **tables** card in **Graph View** and click **View in table**. Alternatively, under the default **Fabric user experience**, click the **Table View** icon on the **Resources** card.

The **Table View** opens:

Resource Name	Type	Provider	Size	Owner	Action
ez-df-cluster	Fabric	OnPrem	0% 1 GiB of 267 GiB	--	⋮
/data/monthly	Table	--	1.08 GiB	admin	⋮
/data/quarterly	Table	--	0.11 GiB	admin	⋮
/ezmera...s/accounts	Table	--	0.00 GiB	admin	⋮
data	Volume	--	1.17 GiB	admin	⋮

4. To view information about a specific table, click the name of the table in **Table View**.

### Viewing Column Families

View column families created on a table.

### About this task

You can view column families from the Data Fabric UI.

Use the following steps to view column families.

### Procedure

1. Log on to the Data Fabric UI.

2. Under the default **Fabric user**, click the **Table View** icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the table.
3. Click the name of the table for which you want to view column families.  
The table **Overview** screen opens.
4. Click the **Column families** tab.

### Results

You can now view the list of column families.

## Managing Table Replication

The topics in this section describe managing table replication.

You can use table replication to improve data availability and load balancing.

- **Availability:** Replicate tables between different fabrics within your cluster.
- **Load balancing:** Use table replicas to reduce the load on the primary table. For example, you can run data analysis tasks on a replica of the primary table rather than the primary table itself.

### Adding a Table Replica

This topic describes adding a table replica.

#### Prerequisites

- A table must be available for you to create a replica.
- You must have the permissions to create table replicas.

#### About this task

Follow the steps given below to create a table replica.

#### Procedure

1. Log on to the Data Fabric UI.
2. Click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the table for which you want to create a replica.
3. Click the table name seen under **Resource Name**.
4. Click **Add replica** on the **Replication** tab.
5. The wizard automatically fills the name of the **Source fabric** and **Path to source table**.
6. Select the **Destination fabric** from the dropdown menu.
7. Enter the **Replica path**.
8. **(Optional):** Click the toggle to enable **Advanced options**:
  - **Throttle:** Enable throttling to limit the speed of connection. Use this option to minimize the impact on the primary table, especially when under heavy load.
  - **Synchronous:** Enable synchronous replication.
  - **Encrypt on wire:** Enable this option to encrypt the replicated data during transfer.

**9. Click **Column families**.**

The **Add replica** wizard opens the **Column families** window.

**10. Select the column families you want to replicate.**

You can either:

- **Replicate all column families.**
- **Replicate specific column families:**
  - a. Click the toggle for the column families you want to replicate.
  - b. For each column family, you can either replicate **All columns** in the column family or you can **Assign columns** to replicate.
  - c. If you select **Assign columns**, enter the column name in the **Column name** field. Click **Add** to add another column to replicate.

**11. Click **Add**.****Results**

The table replica is created in the specified location.

**Viewing Table Replicas**

View the list of created table replicas.

**Prerequisites**

You must have permission to view table replicas.

**About this task**

You can view table replicas from the Data Fabric UI.

Use the following steps to view table replicas.

**Procedure**

1. Log on to the Data Fabric UI.
2. Under the default **Fabric user experience**, click the **Table View** icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the table replica.
3. Click the name of the table for which you want to view replicas  
The table **Overview** screen opens.
4. Click the **Replication** tab.

**Results**

You can now view the list of table replicas.

**Administering Access Controls for Tables**

This topic describes administering access controls for tables.

**Prerequisites**

- To add or edit user access controls, you must have Admin permissions on the table.

**About this task**

Follow the steps given below to create user access controls for a table.

**Procedure**

1. Log on to the Data Fabric UI.
2. Under the default **Fabric user**, click the **Table View** icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the table for which you want to manage access controls.
3. Click the name of the table for which you want to manage access controls. The table **Overview** screen opens.
4. Click the **Settings** tab.
5. Click the **Access Control** button. The **Access Control** wizard opens.
6. Select the user type for the access control.
7. Enter the user name.
8. Select the permissions you want the user to have:
  - **Admin (Access control)**: The user can create, edit, and delete access controls for the table.
  - **Create/rename column family**
  - **Delete column family**
9. **(Optional)** Click **Add** to create another access control for the table.
10. Click **Save**.

## Administering Topics

---

Administer topics for Apache Kafka Wire Protocol with HPE Ezmeral Data Fabric.

Data Fabric supports creation of topics for Apache Kafka Wire Protocol via the Data Fabric UI.

A topic is given a unique name that is used to categorize and organize messages based on common properties.

One or more producers publish messages to topics. Individual consumers subscribe to the topics of their choice to consume the messages that are published to such topics.

For example, a producer could be publishing weather-related data such as daily rainfall to a topic called dailyrain. A consumer could be subscribing to the topic containing the daily rainfall.

### Creating a Topic

Create a topic for Apache Kafka Wire Protocol.

**Prerequisites**

- Topic with the name must not already exist on the fabric.
- You must be a fabric user to create a topic on the fabric.



**About this task**

One or more topics for Apache Kafka Wire Protocol can be created on a fabric, via the Data Fabric UI.

Data is stored in topics in the form of messages for the retention period specified while creating the topic. The retention period can be specified in seconds, minutes, hours, or days. The default unit for retention period is days.

Data can be stored in uncompressed and compressed formats in topics. By default, data on topics is uncompressed.

Data compression helps reduce network bandwidth usage and saves disk space. Data compression results in higher CPU utilization.

Data Fabric supports the following algorithms to compress data in a topic.

- LZ4 - LZ4 is an extremely fast /high-speed lossless compression algorithm.
- LZF – java-based compression algorithm that is optimized for speed with modest data compression
- ZLib – Zlib is an open-source mechanism to provide for lossless compression.

If you wish to store data in a compressed format on the topic, you can select the compression algorithm based on your environment and requirement.

Follow the steps given below to create a topic.

**Procedure**

1. Log on to the Data Fabric UI.
2. Select **Fabric user** option from the dropdown on the Home page.
3. Click **Topic** seen next to Create on the **Resources** card.
4. Enter the topic name.
5. Select the fabric on which you wish to create the topic.
6. Enter the number of partitions for the topic
7. Enter time to live for the topic. This is the duration for which messages are retained within the topic.
8. Choose the compression type.
9. Click Create.

**Results**

The topic is created on the specified fabric. A message regarding the topic creation is displayed on the Data Fabric UI.

**Related maprcli Commands**

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `stream topic create`

**Editing a Topic**

Edit a topic for Apache Kafka Wire Protocol.

### Prerequisites

- Topic to edit must exist.
- You must have permission to edit a topic.

### About this task

You can edit a topic for Apache Kafka Wire Protocol via the Data Fabric UI.

The following fields related to a topic are editable.

- Number of partitions
- Time to live
- Compression scheme

Follow the steps given below to edit a topic.

### Procedure

1. Log on to the Data Fabric UI.
2. Under the default **Fabric user**, click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the topic to edit.
3. Click the topic name seen under **Resource Name**.
4. Edit the topic as required.
5. Click **Save**.

### Results

The changes made to the topic are saved. A message regarding the successful modification is displayed on the Data Fabric UI.

### Related `maprcli` Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `stream topic edit`

## Deleting a Topic

Delete a topic for Apache Kafka Wire Protocol.

### Prerequisites

- The topic to delete must exist.
- You must be a fabric user to delete a topic.

### About this task

You can delete a topic associated with a fabric.

Follow the steps given below to delete a topic.

**Procedure**

1. Log on to the Data Fabric UI.
2. Select **Fabric user** option from the dropdown on the Home page.
3. Click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the topic to delete.
4. Click the topic name seen under **Resource Name**.
5. Locate the topic to delete.
6. Click the ellipsis under **Actions** for the topic row in the tabular list of resources.
7. Click the **Delete** menu option.
8. Click **Delete** to confirm topic deletion.

**Results**

The topic is permanently deleted, along with the data stored on the topic. A message regarding the topic deletion is displayed on the Data Fabric UI.

**Related `maprcli` Commands**

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `stream topic delete`

**Viewing or Downloading Topic Connection Properties**

View and download connection properties related a topic for Apache Kafka Wire Protocol.

**Prerequisites**

You must be a fabric user or a fabric manager to view and download connection properties for topic.

**About this task**

You can view and/or download topic connection properties from the Data Fabric UI.

Use the following steps to view or download connection properties.

**Procedure**

1. Log on to the Data Fabric UI.
2. Under the default **Fabric user**, click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the topic.
3. Click the **View Connection Properties** in the **Endpoint** column for the topic, to view the connection properties.
4. If you want to download the connection properties, click **Download** on the **Connection Properties** dialog box.

**Results**

The connection properties file is downloaded to the default downloads folder on the local machine.

## Administering Volumes

---

Administer volumes on HPE Ezmeral Data Fabric.

HPE Ezmeral Data Fabric provides volumes as a way to organize data and manage fabric performance.

A volume is a logical storage unit that allows you to apply policies to a set of files, directories, and sub-volumes. You can use volumes to enforce disk usage limits, establish ownership and accountability, and measure the cost generated by different projects or departments. For example, you could create a volume for each user, department, or project.

A volume consumes storage space only when data is written to the volume.

### Volume Types

Volumes can be of the following types.

- **Standard volume:** A initial or original volume that contains data.
- **Mirror volume:** A replica of a standard volume that replicates the data present on a standard volume.

### Volume Data Backup

Following are the ways in which data on volumes can be backed up for restoration purpose.

- **Creating mirrors or mirror volumes:** Data in a volume is replicated when you create mirror volumes. A volume that is mirrored is called a standard volume. The mirrored version of a volume is called a mirror volume. You can create mirrors of mirror volumes to replicate the data present on mirror volumes. Such volumes are called cascading mirrors.
- **Creating volume snapshots:** The state of data in a volume is recorded at the point-in-time of the volume snapshot creation. A snapshot is a read-only image of a volume at a specific point in time. Snapshots are useful any time you need to be able to roll back to a known good data set at a specific point in time. For example, before performing a risky operation on a volume, you can create a snapshot to enable rollback capability for the entire volume.

A snapshot takes no time to create, and initially uses no disk space, because it stores only the incremental changes needed to roll the volume back to the state at the time the snapshot was created. The storage used by a volume's snapshots does not count against the volume's quota.

Security policies can be applied to volumes to control access to the volume data. See [Administering Security Policies](#) on page 206 for more information.

## Creating a Standard Volume

Procedure to create standard volume.

### Prerequisites

- The fabric for which you wish to create the volume must have been created on the cloud or on-premises, or in an air-gapped environment.
- You must have the permission to create a volume.
- If you wish to enable data at rest encryption for the volume, the fabric must have data-at-rest encryption enabled.
- If you wish to associate a security policy with the volume, the security policy must have been created by the fabric. manager.

**About this task**

The task describes how to add standard volume via the Data Fabric UI.

A standard volumes is a volume where data is originally written. A standard volume is read-write, default. It can marked as read-only, if required.

You have the option to create a new storage policy or create a new remote target to cold tiering of the volume. You can create a new storage policy or a remote target if you have not already created one or wish to create a new storage policy or remote target for the cold data that would be stored on the volume to create.

Follow the steps given below to create a standard volume.

**Procedure**

1. Log on to the Data Fabric UI.
2. Under the default **Fabric user**, click Create **Volume** on the **Resources** card.
3. Enter the volume name.
4. Select the fabric on which you wish to create the volume.
5. Select the Standard option for **Type**.
6. Specify the **Mount Path** for the volume.
7. Enter the **Volume Limit**. The value specified for the volume limit must be less than or equal to the fabric capacity.
8. Select a security policy for the volume. This is an optional step.
9. Turn on the **Data at rest encryption** toggle to enable encryption for data at rest. This is an optional step.
10. Turn on the Data tiering toggle to enable data tiering to be able to offload data to cold tier. This is an optional step.
  - a) If you enable data tiering, select tiering type. On selecting tiering type as **Remote archiving(Cold)**, you can offload data to a cold tier.
  - b) If you have selected **Remote archiving (cold)**, select the **Storage policy** and the **Remote Target**. You can create a new storage policy and/or a remote target by clicking Create new under the respective dropdowns.
11. Click **Create**.

**Results**

The volume is successfully created on the selected fabric and the details are displayed in the list of resources under the respective fabric on the **Resources** card.

If you have turned on data tiering, tiering is enabled on the volume and you can configure tiering options on the **Settings** tab for the volume.

**Related maprcli Commands**

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `volume create`

## Creating a Mirror Volume

Procedure to create mirror volume.

### Prerequisites

- The fabric for which you wish to create the volume must have been created on the cloud or on-premises or in an air-gapped environment.
- The standard volume that you wish to mirror must have been created, before you can create a mirror volume.
- You must have the permission and valid activation key to create a mirror volume on the fabric.

### About this task

A mirror is a replica of a standard volume and consumes space on the disk. You can create a mirror of an existing volume by using the Data Fabric UI.

Mirror volumes can be used for data recovery to recover data on corrupted or lost volume data.

A mirror volume can be created on-premise or in the cloud, based on your fabric setup. Mirror volumes can be created on a fabric other than the fabric to which the source volume belongs.

Mirror volumes created on-premise are called local mirrors.

Mirror volumes can be local or remote.

Remote mirrors are used for disaster recovery, while local mirrors are used for load balancing.

Mirror volumes that are created from mirror volumes are called cascading mirrors.

Follow the steps given below to create a mirror volume.

### Procedure

1. Log on to the Data Fabric UI.
2. Under the default **Fabric user**, click Create **Volume** on the **Resources** card. Alternatively, you can click the ellipsis under **Actions** for the volume to mirror, on the Resources card, and click the **Create mirror** option to create a mirror for the volume.
3. Enter the volume name.
4. Select the fabric on which you wish to create the volume from **Source Fabric**.
5. Select the Mirror option for **Type**.
6. Select the **Source Cluster** that represents the cluster/fabric that contains the volume to replicate.
7. Select the **Source Volume** to replicate. This is non-editable if you have already chosen the volume to mirror.
8. Specify the **Mount Path** for the volume.
9. Enter the **Volume Limit**. The value specified for the volume limit must be less than or equal to the fabric capacity.
- 10.
11. Select a security policy for the volume. This is an optional step.

12. Turn on the **Data at rest encryption** toggle to enable encryption for data at rest. This is an optional step.
13. Turn on the Data tiering toggle to enable data tiering to be able to offload data to cold tier. This is an optional step.
  - a) If you enable data tiering, select tiering type. On selecting tiering type as **Remote archiving(Cold)**, you can offload data to a cold tier.
  - b) If you have selected **Remote archiving (cold)**, select the **Storage policy** and the **Remote Target**. You can create a new storage policy and/or a remote target by clicking Create new under the respective dropdowns.
14. Click **Create**.

## Results

The mirror volume is created.



**NOTE:** Alternatively, you can click the volume to mirror in Resources list, and click **Actions > Create mirror**

## Related maprcli Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `volume create`

## Converting Standard Volume to Mirror Volume

### Prerequisites

- You must be a fabric user to convert a standard volume to a mirror volume.
- The standard volume must have an associated mirror volume that can serve as the source volume.

### About this task

You can convert a standard volume to a mirror volume and set it up to mirror one of its associated mirror volumes.



**NOTE:** The standard volume, when converted, can only be a mirror of one of its associated mirror volumes.

Follow the steps given below to convert a standard volume to a mirror volume.

### Procedure

1. Log on to the Data Fabric UI.
2. On the **Resources** card, click the standard volume name that you wish to convert to a mirror volume.
3. Click the Actions menu seen on the top right of the tabs displaying the volume details.
4. Click **Make Mirror**.
5. Select the **Source fabric** that has the mirror volume that the converted volume would mirror.

6. Select the **Source volume** that the converted volume will mirror.
7. Click **Save**.

### Results

The source fabric is converted into a read-only mirror of the selected source volume on the selected source fabric with the existing name.

You can associate a mirroring schedule with this volume to ensure that data on the volume is in sync with the source volume.

## Editing a Volume

Edit accountable entity, volume access for accountable entity and volume hard quota.

### Prerequisites

- You must have the permission to edit the volume.

### About this task

You can edit details such as hard quota, read/write access and accountable entity, when you edit a volume via the Data Fabric UI.

Follow the steps given below to edit volume properties.

### Procedure

1. Log on to the Data Fabric UI.
2. Under the default **Fabric user experience**, click the Table view icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume.
3. Click the volume name seen under **Resource Name**.
4. Navigate to **Settings** tab for the volume.
5. Click the pencil/Edit icon next to **Properties** to edit the accountable entity and/or volume access as read or read/write. Click Save. This is an optional step.
6. Click the pencil/Edit icon next to **Quota** to edit the hard quota for the volume. This is an optional step.
7. Click **Save**.

### Results

The volume is edited successfully.

### Related tasks

[Configuring Data Access Control for Volume](#) on page 169

[Configuring Volume Administration Settings](#) on page 170

Configure volume access control for various user types.

### Setting a Volume Quota

Set space quota for volume.

### Prerequisites

You must have be a fabric user to configure or set volume quota.



**About this task**

You can set hard quota for volume via the Data Fabric UI.



**NOTE:** It is recommended to set advisory quota for a volume. See [Setting advisory quota via CLI](#) to set advisory quota. An alarm is set off when advisory quota is reached or exceeded.

Use the following steps to set the volume quota.

**Procedure**

1. Log on to the Data Fabric UI.
2. Select the **Fabric user** on the Home page.
3. Click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume to set the quota for.
4. Click the volume name seen under **Resource Name**.
5. Navigate to the **Settings** tab.
6. Click the Edit icon seen next to **Quota**.
7. Enter the value and select the unit of measurement for the quota.
8. Click **Save**.

**Results**

The hard quota for the volume is set.

Data on the volume cannot occupy space higher than the quota. If the data on the volume exceeds the hard quota, an alarm is raised.

**Configuring Data Access Control for Volume****Prerequisites**

The following prerequisites must be met before attempting to configure data access control for a volume.

- The user or group, for which you wish to assign permissions, must exist for the fabric.
- You must have be a fabric user or fabric manager to configure data access for the volume.

**About this task**

Configure or control access to data on a volume.

There are three user types that can be assigned access to data on a volume. The user types are as follows.

- **Public** - Public refers to all users.
- **User** - A user is an individual user that uses the Data Fabric UI.
- **Group** - Groups are collection of users that are categorized based on a commonality such as department or location.

You can assign read or write permissions to public or to specific users and/or groups. By default, read and write permissions are assigned to all users.

When a read or write permission is assigned to public, it implies that all users and groups are able to read data from and write data to the volume in question.



**NOTE:** A write permission assigned to a user type implies that the respective user type has both read and write permissions on the volume.

If you wish to assign read or write permission to specific users and/or groups, you must first remove any permission assigned to the public user type. You cannot assign read or write permissions to specific users and/or groups when the respective permissions are assigned to public.

To configure data access control for a volume, follow these steps:

### Procedure

1. Log on to the Data Fabric UI.
2. Select the **Fabric user** option on the Home page.
3. Click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume for which you wish to set data access control.
4. Click the volume name seen under **Resource Name**.
5. Click the **Settings** tab.
6. Click the Edit icon seen next to **Data Access Control**.
7. Click **Add** to add permissions for a user or a group.
8. Select user or group from the **Type** dropdown.
9. Enter the name of the user or group to which you wish to assign the permission or permissions.
10. Select read and/or write checkbox depending on the permission you wish to assign to the user or the group.
11. Repeat the steps related to adding user/group and assigning permissions to them, as necessary.
12. Click **Save**.

### Results

Permissions are assigned to the respective users and/or groups.

### Configuring Volume Administration Settings

Configure volume access control for various user types.

### Prerequisites

The following prerequisites must be met before you can configure volume administration settings for users and groups.

- You must have be a fabric manager to perform the operation.
- The users and groups must have been created before you can assign various volume-related permissions.

### About this task

You can assign various permissions related to a volume to specific users and/or groups.

The following permissions with respect to a volume can be assigned to one or more users and/or groups.

- Edit permission

- Admin permission to manage access control to volume and data on the volume
- Restore and mirror permission
- Delete permission
- Full control, which means, all the aforementioned permissions.

Follow the steps given below to configure volume administration settings.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** option from the dropdown on the Home page.
3. Click the Table view icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume.
4. Click the volume name seen under **Resource Name**.
5. Click the **Settings** tab.
6. Click the edit icon seen next to **Volume Admin Control**.
7. Click **Add** to add permissions for a user or a group.
8. Select user or group from the **Type** dropdown.
9. Enter the name of the user or group to which you wish to assign the permission or permissions.
10. Select the checkbox or checkboxes for the permission(s) that you wish to assign to the user or the group.
11. Repeat the steps related to adding user/group and assigning permissions to them, as necessary.
12. Click **Save**.

### Results

Permissions are assigned to the respective users and/or groups.

## Renaming a Volume

Rename a volume.

### Prerequisites

- You must have the permission to rename a volume.

### About this task

You can rename a volume associated with a fabric, via the Data Fabric UI. You can rename only one volume at a time.

Follow the steps given below to rename a volume.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** option from the dropdown on the Home page.

3. Click the Table view icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume to rename.
4. Click the volume name seen under **Resource Name**.
5. Click the ellipsis under **Actions** for the volume.
6. Click the **Rename** option.
7. Enter the new name for the volume and click **Save**.

### Results

The volume is renamed successfully.

### Related maprcli Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `volume rename`

## Viewing Volume Endpoint Info

View volume endpoint information.

### Prerequisites

You must be a fabric user to perform this operation.

### About this task

You can view volume endpoint information from the Data Fabric UI.

This can be used in scripts to make API calls to access the volume.

Follow the steps given below to view volume endpoint information.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** option from the dropdown on the Home page.
3. Click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume.
4. Click the ellipsis under **Actions** for the required volume.
5. Click the **View endpoint** option.

### Results

The volume endpoint information is displayed.

## Viewing Object Endpoint Info to Remotely Access Files as Objects

View endpoint information for files in a volume to be able to access the files as objects when accessed by S3 client.

### Prerequisites

You must be a fabric user to perform this operation.

### About this task

A volume contains one or more files. The files in a volume can be accessed as objects by an S3 client, via the endpoints provided by Data Fabric. You can view the object endpoints from the Data Fabric UI.

The object endpoints for files can be used to make API calls in scripts for S3 clients to access files in the volume.

Follow the steps given below to view object endpoint information for files in a volume.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** option from the dropdown on the Home page.
3. Click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume.
4. Click the ellipsis under **Actions** for the required volume.
5. Click the **View endpoint** option.

### Results

You are able to view the volume endpoint and object endpoint information corresponding to the files on the volume.

## Downloading Volume Endpoint Information

Download JSON file containing endpoint information for the selected volume endpoint information.

### Prerequisites

You must be a fabric user to perform this operation.

### About this task

You can download the endpoint information for a selected volume. The downloaded volume endpoint information is available as a JSON file.



**NOTE:** You can view object endpoints for files on the Data Fabric UI. Only volume endpoint information is downloadable.

Follow the steps given below to download information related to volume endpoint.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** option from the dropdown on the Home page.
3. Click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume.
4. Click the ellipsis under **Actions** for the required volume.
5. Click the **View endpoint** option.
6. Click **Download** on the Endpoints dialog box.

## Results

A JSON file containing the volume endpoint information for the selected volume is downloaded to your local downloads folder.

## Deleting a Volume

Delete a single volume.

### Prerequisites

- You must be a fabric user to perform this operation.

### About this task

You can delete a volume associated with a fabric, via the Data Fabric UI. Once the volume is deleted, the data on the volume is lost unless you have a backup of the data on the volume.

Follow the steps below to delete a volume.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** option from the dropdown on the Home page.
3. Click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume to delete.
4. Click the volume name seen under **Resource Name**.
5. Click the ellipsis under **Actions** for the volume.
6. Click the **Delete** option.
7. Click **Delete** to confirm volume deletion.

## Results

The volume is permanently deleted, along with the data stored on the volume.

### Related maprcli Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `volume delete`

## Administering Volume Snapshots

Snapshot overview and administering snapshots.

### Snapshot Basics

A snapshot is a static or read-only view of a volume that represents the state of the volume at the point in time of the snapshot creation. Because a snapshot is not a replica of a volume, it does not occupy much space on the volume.

A snapshot takes no time to create, and initially uses no disk space, because it stores only the incremental changes needed to roll the volume back to the state at the time the snapshot was created. The storage used by a volume's snapshots does not count against the volume's quota.

A snapshot can be used to restore volume data to the state the volume data at the time of the snapshot creation.

You can use a snapshot for the following purposes.

- Snapshots enable you to roll back to a known good data set and recover data in case of data corruption or accidental deletions, without the help of storage administrators. For example, before performing a risky operation on a volume, you can create a snapshot to enable rollback capability for the entire volume.
- Create static data sets for querying and auditing.

Snapshots are stored in the `.snapshot` directory on the volume mount path.

You can access snapshots via NFS or the Hadoop shell.

You can create a snapshot manually, or automate the process with a schedule. If you wish to schedule the creation of snapshots, you must assign a predefined schedule.

### Related tasks

[Creating a Volume Snapshot](#) on page 176

Create volume snapshot manually via Data Fabric UI.

[Preserving a Volume Snapshot](#) on page 178

Preserve a volume snapshot.

[Restoring a Volume from Volume Snapshot](#) on page 179

Describes how to restore a volume from a volume snapshot.

[Deleting a Volume Snapshot](#) on page 179

Delete a volume snapshot.

### Schedules for Volume Snapshots

Describes schedules for snapshots

#### About Schedule

A schedule to capture snapshots of volume data can be created and assigned to volumes.

Data Fabric provides predefined schedules that can be applied to volumes.

Predefined schedules are classified into three categories depending on the type of data in the volume you wish to back up with a snapshot.

You can select a pre-defined schedule for a volume, depending on the type of data that the volume contains.

If the volume data needs to be backed up very frequently as a snapshot, select the critical data schedule.

If it suffices to back up volume data less frequently, select the normal data schedule.

**Table**

Pre-defined Schedule	Frequency and Retention Period	Comments
Critical Data	<ul style="list-style-type: none"> <li>• Hourly - Retained for 24 hours</li> <li>• Daily at 12:00 AM - Retained for 7 days</li> <li>• Weekly every Sunday at 12:00 AM - Retained for 12 weeks</li> </ul>	<p>Use for volumes with data that might be changing constantly and/or needs to be frequently backed up.</p> <p>If you wish to preserve the snapshot beyond the default retention period per schedule, you can preserve the snapshot.</p>

Table (Continued)

Pre-defined Schedule	Frequency and Retention Period	Comments
Important Data	<ul style="list-style-type: none"> <li>Daily at 6:00 AM - Retained for 24 hours</li> <li>Daily at 12:00 PM - Retained for 24 hours</li> <li>Daily at 6:00 PM - Retained for 24 hours</li> <li>Daily at 12:00 AM - Retained for 7 Days</li> <li>Weekly every Sunday at 12:00 AM - Retained for 4 weeks</li> <li>Monthly every first day of the month at 12:00 AM - Retained for 2 months</li> </ul>	Use for volumes containing data that needs to be backed up frequently during the day and week. If you wish to preserve the snapshot beyond the default retention period per schedule, you can preserve the snapshot.
Normal Data	<ul style="list-style-type: none"> <li>Daily at 12:00 AM - Retained for 7 days</li> <li>weekly every Sunday 12:00 AM - Retained for 4 weeks</li> <li>Monthly every first day of the month at 12:00 AM - Retained for 2 months</li> </ul>	Use for volumes for volumes containing data that changes infrequently or does not need to be backed up frequently. If you wish to preserve the snapshot beyond the default retention period per schedule, you can preserve the snapshot.

When you specify a snapshot schedule on a mirror volume, it specifies how often to take a snapshot of the mirror volume. This snapshot schedule is distinct from the snapshot schedule for the standard volume.

A snapshot schedule for a promotable mirror volume has two purposes:

- The schedule specifies how often to take a snapshot of the mirror volume for the purpose of preserving the state of the mirror before a subsequent mirror operation. This way, if corrupt data is copied from the source volume's snapshot into the mirror volume, the mirror contents can be rolled back to the snapshot.
- If the promotable mirror volume is promoted to a read-write volume, the snapshot schedule specified for the mirror is used for the promoted read-write volume. Once a mirror volume is promoted to a read-write volume, the mirror schedule is disabled.

### Creating a Volume Snapshot

Create volume snapshot manually via Data Fabric UI.

### Prerequisites

- You must be a fabric user to create a volume snapshot.

### About this task

A snapshot is a read-only image of a volume that provides point-in-time recovery of the volume if data on the volume gets corrupted or is lost due to some reason. Snapshots store changes to the data present in the volume. A snapshot preserves access to historical data and facilitates data retrieval, when data is lost due to user errors and application errors.



You can create a snapshot manually, or automate the process with a schedule. Snapshots are stored in the `.snapshots` directory. You can always view snapshots from this directory.

The maximum number of snapshots that you can create for each volume is 4092. Exceeding this limit raises the snapshot failure alarm with an appropriate entry in the CLDB logs.

Follow the steps given below to create a snapshot manually via the Data Fabric UI.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** option from the dropdown on the Home page.
3. Click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume to snapshot.
4. Click the volume name seen under **Resource Name**.
5. Navigate to the **Snapshots** tab.
6. Click **Create** seen on the top right side of the tab.
7. Enter a name for the volume snapshot or retain the default name.
8. Click **Save**.

### Results

The volume snapshot is successfully created. The snapshot is added to the list of snapshots seen on the **Snapshots** tab.

### Related `maprcli` Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `volume snapshot create`

### Scheduling Volume Snapshots

Assign schedule to volume for creation of volume snapshots.

### Prerequisites

- You must be a fabric user to assign a schedule to volume to take a volume snapshot.
- The volume to assign the pre-defined schedule to must be a standard volume.

### About this task

You can assign a pre-defined schedule to a standard volume to take volume snapshots. See [Schedules for Volume Snapshots](#) on page 175 for details on pre-defined schedules.

Follow the steps given below to schedule volume snapshot.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** option from the dropdown on the Home page.

3. Click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume to snapshot.
4. Click the volume name seen under **Resource Name**.
5. Navigate to the **Settings** tab.
6. Under **Schedules**, click the Edit icon seen next to **Snapshots**.
7. Select a suitable schedule option for the volume.
8. Click **Select**.

### Results

The selected schedule is applied to the volume and snapshots are taken per schedule. The snapshots are stored in the `.snapshots` directory of the volume mount path.

### Preserving a Volume Snapshot

Preserve a volume snapshot.

### Prerequisites

- You must be a fabric user to preserve a volume snapshot.

### About this task

You can preserve a volume snapshot if you wish to store a volume snapshot beyond the retention period defined during its creation.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** option from the dropdown on the Home page.
3. Click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume whose snapshot is to be preserved.
4. Click the volume name seen under **Resource Name**.
5. Navigate to the **Snapshots** tab.
6. Select the checkbox for the snapshot to preserve.
7. Click the down arrow next to **Actions** seen on the top right side of the tab. Alternatively,
8. Click the **Preserve** menu option.
9. Click **Preserve** on the message box that appears.

### Results

The volume snapshot is preserved successfully. A message informing you about successful preservation of snapshot is displayed on the Data Fabric UI.

The volume snapshot can be accessed from the `.snapshot` directory on the volume mount path.

### Related maprccli Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `volume snapshot preserve`

### Restoring a Volume from Volume Snapshot

Describes how to restore a volume from a volume snapshot.

#### Prerequisites

- You must be a fabric user to restore a volume from the volume snapshot.

#### About this task

You can restore volume data to a specific point in time with volume snapshots.

Follow the steps given below to restore a volume from a volume snapshot.

#### Procedure

1. Log on to the Data Fabric UI.
2. Select the **Fabric user** option on the Home page.
3. Click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume whose snapshot is to be restored.
4. Click the volume name seen under **Resource Name**.
5. Navigate to the **Snapshots** tab.
6. Select the checkbox for the snapshot from which you wish to restore the volume.
7. Click the ellipsis seen next to the snapshot to restore from.
8. Click the **Restore** menu option.
9. Click **Restore** on the message box that appears.

#### Results

The volume is restored from the snapshot. A message informing you about successful restoration of volume is displayed on the Data Fabric UI.

The volume snapshot is preserved indefinitely in the `.snapshot` directory on the volume mount path.

#### Related maprcli Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `volume snapshot restore`

### Deleting a Volume Snapshot

Delete a volume snapshot.

#### Prerequisites

- You must be a fabric user to delete a volume snapshot.

## About this task

You can delete a volume snapshot manually, if you no longer wish to retain the volume snapshot.



**NOTE:** A volume snapshot is deleted automatically when the retention period for the volume snapshot ends.

## Procedure

1. Log on to the Data Fabric UI.
2. Select the **Fabric user** option on the Home page.
3. Click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume whose snapshot is to be preserved.
4. Click the volume name seen under **Resource Name**.
5. Navigate to the **Snapshots** tab.
6. Select the checkbox for the snapshot to delete.
7. Click the down arrow next to **Actions** seen on the top right side of the tab.
8. Click the **Delete** menu option.
9. Click **Delete** on the message box that appears.

## Results

The volume snapshot is deleted successfully. The volume snapshot is removed from the `.snapshot` folder. You are no longer able to access or use the volume snapshot. The data on the volume referred to by the snapshot remains intact, only the static view pointing to the volume data is deleted.

## Related `maprcli` Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `volume snapshot remove`

## Data Tiering

Conceptual information about data tiering.

Data that is active and frequently accessed is categorized as hot data.

Data that is rarely accessed is categorized as warm data or cold data.

Hot data, warm data, and cold data is identified based on the rules and policies set on by the administrator.

The storage used to store hot data is referred to as the hot-tier. The storage used to store warm data is referred to as an EC-tier, and the mechanism to store cold data is referred to as the cold tier.

Data starts off as hot data when it is first written to local storage on a fabric. Data can be termed as warm or cold based on the storage policies that are configured for the data present on Data Fabric.

Data stored on a fabric requires three times the amount of disk space of the regular volume on premium hardware due to replication (default being 3). After offloading to the cloud, the space used by data (including data in the namespace container) in the volume on the data fabric cluster is freed and only the metadata of the volume in the namespace container is 3-way replicated on the data fabric cluster.

Data can be set up to be automatically offloaded to a volume on a low-cost storage alternative, called a warm tier, on the data fabric cluster. Alternatively, data can be offloaded to a low-cost storage on a third party cloud object store, called a cold tier, like S3.

Data Fabric provides rule-based automated tiering functionality that allows you to seamlessly integrate with:

- Low-cost storage as an additional storage tier in the data fabric cluster for storing file data that is less frequently accessed (warm data) in erasure-coded volume.
- 3rd party cloud object storage as an additional storage tier in the data fabric cluster to store file data that is rarely accessed or archived (cold data).

In this way, valuable on-premise storage resources can be used for more active or hot file data and applications, while warm and/or cold file data can be retained at minimum cost for compliance, historical, or other business reasons. The data fabric provides consistent and simplified access to and management of the data.

Data, once offloaded, is purged on the the data fabric cluster to release the disk space. When you delete an entire file, part of a file, or a snapshot, corresponding objects are removed from the tier

When a client tries to read offloaded data, the data fabric processes the read request of the warm-tiered and cold-tiered standard and mirror volume data differently. Similarly, when a client writes to a tiered volume, the data fabric processes appends and overwrites differently.

To manage data offloading, you must have created storage policies. See [Administering Storage Policies](#) on page 184 to learn more about managing storage policies.

To offload data, you must create remote targets. See [Creating a Remote Target](#) on page 189 to add a new remote target.

You can schedule data offloading. See [<add link to schedules>](#) for further information on creating schedule.

### Schedules for Volume Data Tiering

Describes schedules for data tiering of volume data

#### About Schedule

A schedule to capture snapshots of volume data can be created and assigned to volumes.

Data Fabric provides predefined schedules that can be applied to volumes.

Predefined schedules are classified into three categories depending on the type of data in the volume you wish to back up with a snapshot.

You can select a pre-defined schedule for a volume, depending on the type of data that the volume contains.

If the volume data needs to be backed up very frequently as a snapshot, select the critical data schedule.

If it suffices to back up volume data less frequently, select the normal data schedule.

Table

Pre-defined Schedule	Frequency and Retention Period	Comments
Critical Data	<ul style="list-style-type: none"> <li>• Hourly - Retained for 24 hours</li> <li>• Daily at 12:00 AM - Retained for 7 days</li> <li>• Weekly every Sunday at 12:00 AM - Retained for 12 weeks</li> </ul>	<p>Use for volumes with data that might be changing constantly and/or needs to be frequently backed up.</p> <p>If you wish to preserve the snapshot beyond the default retention period per schedule, you can preserve the snapshot.</p>

Table (Continued)

Pre-defined Schedule	Frequency and Retention Period	Comments
Important Data	<ul style="list-style-type: none"> <li>Daily at 6:00 AM - Retained for 24 hours</li> <li>Daily at 12:00 PM - Retained for 24 hours</li> <li>Daily at 6:00 PM - Retained for 24 hours</li> <li>Daily at 12:00 AM - Retained for 7 Days</li> <li>Weekly every Sunday at 12:00 AM - Retained for 4 weeks</li> <li>Monthly every first day of the month at 12:00 AM - Retained for 2 months</li> </ul>	Use for volumes containing data that needs to be backed up frequently during the day and week. If you wish to preserve the snapshot beyond the default retention period per schedule, you can preserve the snapshot.
Normal Data	<ul style="list-style-type: none"> <li>Daily at 12:00 AM - Retained for 7 days</li> <li>weekly every Sunday 12:00 AM - Retained for 4 weeks</li> <li>Monthly every first day of the month at 12:00 AM - Retained for 2 months</li> </ul>	Use for volumes for volumes containing data that changes infrequently or does not need to be backed up frequently. If you wish to preserve the snapshot beyond the default retention period per schedule, you can preserve the snapshot.

When you specify a snapshot schedule on a mirror volume, it specifies how often to take a snapshot of the mirror volume. This snapshot schedule is distinct from the snapshot schedule for the standard volume.

A snapshot schedule for a promotable mirror volume has two purposes:

- The schedule specifies how often to take a snapshot of the mirror volume for the purpose of preserving the state of the mirror before a subsequent mirror operation. This way, if corrupt data is copied from the source volume's snapshot into the mirror volume, the mirror contents can be rolled back to the snapshot.
- If the promotable mirror volume is promoted to a read-write volume, the snapshot schedule specified for the mirror is used for the promoted read-write volume. Once a mirror volume is promoted to a read-write volume, the mirror schedule is disabled.

### Manually Offloading Data to a Cold Tier

#### Prerequisites

- You must be a fabric user to perform this operation.
- Data tiering must have been enabled on the volume during the volume creation, to be able to offload/recall data.
- To offload data, you must create remote targets. See [Creating a Remote Target](#) on page 189 to add a new remote target.
- To manage data offloading, you must have created storage policies. See [Administering Storage Policies](#) on page 184 to learn more about managing storage policies.

**About this task**

Data, once offloaded, is purged on the the data fabric cluster to release the disk space. When you delete an entire file, part of a file, or a snapshot, corresponding objects are removed from the tier.

Data is offloaded to the tier in the same state, compressed or uncompressed, as was stored in the front-end volume. If data encryption is enabled on the front-end volume (using the `data-encryption` parameter), data is encrypted during and after offload.

At the volume level, data can be offloaded manually by triggering the offload operation.

Follow the steps given below to offload data manually from a volume to a cold tier.

**Procedure**

1. Log on to the Data Fabric UI.
2. Select **Fabric user** on the home page.
3. Click the Table view icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume.
4. Click the ellipsis under **Actions** for the required volume.
5. Select **Offload data**.

**Results**

The data offload begins to the designated cold tier.

**Recalling Data to the Data Fabric File System****Prerequisites**

- You must be a fabric manager or a fabric user to perform this operation.
- Data tiering must have been enabled on the volume during the volume creation, to be able to offload/recall data.
- Offloaded data must be present on the cold tier.

**About this task**

When you read data that has been offloaded to a remote target (or cold tier), data is automatically recalled to the file system to allow the read to succeed.

The recalled data is automatically:

- Purged based on the expiration time period set at the volume level for recalled data if there are no changes (for example, read operation).
- Offloaded based on the rule and the expiration time period set at the volume level for recalled data if there are changes (for example, overwrite operation).

For a cold tiering volume, you must explicitly recall the volume before running any analytics jobs.

Follow the steps given below to recall data manually from a volume to a cold tier.

**Procedure**

1. Log on to the Data Fabric UI.
2. Select **Fabric user** on the Home Page.

3. Click the Table view icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume.
4. Click the ellipsis under **Actions** for the required volume.
5. Select **Recall data**.

### Results

The data recall operation begins and data is recalled on to the fabric file system.

### Administering Storage Policies

Manage storage policies related to data tiering.

You can configure a storage policy (or rules) for data at the volume level.

A storage policy simplifies the lifecycle management of data in the volume including automated migration of files to low-cost storage alternatives. A storage policy contains rules for files that have a well-defined lifecycle or for files you want to switch to different storage tiers during their lifecycle.

You can specify the rules, at the volume level, to selectively identify files to offload (such as file size, file owner, and file modification time), the schedule for offloading the data (for example, two months after file modification), and the settings for storing (such as the location and credentials for the tier) and recalling the offloaded data. You can configure one rule per volume. You can also associate a schedule to automatically offload data at scheduled intervals based on the associated rules.

Data offload is driven by rules, which are configured per volume. Data offload rule can be based on size of file (*s*), owner (*u*, *g*, or *p*) of the file, and/or file modification timestamp (*m*). You can apply one rule per volume.

When a rule is associated with a volume, the rule is first applied on the files in the tiering-enabled volume. When applied on the files in the tiering-enabled volume, the offload is triggered for all files in the snapshot chain as well when the criteria in the rule is met. If the file does not exist in the tiering-enabled volume, rule is applied on the latest state of the file in the snapshot chain. If the file exists in the tiering-enabled volume but has no latest state or if the file was deleted in the tiering-enabled volume, offload does not happen.

Rules can be defined using a combination of the following:

u	<p>Username or user ID, as configured in the OS registry (such as <code>/etc/passwd</code> file, LDAP, etc.), of a specific user.</p> <p><b>Usage:</b> <code>u:&lt;username or user ID&gt;</code></p>
g	<p>Group name or group ID, as configured in the OS registry (such as <code>/etc/group</code> file, LDAP, etc.), of a specific group.</p> <p><b>Usage:</b> <code>g:&lt;groupname or group ID&gt;</code></p>
a	<p>(<i>atime</i>) Time (in seconds or days) since the files were last accessed. The number of seconds can be specified by appending <i>s</i> to value and the number of days can be specified by appending <i>d</i> to the value.</p> <p><b>Usage:</b></p> <p><b>NOTE:</b> If the system time on CLDB and file server nodes are different, the <i>atime</i> rule for offloading data may not work as intended.</p> <p>This tier rule is matched and files are offloaded, when <b>all</b> of the following conditions are met:</p> <p>Assume that the <i>atime</i> feature is enabled on the volume and that the time in the rule is set to <b>a:300s</b>. Based on this rule, all files that are not accessed since 300s, are offloaded. However, this rule is valid only if time since <i>atime</i> tracking is enabled, is more than 300s. The volume level parameter <code>atimeTrackingStartTime</code> denotes the start time of <i>atime</i>.</p> <p>For more information, see <a href="#">Tuning Last Access Time</a>.</p>



m	<p>(mtime) Time (in seconds or days) since the files were last modified. The number of seconds can be specified by appending <i>s</i> to value and the number of days can be specified by appending <i>d</i> to the value.</p> <p><b>Usage:</b></p> <p>All files that are not modified since the specified amount of time, are offloaded.<b>NOTE:</b> If the system time on CLDB and file server nodes are different, the mtime rule for offloading data may not work as intended.</p>
s	<p>The size of the file in bytes, kilobytes, megabytes, or gigabytes. The size of the file can be specified by appending one of the following to the value: <i>b</i> for bytes, <i>k</i> for kilobytes, <i>m</i> for megabytes, or <i>g</i> for gigabytes.</p> <p><b>Usage</b></p> <p>All files whose size exceeds the specified size are offloaded.</p>

Or, use the following:

p	(Default) Specifies all files. Specifies that this operation is applicable to all the files without restriction. This cannot be combined with any other operator.
" "	Indicates none of the files. Specifies that this operation cannot be performed on any of the files.

Use the following to string multiple criteria for offload:

&	AND operation to combine multiple expressions as the criteria for the rule.
	OR operation to indicate either of the expressions as the criteria for the rule.
( )	Delimiters for subexpressions.

## Creating a Storage Policy

### Prerequisites

You must be a fabric manager to perform this task.

### About this task

The expression for a storage policy can comprise the following parameters:

- *u* denotes user.
- *g* denotes group.
- *s* denotes size. Append the value with the unit - *b*, *k*, *m*, *g* for bytes, KB, MB, or GB respectively.
- *m* denotes last modified time. Append the value with the unit - *s*, *d* for seconds or days respectively
- *a* denotes last accessed time. Append the value with the unit - *s*, *d* for seconds or days respectively.

Use & to add an AND condition and | to add an OR condition, to combine multiple parameter-value pairs.

Example: To offload data to cold tier for user *jdoe*, when storage size reaches *100 MB*, and the last accessed time is *7 days*, the expression is *u:jdoe&s:100m&a:7d*

Follow the steps given below to create a storage policy.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric manager** option from the dropdown on the Home page.

3. Click **Fabric Administration** seen on the Home page.
4. Select from the fabrics dropdown the fabric for which you wish to add the storage policy.
5. Scroll down to the **Storage policies** card.



**NOTE:** Alternatively, you can click **Global namespace**, click the fabric link in the table view, and navigate to the **Storage policies** tab for the fabric.

6. Click **Create storage policy**.
7. Enter the **Name** for the storage policy.
8. Enter the expression for the storage policy.
9. Click **Create**.

### Results

The storage policy is created successfully when you enter valid values and units for the parameters specified in the policy.

### Editing a Storage Policy

#### About this task

You can modify the basic rule to:

- Add or remove users and/or groups.
- Change the name of the users and/or groups.
- Change the number of days since the file was last modified for users and/or groups.

If you switch from a basic rule to an advanced rule, all expressions from the basic rule are carried over to the advanced rule. You can modify an advanced rule using a combination of the following expressions:

u	Username or user ID, as configured in the OS registry (such as <code>/etc/passwd</code> file, LDAP, etc.), of a specific user. <b>Usage:</b> u:<username or user ID>
g	Group name or group ID, as configured in the OS registry (such as <code>/etc/group</code> file, LDAP, etc.), of a specific group. <b>Usage:</b> g:<groupname or group ID>

a	<p>(<i>atime</i>) Time (in seconds or days) since the files were last accessed. The number of seconds can be specified by appending <i>s</i> to value and the number of days can be specified by appending <i>d</i> to the value.</p> <p><b>Usage:</b></p> <ul style="list-style-type: none"> <li>"a:&lt;value&gt;s" — specifies <i>atime</i> in seconds</li> <li>"a:&lt;value&gt;d" — specifies <i>atime</i> in days</li> </ul> <p><b>NOTE:</b> If the system time on CLDB and file server nodes are different, the <i>atime</i> rule for offloading data may not work as intended.</p> <p>This tier rule is matched and files are offloaded, when <b>all</b> of the following conditions are met:</p> <ol style="list-style-type: none"> <li><i>atime</i> tracking is enabled at volume level</li> <li>Time since <i>atime</i> that is configured on the volume is more than the time specified in the rule</li> <li>Duration since the file was last accessed is more than the time specified in the rule</li> </ol> <p>Assume that the <i>atime</i> feature is enabled on the volume and that the time in the rule is set to <b>a:300s</b>. Based on this rule, all files that are not accessed since 300s, are offloaded. However, this rule is valid only if time since <i>atime</i> tracking is enabled, is more than 300s. The volume level parameter <i>atimeTrackingStartTime</i> denotes the start time of <i>atime</i>.</p> <p>For more information, see <a href="#">Tuning Last Access Time</a>.</p>
m	<p>(<i>mtime</i>) Time (in seconds or days) since the files were last modified. The number of seconds can be specified by appending <i>s</i> to value and the number of days can be specified by appending <i>d</i> to the value.</p> <p><b>Usage:</b></p> <ul style="list-style-type: none"> <li>"m:&lt;value&gt;s" — specifies <i>mtime</i> in seconds</li> <li>"m:&lt;value&gt;d" — specifies <i>mtime</i> in days</li> </ul> <p>All files that are not modified since the specified amount of time, are offloaded.<b>NOTE:</b> If the system time on CLDB and file server nodes are different, the <i>mtime</i> rule for offloading data may not work as intended.</p>
s	<p>The size of the file in bytes, kilobytes, megabytes, or gigabytes. The size of the file can be specified by appending one of the following to the value: <i>b</i> for bytes, <i>k</i> for kilobytes, <i>m</i> for megabytes, or <i>g</i> for gigabytes.</p> <p><b>Usage</b></p> <ul style="list-style-type: none"> <li>"s:&lt;value&gt;b" — specifies file size in bytes</li> <li>"s:&lt;value&gt;k" — specifies file size in KB</li> <li>"s:&lt;value&gt;m" — specifies file size in MB</li> <li>"s:&lt;value&gt;g" — specifies file size in GB</li> </ul> <p>All files whose size exceeds the specified size are offloaded.</p>

Or, use the following:

p	(Default) Specifies all files. Specifies that this operation is applicable to all the files without restriction. This cannot be combined with any other operator.
"	Indicates none of the files. Specifies that this operation cannot be performed on any of the files.

Use the following to string multiple criteria for offload:

&	AND operation to combine multiple expressions as the criteria for the rule.
	OR operation to indicate either of the expressions as the criteria for the rule.
( )	Delimiters for subexpressions.

You cannot switch from an advanced rule that includes the following to a basic rule because the following are not supported in a basic rule:

- *p* — All the files
- *s* — The size of the file
- & — The AND operation used for specifying multiple users (*u*), groups (*g*), or criteria
- | — The OR operation used with *s* or *m*
- " " — None of the files.
- ( ) — Subexpressions

**NOTE:** The basic rule must contain *mtime* (*m*). It can also include one or more users or groups separated by the OR operation (|).

Follow the steps given below to edit a storage policy.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric manager** option from the dropdown on the Home page.
3. Click **Fabric Administration** seen on the Home page.
4. Select from the fabrics dropdown the fabric for which you wish to edit the storage policy.
5. Scroll down to the **Storage policies** card.



**NOTE:** Alternatively, you can click **Global namespace**, click the fabric link in the table view, and navigate to the **Storage policies** tab for the fabric.

6. Click the ellipsis seen under **Actions** for the storage policy to edit.
7. Click **Edit**.
8. Make the required changes.
9. Click **Save**.

### Results

The changes to the storage policy are saved successfully.

### Deleting Storage Policy

Delete storage policy.

### Prerequisites

The storage policy must not be associated with a volume. A storage policy that is associated with a volume cannot be deleted.

## About this task

You can delete a storage policy that is no longer required, and has no association with any volume on any fabric.

Follow the steps given below to delete a storage policy.

## Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric manager** option from the dropdown on the Home page.
3. Click **Fabric Administration** seen on the Home page.
4. Select from the fabrics dropdown the fabric for which you wish to delete a storage policy.
5. Scroll down to the **Storage policies** card.



**NOTE:** Alternatively, you can click **Global namespace**, click the fabric link in the table view, and navigate to the **Storage policies** tab for the fabric.

6. Click the ellipsis seen under **Actions** for the storage policy to edit.
7. Click **Delete**.
8. Confirm the deletion.

## Results

The storage policy is deleted.

## Administering Remote Targets

Data Fabric provides rule-based automated tiering functionality that allows you to seamlessly integrate with third party cloud object storage as an additional storage tier in the data fabric cluster to store file data that is rarely accessed or archived data, which is referred to as cold data. See

A cold tier is referred to as remote target on the Data Fabric UI.

The remote storage where cold data can be offloaded is called a remote target.

A remote target has a bucket on the third party cloud store where volume data is offloaded based on the policy configured by the fabric manager.

Volume data in 64KB data chunks is packed into 8MB sized objects and offloaded to the bucket on the tier and the corresponding volume metadata is stored in a visible tier-volume as HPE Ezmeral Data Fabric Database tables on the data fabric cluster. During writes and reads, volume data is recalled to the data fabric cluster, if necessary. Data written to the volume is periodically moved to the remote target, releasing the disk space on the filesystem.

You can associate a volume with a remote target.

For cold data, you can offload your cluster data to public, private, and hybrid clouds. You can offload data to remote cloud from vendors such as Amazon AWS, Google Cloud Platform, Microsoft Azure, IBM Cleversafe, Hitachi HCP, and Minio. You can tap into cloud-scale capacity for cold data.



**NOTE:** Data Fabric supports tiering for only file and volume data; tiering of tables and streams is not supported.

## Creating a Remote Target

Create a remote target to offload cold data.

### Prerequisites

You must be a fabric manager to perform this task.

### About this task

You can create one or more remote targets to offload cold data to the remote target.

Follow the steps given below to create a remote target.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric manager** option from the dropdown on the Home page.
3. Click **Fabric Administration** seen on the Home page.
4. Select from the fabrics dropdown the fabric for which you wish to add a remote target.
5. Scroll down to the **Remote targets** card.



**NOTE:** Alternatively, you can click **Global namespace**, click the fabric link in the table view, and navigate to the **Remote targets** tab for the fabric.

6. Click **Create remote target**.
7. Enter the **Name** for the remote target.
8. Select a cloud provider from the **Vendor** dropdown.
9. Enter the **URL**.
10. Select the **Bucket**.
11. Enter the **Region**.
12. Enter the **Access key**.
13. Enter the **Secret key**.
14. Click **Create**.

### Results

The remote target is created successfully.

The remote target can be used as a cold tier for data tiering. You can associate a volume with the remote target to offload the cold data from the volume on to the remote target.

### Editing Remote Target Credentials

Edit credentials for a remote target.

### Prerequisites

You must be a fabric manager to perform this operation.

### About this task

You can edit the credentials for a remote target.

Follow the steps given below to edit remote target credentials.

**Procedure**

1. Log on to the Data Fabric UI.
2. Select **Fabric manager** option from the dropdown on the Home page.
3. Click **Fabric Administration**.
4. Select from the fabrics dropdown the fabric for which the remote target has been created.
5. Scroll down to the **Remote targets** card.



**NOTE:** Alternatively, you can click **Global namespace**, click the fabric link in the table view, and navigate to the **Remote targets** tab for the fabric.

6. Click the ellipsis under **Actions** for the remote target whose credentials you wish to change.
7. Click **Edit Credentials**.
8. Modify the **Access key**, **Secret key** or both, as required.
9. Click **Save**.

**Results**

The updated credentials are saved for the remote target.

**Deleting a Remote Target**

Delete a remote target.

**Prerequisites**

You must be a fabric manager or an infrastructure admin to perform this operation.

The remote target to delete must not be associated with a volume.

**About this task**

You can delete a remote target that is not in use. Before you delete a remote target, ensure that the data on the remote target is backed up.

Follow the steps given below to delete a remote target.

**Procedure**

1. Log on to the Data Fabric UI.
2. Select **Fabric manager** option from the dropdown on the Home page.
3. Click **Fabric Administration** seen on the Home page.
4. Select from the fabrics dropdown the fabric for which you wish for which the remote target has been created.
5. Scroll down to the **Remote targets** card.



**NOTE:** Alternatively, you can click **Global namespace**, click the fabric link in the table view, and navigate to the **Remote targets** tab for the fabric.

6. Click the ellipsis under **Actions** for the remote target whose credentials you wish to delete.
7. Click **Delete**. Confirm the deletion.

## Results

The remote target is deleted. The data on the cold tier is inaccessible. The remote target is no longer available for data tiering.

## Administering Schedules

Introduction to schedules.

A schedule is a group of rules that specify recurring points in time at which certain actions are determined to occur. You can use schedules to automate the creation of snapshots and mirrors and the offload of volume data to a storage tier; after you create a schedule, it appears as a choice in the scheduling menu when you are creating or editing a volume.

When you specify a *snapshot* schedule on a mirror volume, it specifies how often to take a snapshot of the mirror volume. This snapshot schedule is distinct from the snapshot schedule for the standard volume. A snapshot schedule for a promotable mirror volume has two purposes:

- It specifies how often to take a snapshot of the mirror volume for the purpose of preserving the state of the mirror before a subsequent mirror operation. This way, if corrupt data is copied from the source volume's snapshot into the mirror volume, the mirror contents can be rolled back to the snapshot.
- If the promotable mirror volume is promoted to a read-write volume, the snapshot schedule specified for the mirror is used for the promoted read-write volume. Once a mirror volume is promoted to a read-write volume, the mirror schedule is disabled.

A *mirror* schedule specifies how frequently the mirror volume is synchronized with the source volume. In case of a disaster (or any type of data loss on a read-write source volume), the data can be recovered from the mirror volume, but any data written to the source volume since the last successful mirror operation will not be on the mirror volume. Therefore, you should set the mirror schedule such that it meets your RPO (Recovery Point Objective).

A *tier offload* schedule specifies how frequently data in the volume on the fabric is offloaded to the tiered storage. This setting to automatically offload data to the storage tier.

## Creating a Schedule

Add a schedule for data tiering.

## Prerequisites

You must be a fabric manager to perform this task.

## About this task

A schedule is a group of rules that specify recurring points in time at which certain actions are determined to occur. You can use schedules to automate the creation of snapshots and mirrors. A schedule can be attached to a volume. If the

A schedule can have one or more rows.

A rule is made up of three elements.

- Frequency of the schedule such as yearly, daily, weekly, or specific time interval in minutes.
- Time at which schedule is to run
- Time period for which data is to be retained

Follow the steps given below to create a schedule.

## Procedure

1. Log on to the Data Fabric UI.



2. Select **Fabric manager** option from the dropdown on the Home page.
3. Click **Fabric Administration**.
4. Select from the fabrics dropdown the fabric for which you wish to create a schedule.
5. Scroll down to the **Schedules** card.
6. Click **Create schedule**.
7. Enter the **Name** for the schedule.
8. Configure the rule. Select the frequency and time to trigger the schedule, and the data retention period.
9. Click **Create**.

### Results

The schedule is created with the specified rule or rules.

### Editing a Schedule

Edit an existing schedule.

### Prerequisites

You must be a fabric manager to edit a schedule.

### About this task

You can make changes to an existing schedule to incorporate any changes you wish to make to the schedule.

Follow the steps given below to edit a schedule.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric manager** option from the dropdown on the Home page.
3. Select from the fabrics dropdown the fabric for which you wish the schedule has been created.
4. Click **Fabric Administration**.
5. Click the pencil icon for the schedule to edit.
6. Make the necessary changes and click **Save**.

### Results

The changes to the schedule are saved and the new schedule is applied to the data on the volumes to which the schedule applies.

### Scheduling Volume Data Tiering

#### About this task

After creating a schedule, you can associate it with the tiering-enabled volume when you create or modify the volume. If a schedule for offloading data is associated with the volume, data is offloaded automatically as scheduled based on the rules associated with the volume for offloading data.

For volumes enabled for cold tiering, you must assign a schedule to automatically offload data; if you do not assign a schedule, data is not offloaded automatically and you must manually run the offload command to offload data. See

[Manually Offloading Data to a Cold Tier](#) on page 182 for details on manually offloading data.

Follow the steps given below to schedule tiering of data on a volume.

### Procedure

1. Log on to the Data Fabric UI.
2. Click the Table view icon on the **Resources** card.
3. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume with which you wish to associate a data tiering schedule.
4. Click the volume name seen under **Resource name**.
5. Navigate to the **Settings** tab.
6. Under **Schedules**, click the pencil icon seen next to **Object tiering**.
7. Select a suitable schedule option for the volume.
8. Click **Select**.

### Results

The selected schedule is applied to the volume and data tiering is done on the designated cold tier per schedule.

### Viewing Schedules

View a list of schedules.

### Prerequisites

You must be an infrastructure admin or a fabric manager to view the list of schedules.

### About this task

You can view a list of all existing schedules to understand what the various schedules are. A fabric manager can decide if the required schedule already exists or it is required to create new schedules.

Follow the steps given below to view a schedule.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric manager** option or the **Infrastructure admin experience** option from the dropdown on the Home page.
3. Click **Fabric Administration**.
4. Select from the fabrics dropdown the fabric for which you wish view the schedules.
5. Scroll down to the **Schedules** card.

### Results

You are able to see a list of existing schedules.

### Deleting a Schedule

Delete a schedule.

**Prerequisites**

You must be a fabric manager to delete a schedule.

**About this task**

You can delete a schedule that is not associated with a volume.

**TIP:** Attach another schedule to a volume to overwrite the existing schedule (schedule to delete) attached to the volume. Once another schedule is attached, you can delete the old schedule.

Follow the steps given below to delete a schedule.

**Procedure**

1. Log on to the Data Fabric UI.
2. Select **Fabric manager** option from the dropdown on the Home page.
3. Click **Fabric Administration**.
4. Select from the fabrics dropdown the fabric for which you wish to delete the schedule.
5. Click the bin icon for a schedule, to delete a schedule. Confirm the deletion.

**Results**

The schedule is deleted.

**Mirroring**

Synopsis of mirrors and mirroring process.

Creating a mirror volume is similar to creating a normal read/write volume. However, when you create a mirror volume, you must specify a source volume from which the mirror retrieves content. This retrieval is called the mirroring operation. Like a normal volume, a mirror volume has a configurable replication factor. Only one copy of the data is transmitted from the source volume to the mirror volume. HPE Ezmeral Data Fabric volumes can only be mirrored and NOT replicated. However, the source and mirror volumes handle their own internal HPE Ezmeral Data Fabric filesystem replication (which is based on the replication factor) independently. file system internally replicates source and mirror volumes independently of each other.

Volume mirroring from a lower HPE Ezmeral Data Fabric version to higher HPE Ezmeral Data Fabric version is supported. Volume mirroring from a higher Data Fabric version to a lower Data Fabric version is not supported.

**Mirroring Process**

The HPE Ezmeral Data Fabric system creates a temporary snapshot of the source volume at the start of a mirroring operation. The mirroring process reads content from the snapshot into the mirror volume. The source volume remains available for read and write operations during the mirroring process.

If the mirroring operation is schedule-based, the snapshot expires according to the value of the schedule's **Retain For** parameter. Snapshots created during manual mirroring persist until they are deleted manually.

The mirroring process transmits only the differences between the source volume and the mirror. The initial mirroring operation copies the entire source volume, but subsequent mirroring operations can be extremely fast.

**Local Mirroring**

A *local mirror volume* is a mirror volume whose source is on the same cluster. Local mirror volumes are useful for load balancing or for providing a read-only copy of a data set.

You can locate your local mirror volumes in specific servers or on racks with particularly high bandwidth, mounted in a public directory separate from the source volume.

The most frequently accessed volumes in a cluster are likely to be the root volume and its immediate children. To load-balance read operations on these volumes, mirror the root volume (typically `mapr.cluster.root`, which is mounted at `/`). By mirroring these volumes, you can serve read requests from the mirrors, and distribute load across the nodes. Less-frequently accessed volumes that are lower in the hierarchy do not need mirror volumes. Since the mount paths for those volumes are not mirrored throughout, those volumes are writable.

If you are creating a local mirror of the root volume, `root(/)` points to the mirror volume, hence `root` is read-only. For read-write copy of `root (/)`, you must use the special path, `/.rw`

## Remote Mirroring

A remote mirror volume is a mirror volume with a source in another cluster. You can use remote mirrors for offsite backup, for data transfer to remote facilities, and for load and latency balancing for large websites. By mirroring the cluster's root volume and all other volumes in the cluster, you can create an entire mirrored cluster that keeps in sync with the source cluster.

Backup mirrors for disaster recovery can be located on physical media outside the cluster, or in a remote cluster. If disaster strikes the source cluster, you can check the time of last successful synchronization to determine the freshness of the backup.

Once data volumes are created in a primary data center, the Data Fabric administrator creates mirror volumes in a remote secondary data center.

## Starting Volume Mirroring

Start mirroring of data on a mirror volume.

### Prerequisites

- You must have created a mirror volume, on which you can mirror data from the source volume associated with the mirror volume.
- Data must be present on the source volume.

### About this task

Data from the associated source volume can be mirrored to a mirror volume.

Follow the steps given below to start mirroring of data on to a mirror volume.

### Procedure

1. Log on to the Data Fabric UI.
2. Under the default **Fabric user experience**, click the Table view icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume.
3. Click the volume name seen under **Resource Name**.
4. Navigate to the **Mirrors** tab.
5. Click the ellipsis under **Actions** for the required volume.
6. Click the **Start mirroring** option, and click Start on the **Start mirroring** message box to confirm that you wish to start mirroring.

**Results**

Mirroring of data on the source volume is triggered. The status under the **Mirroring** column on Mirrors tab for the volume changes to **On**. The percentage of data mirrored from the source volume is displayed for the volume under the **Mirrored** column on the **Mirrors** tab.

**Stopping Volume Mirroring**

Stop mirroring of data that is in progress on a mirror volume.

**Prerequisites**

Data mirroring must be in progress on the mirror volume.

**About this task**

You can stop mirroring of data that is in progress from the associated source volume onto a mirror volume.

Follow the steps given below to stop mirroring of data on to a mirror volume.

**Procedure**

1. Log on to the Data Fabric UI.
2. Under the default **Fabric user experience**, click the Table View icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume.
3. Click the volume name seen under **Resource Name**.
4. Navigate to the **Mirrors** tab.
5. Click the ellipsis under **Actions** for the volume being mirrored.
6. Click the **Stop mirroring** option.

**Results**

Mirroring of data that is in progress from the source volume onto the mirror volume is stopped.. The status under the **Mirroring** column on Mirrors tab for the volume changes to **Off**.

**Scheduling Volume Mirroring****Prerequisites**

- You must have the privilege to schedule mirroring.
- The volume for which you are scheduling mirroring must be a mirror volume.

**About this task**

When you choose the mirror schedule, consider the amount of data on the volume and the load on the cluster. Remember that the mirroring frequency must allow enough time to complete one mirror operation before the next scheduled mirror operation starts. In addition, if you have a cascaded mirror setup (where A mirrors to B which mirrors to C), you cannot set a mirror schedule for C that starts before B finishes mirroring from A.



**WARNING:** Avoid setting a mirror schedule for more often than every 30 minutes.

Follow the steps given below to schedule mirroring on a volume.

**Procedure**

1. Log on to the Data Fabric UI.

2. Under the default **Fabric user experience**, click the Table view icon on the **Resources** card.
3. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume with which you wish to associate a mirroring schedule.
4. Click the volume name seen under **Resource name**.
5. Navigate to the **Settings** tab.
6. Under **Schedules**, click the pencil icon seen next to **Mirror**.
7. Select a suitable schedule option for the volume.
8. Click **Select**.

### Results

The selected schedule is applied to the volume and data mirroring is performed per the schedule.

## Administering Nodes

---

This section describes node-administration and maintenance operations that you can perform using the Data Fabric UI.

A *node* is an individual physical or virtual machine in a Data Fabric where a set of services is running.

### Viewing Node Information

This section describes how to use the Data Fabric UI to view the node status and other information on the **Nodes** tab.

You must be a Fabric manager to view node information. In addition, node information is available only for on-premises fabrics; it is not available for cloud fabrics.

To view node information for an on-premises fabric:

1. Log on to the Data Fabric UI.
2. Select the **Fabric manager** option.
3. Click **Global namespace**, and check the table view.
4. Under the **Resource Name** column, click the link for the fabric to which you want to display node information.
5. Click the **Nodes** tab. The node list table appears.

### Adding Nodes (On-premises Deployment)

This section describes how to add one or more nodes to an installed on-premises deployment by using the Data Fabric UI.

Adding nodes to an installed fabric can help in several ways. For example, you can add nodes to increase the total storage, redistribute resources, boost the compute capacity, or improve the availability of certain services. Adding nodes is sometimes referred to as a *scaling operation* because it allows you to scale the storage and performance of the fabric.

#### Limitations for Adding Nodes

Note the following limitations for adding nodes by using the Data Fabric UI:

- The Data Fabric UI supports adding nodes to both on-premises and cloud deployments. To add nodes to cloud deployments, such as AWS, Azure, and GCP, see [Adding Nodes \(Cloud Deployment\)](#) on page 202.
- You can use the Data Fabric UI to add nodes to an on-premises deployment only if the fabric is running release 7.7.0 or later.
- You cannot use the Data Fabric UI to add nodes to a customer-managed cluster.  
To add nodes to a customer-managed cluster, you must use the Installer for customer-managed platforms or manual steps. See [Extending a Cluster by Adding Nodes](#) or [Adding Nodes to a Cluster](#).
- If the fabric has fewer than three nodes, adding nodes requires restarting ZooKeeper and Warden one node at a time on the previously installed nodes. During the restart phase, access to each node is lost temporarily. If the fabric has three or more nodes, the restart is not needed.
- On any given fabric, only one scaling operation can be done at a time. If a scaling operation is already in progress, you cannot initiate another scaling operation on the same fabric. However, scaling operations can be performed at the same time on different fabrics belonging to the same global namespace.
- **Add node** operation is not possible when an upgrade operation is in progress. In addition, [Upgrading a Data Fabric](#) on page 80 is not supported when an **Add node** operation is in progress.
- The Data Fabric UI does not currently provide an option to remove a node from a fabric.

### Prerequisites for Adding Nodes


Note the following prerequisites for adding nodes by using the Data Fabric UI:

- To add nodes, you must have fabric manager [credentials](#).
- There is no restriction on the size of a fabric for adding nodes. You can add nodes to any fabric, including a one-node fabric.
- Adding nodes to an on-premises installation requires that you provide the host nodes *before* starting the scaling operation. There is no restriction on the number of nodes you can add, but the nodes must meet certain requirements. Before beginning the operation, review [Prerequisites for On-Premises Installation](#) on page 27.

### Steps for Adding Nodes


Use the following steps to add nodes:

1. Log on to the Data Fabric UI.
2. Select the **Fabric manager** option.
3. On the **Global namespace** card, check the table view.
4. Under the **Resource Name** column, click the name of the fabric for which you want to add nodes. The **Fabric details** are displayed.
5. Click the **Nodes** tab. For more information, see [Viewing Node Information](#) on page 198.
6. Click **Add node**. The **Add node** side drawer appears.
7. Fill in the required node parameters:

For this parameter . . .	Do this
<b>SSH credentials*</b>	Enter your <b>Username</b> for SSH access to the node. Enter your <b>Password</b> .
<b>Nodes*</b>	Enter the fully qualified domain name (FQDN) of the first node you want to add. To add another node, click <b>+ Add node</b> , and enter the next FQDN. If you enter a node name by mistake, click the trash can (  ) icon to remove it.
<b>Network settings</b>	If necessary, specify the optional network settings: <b>EDF subnet</b> – Specify one or more comma-separated subnet masks. For example: <pre>10.10.15.0/24,10.10.16.0/24</pre> <b>EDF external</b> – Specify a comma-separated list of tuples of host names and external IP addresses. For example: <pre>node1:1.1.1.1,node2:1.1.1.2,node3:1.1.1.3</pre> For more information about the network settings, see <a href="#">Using the MapR Subnet and MapR External Advanced Options</a> .

\*This field is mandatory.


- Click **Add node**. The Data Fabric UI displays a **Nodes in scale operation** table that includes a **Status** column.
- Check the status of the scaling operation:

To	Do this
View the status for a single node	Hold your cursor over the information icon (  ). The icon is displayed when the node status is <code>Installing</code> . Hovering over the tooltip displays the installation status as a percentage. Or, in the <b>Actions</b> column, click the ellipsis for a node, and select <b>View progress</b> to display <b>Node details</b> information.
View the status for all nodes	Click <b>View progress</b> .

If the scaling operation is successful, the Data Fabric UI displays a success message, and the table of active nodes is updated.

### Troubleshooting the Scaling Operation

If the scaling operation fails, the Data Fabric UI displays a status of **Failed** in the **Nodes in scale operation** table. To obtain more information:

- Examine failure data for the entire scaling operation by clicking **View progress** above the table. Or, click the ellipsis (  ) in the **Action** column, and select **View progress** (for a single node).
- If possible correct the root cause of the failure.
- Click **Reinitiate** to retry adding the node. The original side drawer is displayed, allowing you to make any needed configuration changes.
- Click **Reinitiate**.

If you are not able to resolve the issue that caused the scaling operation to fail, contact [HPE Support](#).



## Related maprcli Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` commands. Links to these commands are provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `installer clusterscale`
- `installer clusterscalestatus`

## Removing Nodes (On-premises Deployment)

This section describes how to remove one or more nodes from an installed on-premises deployment by using the Data Fabric UI.

Removing nodes from an installed fabric is necessary in many situations, such as:

- To remove the nodes that have reached End of Life.
- To perform maintenance activity, for example, applying security patches, or upgrading the operating system, and so on.
- To perform some activities mandated by organization that need node removal.
- To scale down a fabric that was scaled up to meet a temporary requirement of increased storage, and compute.

### Limitations for Removing Nodes

Consider the following limitations, for removing nodes by using the Data Fabric UI:

- You can use the Data Fabric UI to remove nodes from on-premises deployment only if the fabric is running release 7.9.0 or later.
- You cannot use the Data Fabric UI to remove nodes from a customer-managed cluster.
- You can perform **Delete** node operation only if the fabric has minimum of three nodes.
- You can perform **Delete** node operation only if the cluster storage consumption is less than 80 percent.
- You cannot delete the installer host, that was used to configure the fabric.
- Before proceeding to remove CLDB Master node, perform *Failover of CLDB master* first, and then proceed with **Delete** node operation.
- If you remove a control node (that have CLDB and ZooKeeper running), cluster will be reconfigured with the remaining CLDBs, and Warden will be restarted on each node of the fabric, after removing the control node.
- On any given fabric, only one *Remove node* operation can be done at a time. If a *Remove node* operation is already in progress, you cannot initiate another *Add Node* or *Remove node* operation on the same fabric. However, multiple *Remove node* operations can be performed at the same time on different fabrics belonging to the same global namespace.
- *Remove node* operation is not possible when an *Upgrade fabric*, *Remove fabric*, or *Add node* operation is already in progress.


### Prerequisites for Deleting Nodes

Note the following prerequisite for removing nodes by using the Data Fabric UI:

- To remove nodes, you must have fabric manager [credentials](#).

## Steps for Deleting Nodes

Use the following steps to remove nodes:

1. Log on to the Data Fabric UI.
2. Select the **Fabric manager** option.
3. On the **Global namespace** card, check the table view.
4. Under the **Resource Name** column, click the name of the fabric for which you want to remove nodes. The **Fabric details** are displayed.
5. Click the **Nodes** tab. For more information, see [Viewing Node Information](#) on page 198.
6. Click the ellipsis (  ) in the **Action** column, and select **Delete** . The **Delete node** side drawer appears.
7. Fill in the following node parameters:

For this parameter . . .	Do this
<b>SSH credentials*</b>	Enter your <b>Username</b> for SSH access to the node. Enter your <b>Password</b> .

8. Click **Delete**. The Data Fabric UI displays **Deleting** in **Status** column for the node that is being deleted.  
If the *Remove node* operation is successful, the Data Fabric UI displays a success message, and the table of active nodes is updated

## Related maprcli Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` commands. Links to these commands are provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- [installer clusterscalestatus](#)
- [Removing Nodes \(On-premises Deployment\)](#) on page 201

## Adding Nodes (Cloud Deployment)

This section describes how to add one or more nodes to a cloud deployment by using the Data Fabric UI.

Adding nodes to a cloud-deployed fabric can help in several ways. For example, you can add nodes to increase the total storage, redistribute resources, boost the compute capacity, or improve the availability of certain services. Adding nodes is sometimes referred to as a *scaling operation* because it allows you to scale the storage and performance of the fabric.

## Limitations for Adding Nodes

Note the following limitations for adding nodes by using the Data Fabric UI:

- You can use the Data Fabric UI to add nodes to a cloud deployment only if the fabric is running release 7.8.0 or later.
- Only one **Add node** operation can be done at a time. Multiple **Add node** operation is not allowed (**Add node** button is disabled when a scaling operation is already in progress).
- Scaling operations can be performed at the same time on different fabrics belonging to the same global namespace.

- **Add node** operation is not possible when an upgrade operation is in progress. In addition, [Upgrading a Data Fabric](#) on page 80 is not supported when an **Add node** operation is in progress.
- Each node has a default host name, such as `aws_instance`, `gcp_instance`, or `aws_instance`. For example, a scale operation to add two nodes on AWS may have default name as `aws_instance1`, `aws_instance2` in the UI.
- During each scale operation, maximum number of nodes you can add, is equal to the number of nodes you already have in the fabric. For example, If a fabric has five nodes already, you can scale the fabric by one node to five nodes maximum, and not more than five nodes. However, you can add more nodes by performing the scaling operation again.
- The Data Fabric UI does not currently support scaling down (removing a node from) a fabric.

### Prerequisites for Adding Nodes

Note the following prerequisites for adding nodes by using the Data Fabric UI:


- To add nodes, you must have logged in with fabric manager [credentials](#).
- You must have the account credentials of the cloud service that you want to scale.

### Steps for Adding Nodes

Use the following steps to add nodes:

1. Log on to the Data Fabric UI.
2. Select the **Fabric manager** option.
3. On the **Global namespace** card, check the table view.
4. Under the **Resource Name** column, click the name of the fabric for which you want to add nodes. The **Fabric details** are displayed.
5. Click the **Nodes** tab. For more information, see [Viewing Node Information](#) on page 198.
6. Click **Add node**. The **Add node** side drawer appears. Depending on the cloud provider of the fabric, such as AWS, Azure, or GCP, respective fields appear.
7. Use the one of the suitable table (GCP, Azure, or AWS) to fill in the required node parameters:

Table

For this parameter . . .	Do this
<b>GCP access credentials*</b>	Click on <b>Select file*</b> , and drag and drop the <b>Service account key file</b> .  <b>NOTE:</b> If you do not have a Service account key file already, you can first create one in Google Cloud, and then upload the file.
<b>Fabric details</b>	<ul style="list-style-type: none"> <li>• <b>Node*</b> - Select the number of nodes to be added.</li> </ul>


\*This field is mandatory.

Table

For this parameter . . .	Do this
<b>Azure access credentials*</b>	Enter the following mandatory information: <ul style="list-style-type: none"> <li>• <b>Subscription ID*</b></li> <li>• <b>Tenant ID*</b></li> <li>• <b>Client ID*</b></li> <li>• <b>Client Secret*</b></li> </ul>
<b>Fabric details</b>	<ul style="list-style-type: none"> <li>• <b>Node*</b> - Select the number of nodes to be added.</li> </ul>

\*This field is mandatory.

Table

For this parameter . . .	Do this
<b>AWS access credentials*</b>	Enter the following mandatory information: <ul style="list-style-type: none"> <li>• <b>Access Key*</b></li> <li>• <b>Secret Key*</b></li> </ul> <p> <b>NOTE:</b> User must have <b>AmazonEC2FullAccess</b> permission.</p>
<b>Fabric details</b>	<ul style="list-style-type: none"> <li>• <b>Node*</b> - Select the number of nodes to be added.</li> </ul>

\*This field is mandatory.

- When you have filled in the required node parameters, click **Submit**.
- Click **Add node**. The Data Fabric UI displays a **Nodes in scale operation** table that includes a **Status** column.
- Check the status of the scaling operation:

To	Do this
View the status for a single node	Hold your cursor over the information icon ( ⓘ ). The icon is displayed when the node status is <i>Installing</i> . Hovering over the tooltip displays the installation status as a percentage. Or, in the <b>Actions</b> column, click the ellipsis for a node, and select <b>View progress</b> to display <b>Node details</b> information.
View the status for all nodes	Click <b>View progress</b> .

If the scaling operation is successful, the Data Fabric UI displays a success message, and the table of active nodes is updated.

### Troubleshooting the Scaling Operation

If the scaling operation fails, the Data Fabric UI displays a status of **Failed** in the **Nodes in scale operation** table. To obtain more information:

- Examine failure data for the entire scaling operation by clicking **View progress** above the table. Or, click the ellipsis ( ⓘ ) in the **Action** column, and select **View progress** (for a single node).

2. If possible correct the root cause of the failure.
3. Click **Reinitiate** to retry adding the node. The original side drawer is displayed, allowing you to make any needed configuration changes.
4. Click **Reinitiate**.

If you are not able to resolve the issue that caused the scaling operation to fail, contact [HPE Support](#).

### Related maprcli Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` commands. Links to these commands are provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `installer clusterscale`
- `installer clusterscalestatus`

### Related concepts

[Adding Nodes \(On-premises Deployment\)](#) on page 198

This section describes how to add one or more nodes to an installed on-premises deployment by using the Data Fabric UI.

## Auditing Fabric and Fabric Data

---

Auditing in Data Fabric

You can enable auditing of cluster administration and data-access operations using the Data Fabric UI.

### Levels of Auditing

There are two levels of auditing:

- Auditing for fabric-level administration operations
- Auditing of data access operations on the fabric

## Enabling/Disabling Fabric Auditing

Enable/disable fabric auditing

### About this task

Follow the steps given below to enable/disable auditing of fabric administration operations.

### Procedure

1. Log on to the Data Fabric UI with Admin or Fabric Manager [credentials](#).
2. Under the default **Fabric user experience**, click the Table view icon on the **Resources** card. Click the link for the fabric in the Fabrics list for which you wish to enable auditing.
3. Navigate to the **Settings** tab on the fabrics page.
4. Under **Fabric Settings**, click the Edit icon.
5. Toggle **Cluster Auditing** to enable auditing on the fabric

**Results**

Auditing of fabric administration operations is enabled on the fabric. After auditing is enabled, audit log entries are generated when fabric administration operations are performed.

**Configuring Auditing for Data Access Operation**

Enable or disable auditing for data access operations on fabric.

**About this task**

Data access auditing can be enabled via the Data Fabric UI.

Follow the steps given below to audit data access operation on a fabric.

**Procedure**

1. Log on to the Data Fabric UI with Admin or Fabric Manager [credentials](#).
2. Click the Table View icon on the **Resources** card. Click the link for the fabric in the Fabrics list for which you wish to enable auditing.
3. Navigate to the **Settings** tab on the fabrics page.
4. Under **Fabric Settings**, click the Edit icon.
5. Select the **On** option for **Configure auditing of data-access operations**.
6. Enter the maximum size for the audit log.
7. Enter the retention period.
8. Click **Update**

**Results**

Auditing for data access operations is enabled on the fabric with the specified audit log size and retention period.

**Administering Security Policies**

---

Add, edit, delete, and manage state of security policies.

A security policy is an access control mechanism that can be applied to data objects on a fabric. Once a security policy is applied, it governs how a user can access data objects on the volume to which the security policy is applied.

A security policy can be associated with a volume.

**TIP:** A security policy is an access control mechanism for data stored on Data Fabric volumes, while a bucket policy is an access control mechanism applied to objects in an S3 object store associated with Data Fabric.

**Security Policy Life Cycle**

The state of a security policy is interpreted as a combination of two parameters:

- allow tagging
- access control


The following table explains the various values of the allow tagging and access control parameters.

Parameter	Accepted Values and Description	Default value
allow tagging	<p>false</p> <ul style="list-style-type: none"> <li>• Disables tagging; users cannot apply the security policy to data objects.</li> <li>• This is the default setting when the fabric manager creates a security policy. The fabric manager can specify the setting explicitly when creating the security policy.</li> <li>• When a security policy is active (allow tagging=true) but needs to be deprecated, modify the policy and set allow tagging=false. This prevents users from tagging any other data objects with the policy. Note that the system continues to enforce the security controls set in the security policy for data objects that were already tagged with the security policy.</li> </ul> <p>true</p> <ul style="list-style-type: none"> <li>• Enables tagging; users can apply the security policy to data objects.</li> <li>• When creating or modifying a security policy, a fabric manager can set allowtagging to true.</li> <li>• When creating a security policy, as a fabric manager, you may want to set this parameter to true to test the security settings in the policy or to use tagging tools to discover data content and tag the data.</li> <li>• To enable a deprecated security policy, set allow tagging to true.</li> </ul>	false

Parameter	Accepted Values and Description	Default value
access control	<p>Disarmed</p> <ul style="list-style-type: none"> <li>Unless the fabric manager changes the setting when creating the security policy, this is the default setting if the fabric manager creates a security policy.</li> <li>The system does not enforce the access permissions set in the security policy during data operations on the data objects tagged with the security policy.</li> </ul> <p>Armed</p> <ul style="list-style-type: none"> <li>The system enforces the permissions set in the security policy during data operations on the data objects tagged with the security policy.</li> <li>When creating or modifying a security policy, as a fabric manager, you can set access control to Armed.</li> <li>To enforce access permissions set in a deprecated security policy, the fabric manager can set access control to Armed. The system continues to enforce access permissions set in the security policy for all data operations on the data objects tagged with the policy.</li> </ul> <p>Denied</p> <ul style="list-style-type: none"> <li>Denies all access to data objects tagged with the security policy.</li> </ul>	Disarmed

You can change the state of a security policy through the `allow tagging and access control` parameters to move a security policy through a life cycle, as shown in the following image where the security policy moves from new to retired.





The following table describes each of the stages in the security policy life cycle:

Stage	Description
new (default)	<ul style="list-style-type: none"> <li>Default upon security policy creation.</li> <li>Users cannot tag data objects with the security policy.</li> <li>The system does not enforce access permissions set in the security policy</li> </ul>
in use	<ul style="list-style-type: none"> <li>Users can tag data objects with the security policy.</li> <li>The system enforces all security controls set in the security policy during data operations on data objects tagged with the security policy.</li> <li>Security controls set in the policy can include access permissions, auditing, and wire-level encryption.</li> </ul>
deprecated	<ul style="list-style-type: none"> <li>Users can no longer tag the security policy to data objects.</li> <li>The system still enforces the security controls set in the security policy for all data operations on the data objects tagged with the policy. Users cannot tag any additional data objects with the policy.</li> </ul>
retired	<ul style="list-style-type: none"> <li>Users cannot tag the security policy to data objects.</li> <li>All data operations on the data objects tagged with the security policy are denied by the system.</li> </ul>

### About Security Policy Domain

Describes a security policy domain.

A security policy domain is a group of fabrics that directly or indirectly share data and use the same security policies to control access to the data. A security policy domain consists of one master fabric and zero or more member security policy fabrics that create a global security policy namespace.

A global policy master is a prerequisite for the creation of security policies. A global policy master is a fabric on which security policies can be created.

You can create and modify security policies only on the fabric that is designated as the global policy master. When you create or update security policies, the policy server updates the `mapr.pbs.base` volume with the security policy metadata. Subsequently, the security policies are mirrored to other member fabrics in the global namespace.

By default, the first fabric or the primary fabric that you create on the global namespace is designated as the global policy master. Hence, it is not required to explicitly assign an alternate global policy master, unless the primary fabric goes down.


Each fabric, to which a security policy is applied, operates independently and, therefore, does not require network connectivity to the global policy master to enforce policies. A security policy server in each of the fabrics enforces the policies and manages the security policy metadata in an internal volume named `mapr.pbs.base`.

See [Security Policy Implementation Workflow](#) on page 210 for details on how to apply security policies to fabric volumes on Data Fabric.

## Security Policy Implementation Workflow

Describes the security policy workflow, in general, and the steps in implementing a security policy.

In the diagram below, all member fabrics, that is, Member Fabric 1, Member Fabric 2, Member Fabric 3, and Member Fabric 4 get the policy metadata from the global policy master.



The security policies are mirrored to all member fabrics in the global namespace from the global policy master.

### 1 – Create a Security Policy Domain

By default, the primary fabric is set as the global policy master and all secondary fabrics are member fabrics in the global namespace, with respect to security policies. You must create security policies on

the primary fabric as security policies are propagated to secondary fabrics or member fabrics by way of mirroring of the security policies. The mirroring takes place every 15 minutes.

**!** **IMPORTANT:** If a primary fabric needs to go offline or fails, set one of the member fabrics as the new primary fabric.

Use the following command to set a member fabric as the primary fabric.

```
maprcli clustergroup updateprimary -clustername <specify actual fabric name here>
```

Once a new member fabric is converted into a primary cluster, it automatically becomes the global policy master.

To identify if a fabric is the global policy master, run the following command on the command line for the fabric:

```
maprcli config load --keys "cldb.pbs.global.master"
```

The value of `cldb.pbs.global.master` is 1 in the output of the aforementioned command, for the fabric that is designated as the global policy master. For a member fabric, the value is 0.

## 2 - Create and Update Security Policies

You can create and update security policies on the global policy master only. You cannot create or modify security policies on member fabrics. See [Creating a Security Policy](#) on page 222 for instructions to create a security policy.

The following table lists the operations you can and cannot perform on the global policy master and member security-policy fabrics:

Security-policy fabric type	Allowed operations	Prohibited operations
Master (fabric set as the global policy master)	<ul style="list-style-type: none"> <li>• Create</li> <li>• Modify</li> <li>• Export</li> <li>• View</li> <li>• Tagging</li> </ul>	<ul style="list-style-type: none"> <li>• Import</li> </ul>
Member	<ul style="list-style-type: none"> <li>• Import</li> <li>• Export</li> <li>• View</li> <li>• Tagging</li> </ul>	<ul style="list-style-type: none"> <li>• Create</li> <li>• Modify</li> </ul>

## 3 – Propagate Security Policies

Once you create security policies on the global policy master, the policy metadata is automatically mirrored on to the other fabrics that are members of the global namespace.

## Data Movement Considerations

The policy server in each security-policy fabric manages security policies and composite IDs. A composite ID is a unique, internal integer that maps to a security policy or set of security policies. The policy server stores the mapping in an internal volume named `mapr.pbs.composite`.

When you assign a security policy to a filesystem resource, the composite ID for that security policy is stored with the resource. Storing the composite ID with the resource instead of the security policy itself optimizes storage. For example, if a policy named HIPAA maps to composite ID 200, this composite ID is stored with any file you tag with HIPAA.

Security policies are shared across the security policy domain, but composite IDs are not. The same security policy on fabricA will have a different composite ID on fabricB and fabricC, as shown in the following table:

Fabric Name	Security Policy	Fabric ID
fabricA	HIPAA	200
fabricB	HIPAA	500
fabricC	HIPAA	800

By default, up to one million composite IDs can be created instantly after which there is a throttle process in place. The default limit of one million composite IDs is sufficient for about one thousand security policies. Using security policies as intended should not trigger the throttle process. However, using security policies for general tagging purposes can quickly exhaust composite IDs and trigger throttling.

### Important Notes About Composite IDs

- You cannot see or interact with composite IDs. However, if you copy a file from one fabric to another, only the data is copied. The policy server on the destination fabric does not recognize the composite ID associated with the file and therefore cannot enforce the access controls configured in the policy. To avoid this issue, use mirroring to synchronize data. During mirroring, security policies are propagated to the destination fabric. The policy server on the destination fabric assigns new composite IDs to the security policies before data synchronization starts. The composite ID/security policy mappings are present when data synchronizes.
- Do not schedule mirroring for the composite ID internal volume `mapr.pbs.composite`.
- Composite IDs are only used with filesystem resources. The database stores policies as an array of policy IDs in the key-value store. The database policy IDs are unique across the global policy domain, which simplifies table replication. For example, policy IDs in JSON tables can be copied from one fabric to another. The server deals with the policy ID, not the policy name. Policy IDs are evaluated and translated to the policy name on the client side.

## Security Policy Enforcement Process

Describes the steps followed during security policy enforcement on volumes.

### Order of Enforcement

Data Fabric File System enforce security policies hierarchically, starting at the volume level.

If the volume-level enforcement mode is set to `PolicyAceAndDataAce` (default setting), the system evaluates and enforces the ACEs directly applied to data objects AND the ACEs defined in the security policies applied to data objects. When a user submits a data-operation request, the system evaluates and enforces the ACEs hierarchically, starting with the volume in which the data resides.

For example, to perform a write operation on a file, the system first evaluates permissions on the volume in which the file resides. If at least one security policy is applied to the volume, the system evaluates the ACEs set in the security policy AND the ACEs or POSIX mode bits directly applied to the volume. Both

sets of ACEs must allow the user to access the volume. If one set of ACEs does not permit access to the volume, the system denies the user permission to perform the operation. If both sets of ACEs permit access to the volume, the system checks access permissions on the file. The system evaluates security policies applied to the file AND any ACEs or POSIX mode bits applied directly to the file. Both sets of ACEs must permit the user write access on the file. If they both allow access (`writelfileeace`), the user can perform the data operation on the file. If not, the system denies access.

Note the following behaviors related to the enforcement mode setting:

- When set to `PolicyAceOnly`, the system only enforces the ACEs set in security policies. A user can only perform data operations on a data object if the security policies associated with the data object allow the user access. However, if a data object is not associated with at least one security policy, the system enforces any ACEs or POSIX mode bits set directly on the data object. In this case, a user can only access the data object if the ACEs or POSIX mode bits set directly on the data object allow the user access.
- In `PolicyAceOnly` and `PolicyAceAndDataAce` modes, if a security policy is applied to a data object, and ACEs are not defined in the policy (" "), the system continues to the next level data object to evaluate permissions.

### Data Fabric File System Enforcement Process

The Data Fabric filesystem enforces security policies on data objects, in the following order:

- Volumes
- Files/Directories

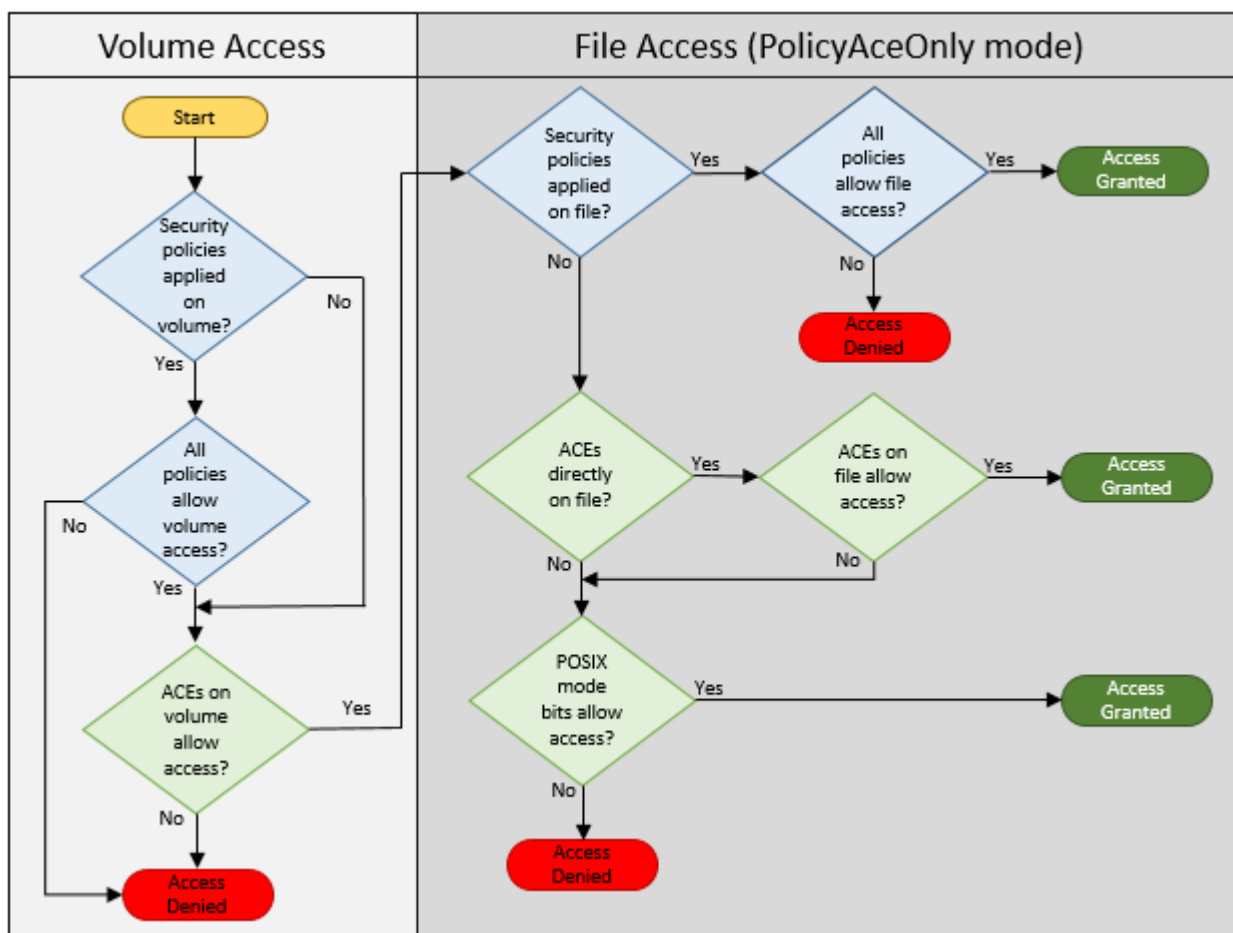




**NOTE:** The system only enforces directory ACEs when determining access to the directory during directory operations. For read and write operations, directory ACEs are enforced during the path-walk operation when opening a file. If the user has a handle (FID) to the file, the user can access the file directly with the FID. In that case, the system ignores directory ACEs. See [Managing File and Directory ACEs](#) on page 217 for details on directory ACEs.

The following diagram shows the order in which the Data Fabric filesystem evaluates and enforces data operations on data objects when the enforcement mode is set to `PolicyAceOnly`:

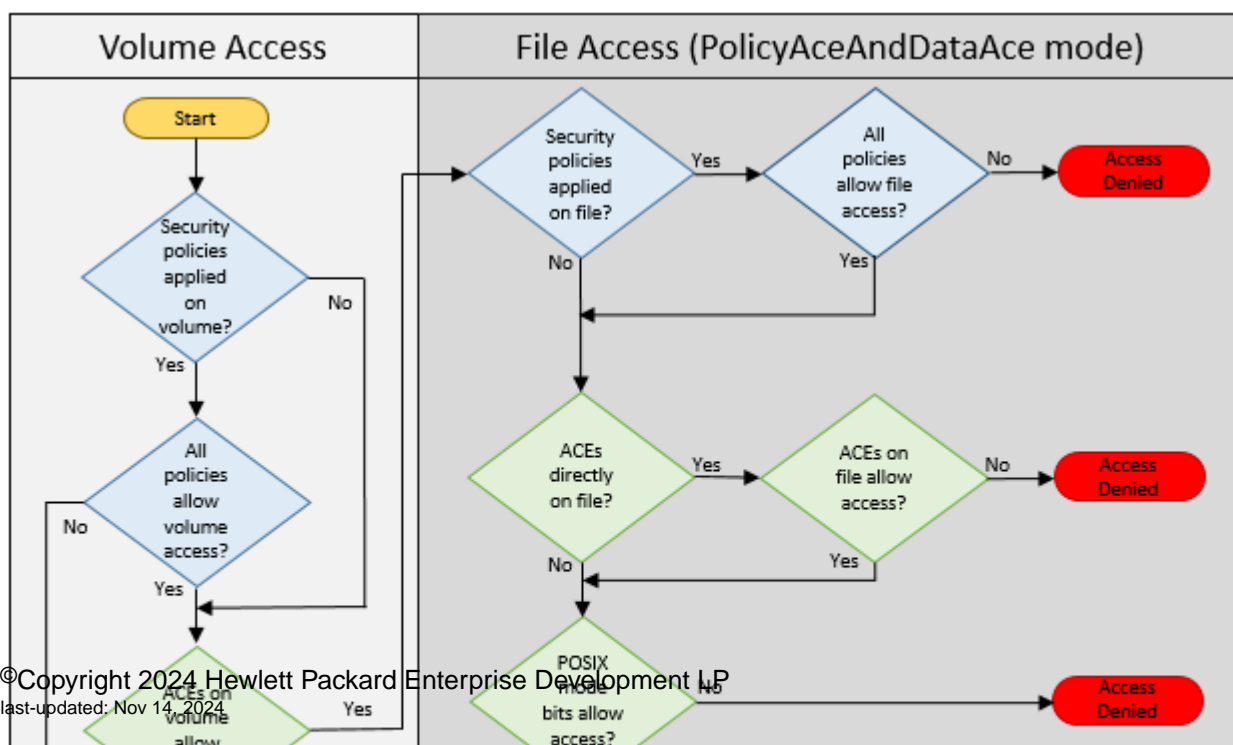


**NOTE:** If no policy is applied at the volume or file/directory level, the system will enforce DataAces (mode and ACEs applied directly on data object) to protect the data.



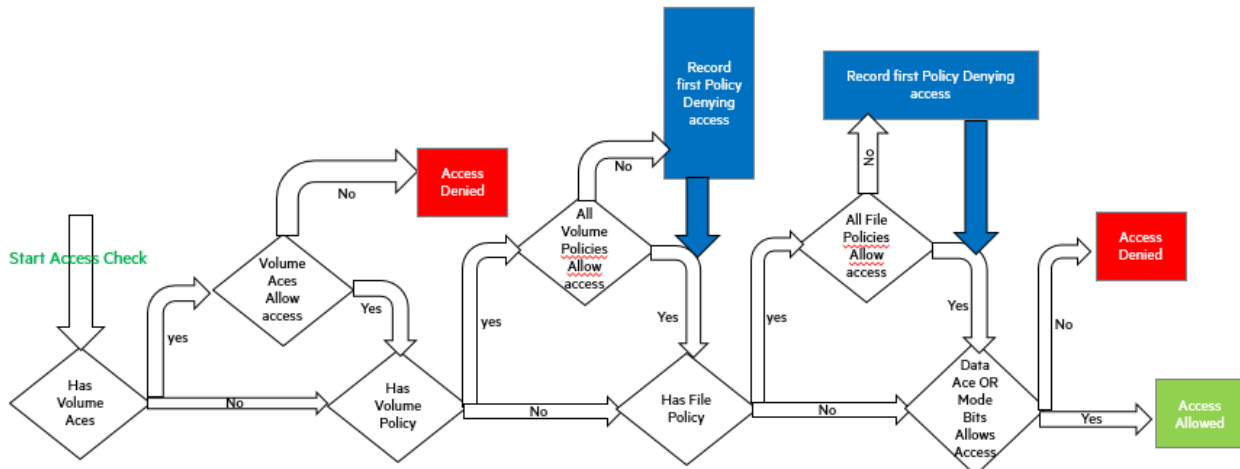
-  Indicates system evaluation and enforcement of ACEs defined in the security policies applied to data objects
-  Indicates system evaluation and enforcement of ACEs defined directly on data objects

The following diagram shows the order in which the Data Fabric file system evaluates and enforces data operations on data objects when the enforcement mode is set to `PolicyAceAndDataAce` (default mode):



The following diagram shows the order in which the Data Fabric file system evaluates and audits data operations on data objects when the enforcement mode is set to `PolicyAceAuditAndDataAce` (permissive mode):

**NOTE:** The system does not enforce denied-access checks, but does log the information about the denied check in the audit logs.



### Understanding Access Control in a Security Policy

The implications of permissions assigned to users and groups in a security policy.

The following types of access can be granted to all (Public) or specific users or groups:

Entity	Permission
Directories	<ul style="list-style-type: none"> <li>• <b>Read</b> the contents of a directory. If you do not select this permission, mode bits are used to determine read access. To read the contents of a directory that is tagged with this security policy, the user must also have read permissions on the volume, the parent directory (if any), and the file.</li> <li>• <b>Lookup</b> or list the contents in a directory. If you do not select this permission, mode bits are used to determine lookup access. To lookup a file of directory that is tagged with this security policy, the user must also have read permissions on the volume and the lookup permission on the directory.</li> <li>• <b>List</b> the contents of a directory. If you do not select this permission, mode bits are used to determine <b>directorylist</b> access. To <b>list</b> the contents of a directory that is tagged with this security policy, the user must also have read permissions on the volume, and lookup permission on all directories in the path (if any).</li> <li>• <b>Add</b> a file or subdirectory. If you do not select this permission, mode bits are used to determine permissions to create files or subdirectories. To add a child to a directory that is tagged with this security policy, the user must also have write permissions on the volume and the parent directory, add child permission on the parent directory, and read and execute permissions on all directories in the path.</li> <li>• <b>Delete</b> a file or subdirectory. If you do not select this permission, mode bits are used to determine permissions to create files and/or subdirectories. To delete a child of a directory that is tagged with this security policy, the user must also have write permissions on the volume and delete child permission on the parent directory, and lookup permissions on all directories in the path.</li> </ul> <p>For more information, see <a href="#">Managing File and Directory ACEs</a> on page 217.</p>
Files	<ul style="list-style-type: none"> <li>• <b>Read</b> a file. If you do not select this permission, mode bits are used to determine read access to file. To read a file that is tagged with this security policy, the user must also have read permissions on the volume, and lookup permission on all directories in path.</li> <li>• <b>Write</b> to a file. If you do not select this permission, mode bits are used to determine read access to the file. To write to a file that is tagged with this security policy, the user must also have write permissions on the volume, and lookup permission on all directories in the path.</li> <li>• <b>Execute</b> a file. If you do not select this permission, mode bits are used to determine execute access to the file. To execute a file that is tagged with this security policy, the user must also have read permissions on the volume, and lookup permission on all directories in the path.</li> </ul> <p>For more information, see <a href="#">Managing File and Directory ACEs</a> on page 217.</p>



## Managing File and Directory ACEs

Describes the implications of setting access control expressions on files and directories.

A file [Access Control Expression \(ACE\)](#) allows you to define access (allowlist and denylist) to files and directories for a combination of users, groups, and roles. If ACEs are not set, POSIX mode bits for the file or directory are used to grant or deny access to the file or directory.

When you set ACEs, Data Fabric sets or resets the corresponding POSIX mode bits to match the permissions granted through ACEs.

- If both ACEs and POSIX mode bits are set, access is granted if access is allowed through ACEs or POSIX mode bits.
- If ACEs are not set, POSIX mode bits are used to grant access.
- If neither ACEs nor POSIX mode bits are set, access is denied.

The owner of the file or directory can set, modify, and remove ACEs for that file or directory using `hadoop mfs` commands.

### File ACEs

You can set and modify permissions to read, write, and execute files using the `hadoop mfs` command or the File ACE Java APIs and C ACE APIs. Specifically, the following access types are supported.

Access Type		Description
Command Line	Java API (Enum)	
<code>-readfile</code>	READFILE	Read a file.
<code>-writefile</code>	WRITEFILE	Write to a file.
<code>-executefile</code>	EXECUTEFILE	Execute a file.

### Directory ACEs

You can set the same ACEs on directories as for files. In addition, directory ACEs support permissions to list, add child, delete child, and lookup directories using `hadoop mfs` command. Specifically, the following access types are supported.

Access Type		Description
Command Line	Java API (Enum)	
<code>-readfile</code>	READFILE	Read a file.
<code>-writefile</code>	WRITEFILE	Write to a file.
<code>-executefile</code>	EXECUTEFILE	Execute a file.
<code>-readdir</code>	READDIR	List the contents of a directory. This access is required to write and/or execute files in the directory.
<code>-lookupdir</code>	LOOKUPDIR	Lookup a file in a directory. This access is required to find, read, write, and/or execute files in the directory.
<code>-addchild</code>	ADDCHILD	Add a file or subdirectory.
<code>-deletechild</code>	DELETECHILD	Delete a file or subdirectory.

Although you can set both file and directory ACEs on directories, only the directory ACEs are used for determining access to the directory. The file ACE on the directory is used as the default ACE setting for new files under that directory.

By default, when you set ACEs on a parent directory:

- Permissions for existing files and subdirectories under that parent remain unchanged.
- New files under that parent inherit the file ACEs and corresponding POSIX mode bits of the parent directory, if available. Otherwise, new files get the default ACE, the empty string (""), which indicates that no one has permissions to read, write, or execute the file. POSIX mode bits are set on the file in the traditional way.
- New subdirectories under the parent inherit both the directory and file ACEs and corresponding POSIX mode bits from the parent directory.



**NOTE:** When accessing files and directories, the ACEs on files have no effect on accessing the parent directory.

### Workaround for Execute Operation when ACES are set on an executable file

When ACEs are set on any file, mode bits are cleared. For a binary to execute, the kernel checks whether the execute bit is set or not, and restricts execution if it is not set. To run an executable file with ACEs set on it, use one of the following workarounds:

1. Set owner mode exec bit on binaries/shell scripts.
2. Set group mode exec bit on binaries/shell scripts.
3. Change owning group for the files to the group used in MapRAces, and set the executable group mode bit.

## Security Policy Permissions

Permissions define which administrative users can create, view, and modify security policies. Administrators set the permissions on security policies through cluster-level and security policy-level ACLs.

### Permission Levels

Policy-based security supports cluster-level and policy-level permissions.

The following table describes the two permission levels:

Permission Level	Description
Cluster-level	<ul style="list-style-type: none"> <li>• Controls which administrators can create and view security policies in a cluster.</li> <li>• Administrators with cluster-level <code>cp</code> permission can create security policies.</li> <li>• Administrators with cluster-level <code>fc</code> permission can view all the security policies created.</li> </ul>
Policy-level	<ul style="list-style-type: none"> <li>• Controls which administrators can view and modify security policies.</li> <li>• Policy-level permissions are set on a per-policy basis.</li> <li>• Permissions set on one security policy do not apply to other security policies.</li> </ul>

Administrators with cluster-level permissions can set cluster-level and security policy-level permissions through any of the following tools:

- Data Fabric UI
- `maprcli acl set|edit` commands

- `maprcli security policy create` commands

**IMPORTANT:** Note these important considerations for security-policy permissions:

- On a fresh cluster install, the `root` user and the Data Fabric user (typically named `mapr` or `hadoop` on each node) have `cp` permission. On an upgraded cluster, only the Data Fabric user has `cp` permission.
- As the cluster owner, the Data Fabric user (typically named `mapr` or `hadoop` on each node), has overriding permission on security policies, including the administrative ACLs. The Data Fabric user can create, view, and modify security policies, regardless of the cluster-level and policy-level permission specified.
- By default, [administrators](#) do not have permission to create security policies. Administrators need cluster-level `cp` (`create security policy`) permission to create security policies. Administrators with cluster-level `a` (`admin`) permission can grant `cp` permission to themselves or other administrators.

**TIP:** You must designate a cluster as the global policy master before you create security policies. Setting a global policy master creates a global namespace for security policies. See [Designating a Fabric as Global Policy Master](#) on page 222.

- Any user with a valid Data Fabric ticket can view security policy IDs and names. This allows non-administrative users to determine which security policies to apply to data objects.

### Permission Codes


Cluster-level and security policy-level permission codes that are set through ACLs grant security policy access to administrators. An administrator (with cluster-level `a` (`admin`) and `cp` (`create security policy`) permissions) that creates a security policy has full control over the security policy unless they specifically grant other administrators access to the security policy through policy-level permissions.

The following sections describe the cluster-level and policy-level permission codes for security policy access:

#### Cluster-Level Permission Codes

The following table lists some cluster-level permission codes and how they relate to security policies.

Cluster-level permission code	Description
<code>a</code> ( <code>admin</code> )	<ul style="list-style-type: none"> <li>• Grants administrative access to cluster ACLs.</li> <li>• Can grant <code>create security policy</code> (<code>cp</code>) permission to themselves or other administrators.</li> <li>• Cannot view or edit the details of any security policy created by other admins. Can only view the security policy ID and name.</li> <li>• Needs security policy-level permissions to view or edit security policies created by other admins.</li> </ul>

Cluster-level permission code	Description
cp (create security policy)	<p> <b>ATTENTION:</b> Administrators need this permission to create security policies.</p> <ul style="list-style-type: none"> <li>Administrators with a (admin) cluster-level permission can grant cp permission to themselves or other administrators.</li> <li>Administrators can view and edit all parts of the security policies they create, including the ACEs and permissions on the security policies.</li> <li>Grants the administrator that creates a security policy the following security <i>policy-level</i> permissions on the security policy:                             <ul style="list-style-type: none"> <li>Full Control (fc)</li> <li>Admin (a)</li> <li>Read (r)</li> </ul> </li> <li>Administrators who create security policies can override their access to the security policies by designating policy owners who can then manage the security policies.</li> </ul>
fc (full control)	<ul style="list-style-type: none"> <li>Grants full control over the cluster and enables all cluster-level administrative options.</li> <li>Cannot change the cluster-level ACLs.</li> <li>Can view all security policies.</li> <li>Cannot create security policies.</li> <li>Cannot edit the details of any security policy unless specifically granted access to a security through policy-level permissions.</li> </ul>

**Policy-Level Permission Codes**

Separate read (r) and edit (fc) permissions for policy owners allow some policy owners to view policy information while others can edit policy information. This allows most administrators to administer the system without seeing the data and also prevents some policy owners from adding their credentials to the administrative ACLs to manipulate the data access ACEs.

Policy-level permissions are set on a per-policy basis. Permissions set on one security policy do not apply to other security policies.

The following table lists the policy-level permission codes needed to perform actions on security policies.

Policy-level permission code	Description
a (admin)	<ul style="list-style-type: none"> <li>Can view and modify permissions on the security policy.</li> <li>Cannot view or modify the security policy; can only view the security policy name and ID.</li> </ul>
fc (full control)	<ul style="list-style-type: none"> <li>Can view and edit any part of the security policy, including the data access ACEs.</li> <li>Cannot view or modify permissions on the security policy.</li> </ul>
r (read)	Can view all parts of a security policy, but cannot modify any part of the security policy.

**Permissions Table**

The following table lists the cluster-level and policy-level permissions needed to perform specific actions on security policies:



**NOTE:** Administrators who create a security policy have policy-level r, a, and fc permission on the security policy.

Action	Cluster-Level	Policy-Level
Create a security policy	cp	--
View details of all security policies	fc	--
View details of a security policy	--	r
View and edit permissions on a security policy (ACLs)	--	a
View and edit the details of a security policy (ACEs, auditing, wire-level encryption)	--	fc

For more information, see

- [Creating a Security Policy](#) on page 222

- [Editing a Security Policy](#) on page 224
- [Enabling a Security Policy](#) on page 228
- [Assigning Multiple Security Policies to One or More Volumes](#) on page 225
- [About ACL](#)

## Designating a Fabric as Global Policy Master

Designate a fabric in the global namespace as the global policy master.

### About this task

A security policy domain is a group of fabrics that share data and use the same security policies to control access to the data. A security policy domain consists of a global policy master and zero or more member fabrics that constitute a global security policy namespace. Before you can create security policies, one fabric must be set as the global policy master, and security policies must be created and managed only on the fabric that is designated as the global policy master.

The primary fabric is auto-designated as the global policy master. Hence, it is not required to explicitly designate a fabric as a global policy master.



**NOTE:** Every 15 minutes, the policies created on the global policy master are mirrored on the member fabrics in the same global namespace.

## Creating a Security Policy

Add a security policy on the global policy master.

### Prerequisites

You must have the permission to create a security policy.

### About this task

Security policies can be created on a fabric that is designated as the global policy master.

See [Designating a Fabric as Global Policy Master](#) on page 222 to designate a fabric as the global policy master.

A security policy is a common set of access permissions on the Data Fabric file system that can be assigned to users and/or groups, or to public (all users).

The following permissions can be assigned on files and directories on the Data Fabric file system.

- Read, write, execute permissions on files
- Read, lookup, add child directory, delete child directory on directories

A security policy can be assigned to volumes when tagging is allowed.

See [Administering Security Policies](#) on page 206 for details on values for the `allow tagging` and `access control` fields for a security policy.



**NOTE:** When you allow tagging, you can assign the security policy to a volume on the fabric.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric manager** from the dropdown on the Home page.

3. Click **Security Administration** seen on the Home page.
4. Click Create Policy on the **Global policies** card.
5. Enter the **Name** of the security policy.
6. Enter the **Description**.
7. Select the option for **Access Control**.
8. Toggle **Allow Tagging** to allow or disallow tagging.
9. Click **Add access permissions** to add access permissions to directories and files for selected users or groups.
10. To grant permission to all users and groups, turn on the **Public** toggle. To grant permissions to specific users or groups, turn off the **Public** toggle, and enter a comma-separated list of users or groups.
11. Select the permissions to be granted on directories and files to the specified users or groups.
12. Click **Add**.
13. Click **Create**.

### Results

The security policy is created and is displayed on the **Global Policies** card for the fabric.

### Related `maprcli` Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `policy create`

## Viewing a Security Policy

View security policy details.

### About this task

You can view security policy details on the Data Fabric UI.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric manager** from the dropdown on the Home page.
3. Click **Security Administration** on the Home page.
4. Scroll down to the **Global policies** card.
5. On the list of policies, click the ellipsis seen under **Actions** for the security policy to edit.
6. Click **View details** to view the security policy details.

### Results

The security policy is displayed on the Data Fabric UI.

### Related `maprcli` Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- [policy info](#)

## Viewing All Security Policies

View all security policies on the Data Fabric UI.

### About this task



**NOTE:** The **View All** button is visible when you have configured five or more security policies on the global policy master.

Use the following steps to view all security policies on the Data Fabric UI.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric manager** from the dropdown on the Home page.
3. Click **Security Administration** on the Home page.
4. Click **View All** on the **Global policies** card.

### Results

All existing security policies are displayed on the Data Fabric UI.

### Related maprcli Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- [policy list](#)

## Editing a Security Policy

Make changes to a security policy.

### Prerequisites

You must be a fabric manager to perform this operation.

### About this task

You can make changes to a security policy for purposes such as deprecating or retiring the security policy, and/or changing access permissions.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric manager** from the dropdown on the Home page.
3. Click **Security Administration** on the Home page.
4. Scroll down to the **Global policies** card.
5. On the list of security policies, click the ellipsis under **Actions** for the security policy to edit.



6. Click **Edit policy** to make changes to the policy.
7. Make the required changes.
8. Click **Save**.

### Results

The changes are saved and applied.

### Related `maprcli` Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `policy modify`

## Assigning a Security Policy to One or More Volumes

### About this task

You can assign security policy to one or more volumes associated with a fabric, via the Data Fabric UI.

Follow the steps given below to assign a security policy to one or more volumes.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric manager** from the dropdown on the Home page.
3. Click **Security Administration** on the Home page.
4. Scroll down to the **Global policies** card.
5. On the list of security policies, click the ellipsis under **Actions** for the security policy to assign to one or more volumes.
6. Click **Assign Policy**.
7. Select the fabric and one or more volumes to assign the policy to.
8. Click **Save**.

### Results

The security policy is assigned to the selected volumes.

### Related `maprcli` Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `policy attach`

## Assigning Multiple Security Policies to One or More Volumes

Describes how to assign multiple security policies to volumes.

### Prerequisites

You must be a fabric manager to perform this operation.

### About this task

You can assign multiple security policies to one or more volumes associated with a fabric at one go, via the Data Fabric UI.

Follow the steps given below to assign a security policy to one or more volumes.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric manager** from the dropdown on the Home page.
3. Click **Security Administration** on the Home page.
4. Click **View All** on the **Global policies** card.
5. Click **Assign Policy**.
6. Search for policies in the search bar and select the policies to apply to a volume or a common set of volumes.
7. Click **Select Resources**.
8. Select the fabric and one or more volumes to assign the selected policies to.
9. Click **Save**.

### Results

The selected security policies are assigned to the selected volumes.

### Related `maprcli` Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `policy attach`

## Unassigning One or More Security Policies from a Volume

Unassign a policy from a volume to which it has been previously assigned.

### Prerequisites

You must be a fabric user to perform this operation.

### About this task

You can unassign a security policy from a volume to which it has been assigned.

Follow the steps given below to unassign one or more security policies from a volume.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** from the dropdown on the Home page.
3. Click the Table view icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the volume to rename.

4. Click the volume name seen under **Resource Name**.
5. Go to **Settings** tab for the volume.
6. Click the edit icon for **Security Policy**.
7. On the **Edit Policy** dialog box, deselect the check boxes for the policy or policies that you wish to unassign. You can use the search bar to search for the required policy if there are multiple policies attached to the volume.
8. Click **Save**.

### Results

The policy is unassigned from the volume.

## Disabling a Security Policy

Describes how to disable a security policy.

### Prerequisites

You must be a fabric manager to perform this operation.

### About this task

You can disable a security policy instead of deleting it completely from Data Fabric. When you disable a security policy, it does not apply to any volume that it has been assigned to. A disabled security policy cannot be assigned to any other volume, unless the policy is enabled again.

Follow the steps given below to disable a security policy.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric manager** from the dropdown on the Home page.
3. Click **Security Administration** on the Home page.
4. Scroll down to the **Global policies** card.
5. On the list of security policies, click the ellipsis under **Actions** for the security policy to disable.
6. Click **Edit policy** to make changes to the policy.
7. Select **Disarmed** from the **Access Control** dropdown.
8. Click **Save**.

### Results

The security policy is disabled. You can enable the security policy that has been disabled.

### Related `maprcli` Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `policy modify`

## Enabling a Security Policy

Describes how to enable a security policy.

### Prerequisites

You must be a fabric manager to perform this operation.

### About this task

You can enable a security policy that has been recently created or has been disabled in the past. Once you enable a security policy, it can be assigned to a fabric resource such as a volume.

Follow the steps given below to enable a security policy.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric manager** from the dropdown on the Home page.
3. Click **Security Administration** seen on the Home page.
4. Scroll down to the **Global policies** card.
5. Click the ellipsis under **Actions** for the security policy to enable.
6. Click **Edit policy** to make changes to the policy.
7. Select **Armed** from the **Access Control** dropdown.
8. Click **Save**.

### Results

The security policy is enabled. An enabled security policy can be assigned to one or more volumes.

### Related `maprcli` Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `policy modify`

## Administering Bucket Policies

---

Describes how to manage the bucket policy associated with a bucket.

A bucket policy specifies domain users and the operations they can perform on buckets. Bucket policies override the default bucket policy inherited from the account.

**TIP:** A bucket policy is an access control mechanism applied to objects stored on an S3 object store associated with Data Fabric, while a security policy is an access control mechanism for data stored on Data Fabric volumes.

Typically, a fabric manager applies policies; however, given the proper permissions, domain and IAM users can also apply policies.

A bucket policy comprises the following elements:

- Effect : Allow or deny permission on a resource.

- **Principal:** The user, group that is allowed or denied resource access.
- **Action:** The operation on the resource that is allowed or denied.
- **Resource:** The bucket resource(s) on which the action is allowed or denied.

You can create bucket policies by using the Data Fabric UI. There are two methods:

- Create or upload JSON from a file . See [Creating a Bucket Policy using JSON](#) on page 229
- Use the policy generator to construct a JSON. See [Creating a Bucket Policy using Policy Generator](#) on page 231

## Creating a Bucket Policy using JSON

Describes how to create a bucket policy that is applicable to all objects present in the bucket.

### About this task

You can add a bucket policy by using the Data Fabric UI.

If you are well-versed with JSON, you can directly write a policy JSON or upload the policy JSON from a file to create a bucket policy.

If you are not well-versed with JSON, you can use the policy generator that will generate the policy in JSON format. See [Creating a Bucket Policy using Policy Generator](#) on page 231.

Follow the steps given below to create a bucket policy by directly uploading a JSON file or creating a JSON file.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** from the dropdown on the Home page.
3. Click the Table view icon on the **Resources** card.
4. In the tabular list of fabrics, click the down arrow for the fabric that contains the bucket to which you wish to upload objects.
5. Click the bucket name seen under **Resource Name**.
6. Navigate to the **Settings** tab.
7. Under **Properties**, click the pencil icon for **Policy**.
8. Scroll down to view **Select File** under the blank text area.
9. Click **Select File** and select a pre-defined JSON file from the path where it is stored.
10. Click **Upload JSON**. The policy is visible in the blank text area or form visible on the **Generate Policy** page. Alternatively, you can also type in a JSON in the blank text area provided on the **Generate Policy** page.
11. Click **Save Policy**.

### Results

The bucket policy is saved and applied to the bucket resources mentioned in the bucket policy. See

[Sample Bucket Policy](#) on page 230 for bucket policy JSON examples.

### Related maprcli Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `mc policy set`

### Sample Bucket Policy

Example for bucket policy.

The following bucket policy allows all users in *group1* to get, put, and delete objects, and list the bucket contents. The `${bucket}` keyword is a placeholder that the system automatically replaces with the bucket name.

```
{
  "Version": "2012-10-17",
  "Id": "PolicyContent1",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": "arn:primary:default:group:group1",
      "Action": ["s3:GetObject", "s3:PutObject", "s3:DeleteObject"],
      "Resource": "arn:aws:s3:::${bucket}/*"
    },
    {
      "Effect": "Allow",
      "Principal": "arn:primary:default:group:group1",
      "Action": ["s3:ListBucket"],
      "Resource": "arn:aws:s3:::${bucket}"
    }
  ]
}
```

The following policy allows all users in *group1* to get, put, and delete objects, and list the bucket contents while also denying *user1* and *user2* in *qagroup1* permission to perform get, put, and delete operations.

```
{
  "Version": "2012-10-17",
  "Id": "PolicyContent1",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": "arn:primary:default:group:group1",
      "Action": ["s3:GetObject", "s3:PutObject", "s3:DeleteObject"],
      "Resource": "arn:aws:s3:::${bucket}/*"
    },
    {
      "Effect": "Deny",
      "Principal": {
        "AWS": [
          "arn:primary:default:user:user1",
          "arn:primary:default:user:user2"
        ]
      },
      "Action": ["s3:GetObject", "s3:PutObject", "s3:DeleteObject"],
      "Resource": "arn:aws:s3:::${bucket}/*"
    }
  ]
}
```

```

        "Principal": "arn:primary:default:group:group1",
        "Action": ["s3:ListBucket"],
        "Resource": "arn:aws:s3:::${bucket}"
    }
]
}

```

The following policy allows *user1* to perform all the specified operations:

```

{
  "ID": "PolicyContent1",
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": "arn:primary:default:user:user1",
      "Action": [
        "s3:GetObjectRetention",
        "s3:GetObjectTagging",
        "s3>DeleteObjectTagging",
        "s3>DeleteObjectVersionTagging",
        "s3:GetObject",
        "s3:GetObjectLegalHold",
        "s3:PutObject",
        "s3:PutObjectLegalHold",
        "s3:PutObjectRetention",
        "s3:PutObjectTagging",
        "s3>DeleteObject"
      ],
      "Resource": "arn:aws:s3:::${bucket}/*"
    },
    {
      "Effect": "Allow",
      "Principal": "arn:primary:default:user:user1",
      "Action": [
        "s3>DeleteBucket",
        "s3>DeleteBucketPolicy",
        "s3:GetBucketPolicy",
        "s3:GetBucketTagging",
        "s3:ListBucket",
        "s3:PutBucketPolicy",
        "s3:PutBucketTagging"
      ],
      "Resource": "arn:aws:s3:::${bucket}"
    }
  ]
}

```

## Creating a Bucket Policy using Policy Generator

Describes how to create bucket policy using policy generator.

### About this task

You can add multiple statements to the policy. On saving the policy, a JSON is constructed by the policy generator.

Follow the steps given below to create a bucket policy by using the policy generator.

### Procedure

1. Log on to the Data Fabric UI.

2. Select **Fabric user** from the dropdown on the Home page.
3. Click the Table view icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the bucket to which you wish to upload objects.
4. Click the bucket name seen under **Resource Name**.
5. Navigate to the **Settings** tab.
6. Under **Properties**, click the pencil icon for **Policy**.
7. Click **Switch to form**.
8. Enter the policy name.
9. Select the **Effect**.
10. Select the **Actions** that are allowed or denied to the resource.
11. Enter the users or groups separated by commas in **Principal(s)**.
12. Enter one or more bucket names separated by commas in **Resource(s)**.
13. You can add multiple such statements by clicking the **Add** seen on the right of **Statements**.
14. Once you are done adding the statements, you can **Review policy** to review the auto-generated JSON.
15. Make changes to the JSON, if required. This is an optional step.
16. Click **Save Policy**.

## Results

The bucket policy is saved. See [Sample Bucket Policy using Policy Builder](#) on page 232 for an example.

## Related maprcli Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `mc policy set`

## Sample Bucket Policy using Policy Builder

The following screenshots represent the options to select to create the bucket policy (see JSON below) mentioned in this example.



**Generate policy**

Step 1 of 2

## Policy details

 **Switch to JSON**

Policy name\*

sample-bucket-policy

**Statements**

**+ Add**

Effect

**Statement 1**

Allow  
 Deny

**Statements**

**+ Add**

Effect

**Statement 1**

Allow  
 Deny

Action(s)\*

52 items

- s3:CreateBucket
- s3>DeleteBucket
- s3>DeleteBucketPolicy
- s3>DeleteObject
- s3>DeleteObjectTagging
- s3>DeleteObjectVersion

Principals(s)\*

Use a comma to separate multiple principals.

arn:primary:default:group:group1

Resources(s)\*

Use a comma to separate multiple resources.

arn:aws:s3:::\${bucket}/\*

sample-bucket-policy

**Statements** + Add

Statement 1 🗑️

Statement 2 🗑️

**Effect**

Allow

Deny

**Action(s)\***

52 items

- s3:AbortMultipartUpload
- s3:BypassGovernanceRetention
- s3:CreateBucket
- s3>DeleteBucket
- s3>DeleteBucketPolicy

**Principals(s)\***

Use a comma to separate multiple principals.

arn:primary:default:group:group1

**Resources(s)\***

Use a comma to separate multiple resources.

arn:aws:s3:::\${bucket}/\*

```
{
  "Id": "sample-bucket-policy",
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Statement1713336710809",
      "Principal": {
        "AWS": [
          "arn:primary:default:group:group1"
        ]
      },
      "Action": [
        "s3:CreateBucket",
        "s3>DeleteBucket",
        "s3>DeleteBucketPolicy",
        "s3>DeleteObject",
        "s3>DeleteObjectTagging"
      ],
      "Resource": [
        "arn:aws:s3:::${bucket}/*"
      ],
      "Effect": "Allow"
    }
  ]
}
```

```

    },
    {
      "Sid": "Statement1713337230508",
      "Principal": {
        "AWS": [
          "arn:primary:default:group:group1"
        ]
      },
      "Action": [
        "s3:AbortMultipartUpload"
      ],
      "Resource": [
        "arn:aws:s3:::${bucket}/*"
      ],
      "Effect": "Deny"
    }
  ]
}

```

The aforementioned bucket policy allows all users in *group1* to create bucket, delete bucket, delete bucket policy, delete objects in the bucket, and delete objects, and delete object tagging. The policy disallows or denies permission to abort a multi-part upload of file to the bucket. The `${bucket}` keyword is a placeholder that the system automatically replaces with the bucket name.

## Editing a Bucket Policy

Describes how to make changes to the bucket policy associated with a bucket.

### About this task

You can edit a bucket policy.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** from the dropdown on the Home page.
3. Click the Table view icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the bucket for which you wish to revise the bucket policy.
4. Click the bucket name seen under **Resource Name**.
5. Navigate to the **Settings** tab.
6. Under **Properties**, click the pencil icon for **Policy**.
7. Continue to edit the JSON or select a JSON file from the desired location. Alternatively, click **Switch to form** to use the policy generator.
8. Make the required changes and review the policy.
9. Click **Save Policy**.

### Results

The changes to the bucket policy are saved and applied to the bucket resources mentioned in the bucket policy.

## Delete a Bucket Policy from a Bucket

Describes how to delete a bucket policy from a bucket.

### About this task

You can delete a bucket policy by using the Data Fabric UI.

Follow the steps given below to delete a bucket policy.

### Procedure

1. Log on to the Data Fabric UI.
2. Select **Fabric user** from the dropdown on the Home page.
3. Click the Table view icon on the **Resources** card. In the tabular list of fabrics, click the down arrow for the fabric that contains the bucket to which you wish to upload objects.
4. Click the bucket name seen under **Resource Name**.
5. Navigate to the **Settings** tab.
6. Under **Properties**, click the bin icon for **Policy**.
7. Click **Delete** when asked to confirm the deletion operation.

### Results

The bucket policy is deleted from the bucket.

## Working with an External NFS Server

---

Associate an external NFS server with Data Fabric to share data across clusters in the global namespace.

An external NFSv4 server can be used to share data across clusters in the global namespace.

You can import an external NFSv4 server into Data Fabric global namespace, for a NFSv4 client to be able to access the data present on such external NFSv4 server.

## Importing an External Network File System Server

Import an external NFS server into Data Fabric to be able to transfer data from Data Fabric to the external NFS server to make it shareable across the clusters in the global namespace or cluster group.

### About this task

You can import an external network file system (NFS) into the global namespace so that the external NFS server is available and accessible to all the clusters in the global namespace.

After you import an external NFS server, you are able to transfer data from the Data Fabric cluster on to the external NFS server. The data, thus, transferred is shareable between all clusters present in the global namespace.



**NOTE:** NFSv4 compliant servers can be imported into Data Fabric.

When you are importing an external NFS server, you can specify one of more hostnames or IP addresses that are assigned to the NFS server. If multiple network interface controllers are attached to the NFS server, the NFS server is identified by multiple IP addresses or hostnames.

Follow the steps given below to import an external NFS server into Data Fabric.

**Procedure**

1. Log on to the Data Fabric UI .
2. Select **Fabric manager** on the Home page.
3. Click **Global namespace**.
4. Click **Import External NFS**.
5. Enter the name for the NFS server in **NFS name**.
6. Enter the IP addresses or the hostnames for the external NFS server as a comma-separated string in **IP address or hostname**.
7. Click **Import**.

**Results**

The NFS server is imported into Data Fabric. The NFS server is visible under the list of resources in the global namespace.

You are able to transfer data to the NFS server after importing the NFS server.

**Related maprcli Commands**

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `clustergroup addexternal`

**Viewing the IP Address/Hostname for External NFS Server**

View the IP address or hostname associated with an external NFS server on the Data Fabric UI.

**About this task**

You can view the IP addresses or hostnames that are associated with an external NFS server.

Follow the steps given below to view the list of IP address(es) or hostname(s) associated with an external NFS server.

**Procedure**

1. Log on to the Data Fabric UI .
2. Select **Fabric manager** on the Home page.
3. Click **Global namespace**.
4. On the Table view, click the ellipsis under **Actions** for the external NFS server.
5. Click the **View IP Addresses/hostnames** option under **Action**.

**Results**

The IP addresses or hostnames associated with the external NFS server are displayed.

**Deleting an External NFS Server**

Delete an external NFS server association with Data Fabric.

### About this task

You can delete the association of an external NFS server with Data Fabric via the Data Fabric UI. Follow the steps given below to delete an external NFS server from Data Fabric.

### Procedure

1. Log on to the Data Fabric UI .
2. Select **Fabric manager** on the Home page.
3. Click **Global namespace**.
4. On the Table view, click the icon under **Action** for the external NFS server.
5. Click the **Delete** option under **Action**.
6. Click **Delete** on the confirmation message that appears, to confirm deletion.

### Results

The external NFS server entry is deleted from the Data Fabric UI, and the association between the external NFS server and Data Fabric is removed. You are unable to access or transfer data on the external NFS server via Data Fabric.



**NOTE:** If you wish to access the data on the deleted NFS server, you must import the NFS server into Data Fabric again.

### Related maprcli Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `clustergroup remove cluster`

## Working with an External S3 Object Store

---

A fabric manager and fabric user can import an external S3 object store into the Data Fabric UI.

You can import a third-party S3 object store into the global namespace, to be able to view data from the third-party S3 object store via the Data Fabric UI. Data Fabric supports import of Amazon S3, GCP, VAST, Scality, WEKA, and other such S3-compliant object stores.

See [S3 Federation in Global Namespace](#) on page 89 for more information.

To copy data from fabrics on the global namespace to the external S3 object store, download the data to local folder and then upload to the external S3 object store.

### Importing an External S3 Object Store

Describes how to import an external S3 object store into the global namespace.

#### Prerequisites

If Data Fabric accesses the internet via a proxy server, do the following.

1. Specify the proxy settings in `/opt/mapr/initscripts/mapr-s3server`:

```
export HTTPS_PROXY="http://<proxy server FQDN or proxy server IP>:8080"
export http_proxy="http://<proxy server FQDN or proxy server IP>:8080"
export HTTP_PROXY="http://<proxy server FQDN or proxy server IP>:8080"
export https_proxy="http://<proxy server FQDN or proxy server IP>:8080"
```

2. Run the following command to restart the related Data Fabric component:

```
maprcli node services -name s3server -action restart -nodes $
(hostname -f) -json
```

The Data Fabric is now able to access the internet via the proxy, and able to connect to an external S3 object store.

### About this task

A fabric manager or a fabric user can import an external S3 object store into the global namespace to transfer data from the Data Fabric to the external S3 object store. You can import AWS S3, Google Cloud Platform (GCP), WEKA, Scality, VAST and other S3-compliant object stores into a global namespace to consolidate your data across external S3 object stores on Data Fabric.

Use the following steps to import an external S3 object store into the global namespace by using the Data Fabric UI.

### Procedure

1. Log on to the Data Fabric UI.
2. If you are a fabric manager, select **Fabric manager** on the home page. Skip this step if you are a fabric user.
3. If you are a fabric manager, click **Global namespace**. Optionally, you can switch to the fabric user view, where the **Import External S3** button appears in the **Resources** section. If you are a fabric user, you can access the the **Import External S3** button only in the **Resources** section.
4. Click **Import External S3**.
5. Enter the name for the S3 object store in **Name**.
6. Enter the **S3 vendor type**, selecting from one of these values:
  - **AWS**
  - **GCP**
  - **Generic** (for WEKA, Scality, VAST, and other S3-compliant object stores)
7. Depending on the vendor type you selected, fill in the remaining values by consulting the following tables. An asterisk (\*) indicates a required field:
  - For **AWS**:

Parameter	Description
Name*	Name of the S3 object store.
S3 Vendor	Choose <b>AWS</b> .
Region*	The AWS region.

Parameter	Description
Access type	For the AWS vendor type, you can select from one of the following: <ul style="list-style-type: none"> <li>Access Credentials</li> <li>Secure Token Service</li> </ul>
Access Credentials	Selecting this value requires you to specify an access key and secret for access to the S3 object store. See the descriptions of these keys later in this table.
Secure Token Service	Selecting this value requires you to specify a Web identity role ARN. To configure the ARN, see <a href="#">Configuring STS for Data Fabric</a> on page 245.  STS is an access method that provides an alternative to the traditional access key and secret key. For more information, see <a href="#">Integrating the AWS Security Token Service (STS) with Data Fabric</a> on page 243.
Access key*	A long-term credential for an Amazon user. The key enables access to S3 resources for all fabrics in the global namespace. For more information, see <a href="#">Managing access keys for IAM users</a> in the Amazon documentation.
Secret key*	A long-term credential for an Amazon user. The key enables access to S3 resources for all fabrics in the global namespace. For more information, see <a href="#">Managing access keys for IAM users</a> in the Amazon documentation.
Web identity role ARN*	An Amazon resource name (ARN) that enables STS authentication. See <a href="#">Configuring STS for Data Fabric</a> on page 245.

- For **GCP**:

Parameter	Description
Name*	Name of the S3 object store.
S3 vendor	Choose <b>GCP</b> .
Region*	The GCP region.
Access key*	A key that enables access to S3 resources for all fabrics in the global namespace.
Secret key*	A key that enables access to S3 resources for all fabrics in the global namespace.

- For **Generic**:

Parameter	Description
Name*	Name of the S3 object store.
S3 vendor	Choose <b>Generic</b> .
Access key*	A key that enables access to S3 resources for all fabrics in the global namespace.
Secret key*	A key that enables access to S3 resources for all fabrics in the global namespace.
Hostname / IP Address*	Enter the host names or IP addresses of the external S3 object store as a comma-separated list.
S3 server port	The server port. The default value is 9000.
Use TLS encryption	TLS encryption enables communication over a secure connection. TLS is enabled by default.
S3 server certificate	If the generic S3 object store is not CA certified, you must drag and drop the S3 server certificate into this box to enable secure communication. If the S3 server is CA certified, the certificate is not required.

## 8. Click **Import**.



## Results

The S3 object store is imported into your global namespace, and is visible as part of the global namespace. The S3 object store is visible under the list of resources in the global namespace or list of resources with your username as the owner.

## Related `maprcli` Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `clustergroup addexternal`

## Deleting an External S3 Object Store

Delete reference to an external S3 object store that has been imported into the global namespace.

### Prerequisites

You must be a fabric manager or a fabric user to perform this operation.

If you are a fabric user, you can delete an external S3 object store if you are the owner of the imported S3 object store.

### About this task

As a fabric manager, you can remove the reference to an external S3 object store from the global namespace. Once an imported external S3 object store is deleted, you are unable to access any objects that are present on the external S3 object store. The objects, however, are not deleted from the object store. Only the object store association with the Data Fabric global namespace is removed.

A fabric user that is the owner of an external S3 object store can delete the external S3 object store from the global namespace.

### Procedure

1. Log on to the Data Fabric UI .
2. If you are a fabric manager, click **Global namespace** on the **Fabric manager** view or click **Resources** on the **Fabric user** view. If you are a fabric user, click **Resources**.
3. On the Table View, click the ellipsis under **Actions** for the external S3 server.
4. Click the **Delete** option.
5. Confirm the deletion of the external S3 object store.

## Results

The external S3 object store is removed from the global namespace and the buckets and objects present on the external S3 server are no longer accessible from the global namespace. You are unable to manage the external S3 object store data by using Data Fabric.

## Related `maprcli` Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `clustergroup remove cluster`

## Viewing Object Store Details

Describes how to view the S3 object store details such as the name, S3 vendor and other details on the Data Fabric UI.

### About this task

By using the Data Fabric UI, you can view the S3 vendor, name, region and other details for an external Amazon S3, GCP or a generic object store that has been imported into the global namespace.

### Procedure

1. Log on to the Data Fabric UI .
2. If you are a fabric manager, click **Global namespace** (on the **Fabric manager** view ) or click **Resources** (on the **Fabric user** view). If you are a fabric user, click **Resources**.
3. On the Table View, click the ellipsis under **Actions** for the external S3 server.
4. Click the **Quick view** option.

### Results

The external S3 details are displayed. The owner is the username for the user that imported the external S3 object store.

## Sharing External S3 Object Store with User or Group

Describes how to share S3 object store with other users and groups by using the Data Fabric UI.

### Prerequisites

You must be one of the following users to perform this operation:

- owner of the imported external S3 object store
- fabric manager

### About this task

You can share an imported external object store with other Data Fabric users and groups, so that the users or groups are able to use the data from the external object store.

As a fabric manager or a fabric user, you can share the external S3 object store that has been imported by you into Data Fabric.

Follow the steps given below to share an external S3 object store with other Data Fabric users and/or into Data Fabric.

### Procedure

1. Log on to the Data Fabric UI.
2. If you are a fabric manager, click **Global namespace** (on the **Fabric manager** view ) or click **Resources** (on the **Fabric user** view). If you are a fabric user, click **Resources**.
3. On the Table View, click the ellipsis under **Actions** for the external S3 server row.
4. Click **Share**.
5. Search and click on the users and/or groups with which you wish to share the external S3 object store.

6. Click **Submit**.

### Results

The S3 object store is shared with the selected users and/or groups. The users and/or groups are able to access the data stored on the S3 object store.

## Removing Share for External S3 Object Store

Describes how to remove sharing for an external S3 object store with other users and/or groups, by using the Data Fabric UI.

### Prerequisites

You must be one of the following users to perform this operation:

- the owner of the external S3 object store
- fabric manager

### About this task

As a fabric manager or a fabric user, you can share the external S3 object store that you have imported into Data Fabric.

Follow the steps given below to remove the share on an external S3 object store with other Data Fabric users and/or groups.

### Procedure

1. Log on to the Data Fabric UI.
2. If you are a fabric manager, click **Global namespace** (on the **Fabric manager** view ) or click **Resources** (on the **Fabric user** view). If you are a fabric user, click **Resources**.
3. On the Table View, click the ellipsis under **Actions** for the external S3 server row.
4. Click **Share**.
5. Click the bin icon to remove sharing for the users and/or groups that you wish to stop sharing the external S3 object store with.
6. Click **Submit**.

### Results

The S3 object store share is revoked for the users and/or groups that have been removed. The users and/or groups are unable to access the data stored on the S3 object store.

## Integrating the AWS Security Token Service (STS) with Data Fabric

Describes how the HPE Ezmeral Data Fabric can access AWS services by using the Security Token Service (STS) rather than a secret key and access key.

Data Fabric releases 7.5 and later support importing an external S3 object store into the global namespace. This feature requires the user to provide an access key and secret key to access the external S3 object store.

With release 7.7.0, Data Fabric provides an option for gaining access to AWS S3 object stores. You can import an external AWS S3 server by using the `maprcli clustergroup addexternal` command and

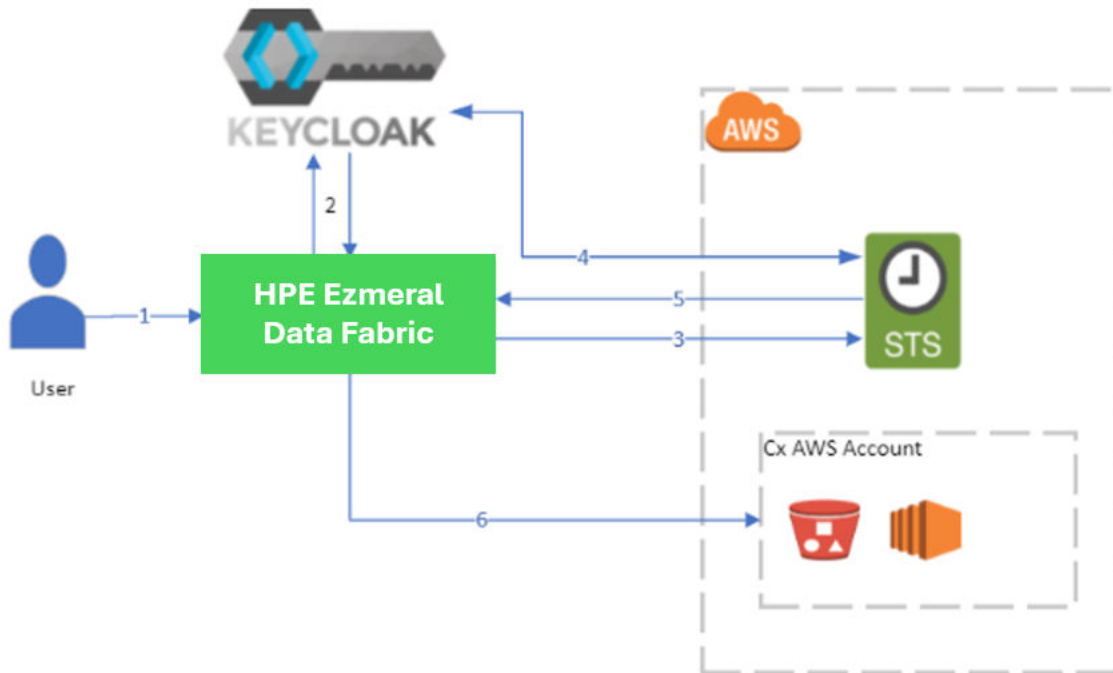
specifying an Amazon Resource Name (ARN) to enable STS authentication. For configuration steps, see [Configuring STS for Data Fabric](#) on page 245.

With release 7.8.0, you can use the Data Fabric UI to enable STS authentication and specify the ARN when you import the external S3 object store. See [Importing an External S3 Object Store](#) on page 238.

Using STS simplifies the process of accessing AWS services by using STS tokens for authentication. With STS tokens, the Data Fabric user can assume an AWS role and get temporary credentials to perform S3 actions. Once the external S3 object store is imported into the global namespace, **all S3 operations** automatically use STS.

### How STS Works with Keycloak and Data Fabric

The following diagram illustrates the authentication flow based on a Keycloak web identity to a user account using AWS STS:



In the diagram:

1. The user logs in to the Data Fabric.
2. Keycloak authenticates the user and generates a JWT token for the user.
3. The Data Fabric requests a temporary access key and secret key for the user using the Keycloak JWT token from STS.
4. STS verifies the token validity.
5. If the token is valid, STS responds with temporary credentials to access the user's AWS account.
6. Data Fabric accesses the user account to perform the infrastructure or S3 actions.

### Limitations for STS Support

Note the following limitations for using STS in the current release:

- The option to use STS when importing an S3 object store is available only for AWS S3 object stores. Non-AWS S3 object stores may not use STS with Data Fabric.

- Enabling STS requires Keycloak to be deployed on a public network IP address so that AWS STS can communicate with Keycloak and verify that the JWT tokens are from the Data Fabric software. If your Keycloak deployment resides on an intranet and is not reachable by a public network, you cannot use STS. However, you can still use the access key and secret key import method.

## Configuring STS for Data Fabric

Describes how to configure the AWS role and enable STS when you import an external S3 object store into the global namespace of the HPE Ezmeral Data Fabric.

Enabling STS is a two-step process:

1. Configure the role in your AWS environment. See [Configuring a Role for Data Fabric in AWS](#) on page 245.
2. Use one of the following procedures to import the external AWS S3 object store:
  - Command line procedure: [Importing the External AWS S3 Object Store by Using the maprcli Command](#) on page 249
  - Data Fabric UI procedure: [Importing an External S3 Object Store](#) on page 238

### Prerequisite for Configuring STS

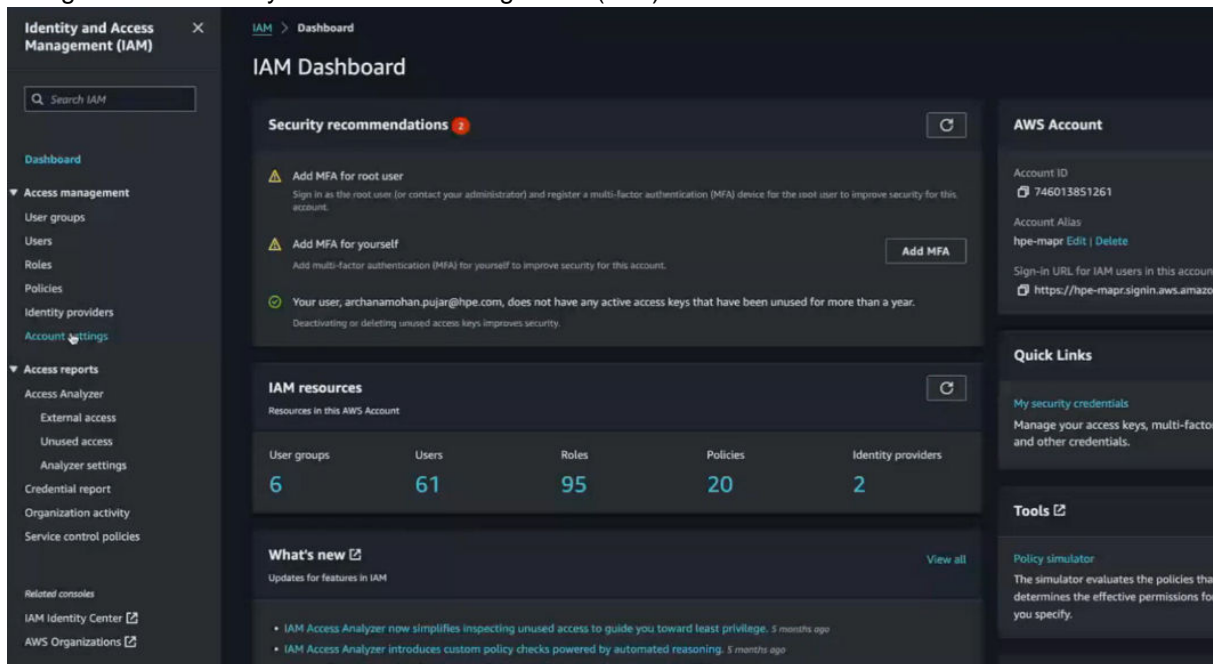
Enabling STS requires Keycloak to be deployed on a public network IP address so that AWS STS can communicate with Keycloak and verify that the JWT tokens are from the Data Fabric software. If your Keycloak deployment resides on an intranet and is not reachable by a public network, you cannot use STS. However, you can still use the access key and secret key import method.

For other STS limitations, see [Integrating the AWS Security Token Service \(STS\) with Data Fabric](#) on page 243.

### Configuring a Role for Data Fabric in AWS

Before importing an external S3 object store into the global namespace, you must configure your AWS environment as follows. In AWS:

1. Navigate to the Identity and Access Management (IAM) dashboard:

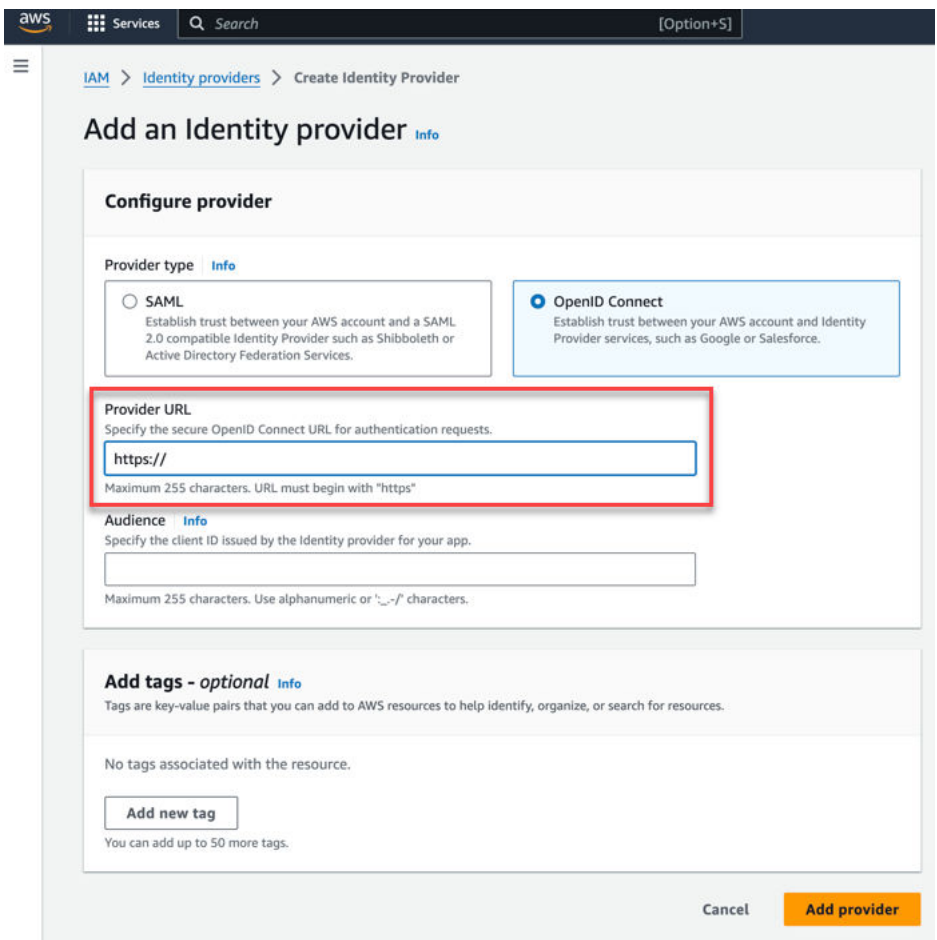


2. In the left-navigation pane, click **Identity providers**.
3. Click **Add provider**.
4. Click **OpenID Connect**:

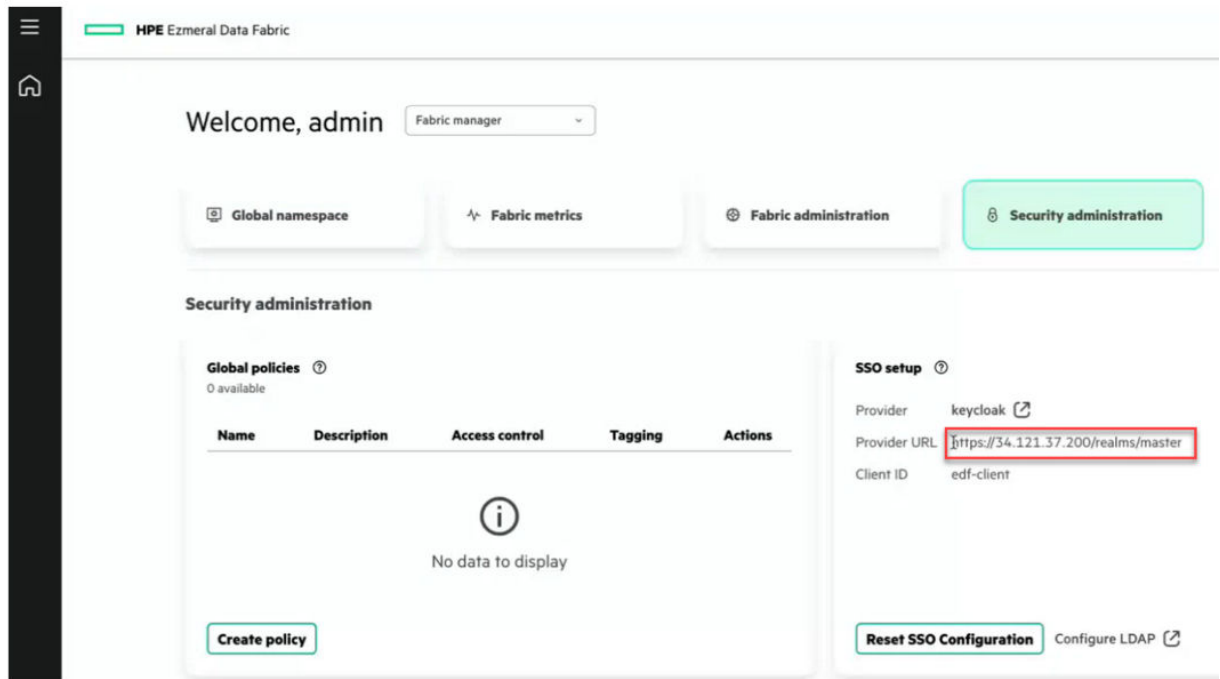
The screenshot shows the AWS IAM console interface for adding an identity provider. The breadcrumb navigation is IAM > Identity providers > Create Identity Provider. The main heading is 'Add an Identity provider' with an 'Info' link. Under the 'Configure provider' section, the 'Provider type' is set to 'OpenID Connect', which is highlighted with a red box. The 'SAML' option is unselected. Below this, the 'Provider URL' field is filled with 'https://54.245.74.65/realms/master' and a 'Get thumbprint' button is visible. The 'Audience' field is filled with 'edf-client'.

This screen enables you to create an external identity provider and specify the type as OIDC.

5. Specify the provider URL:

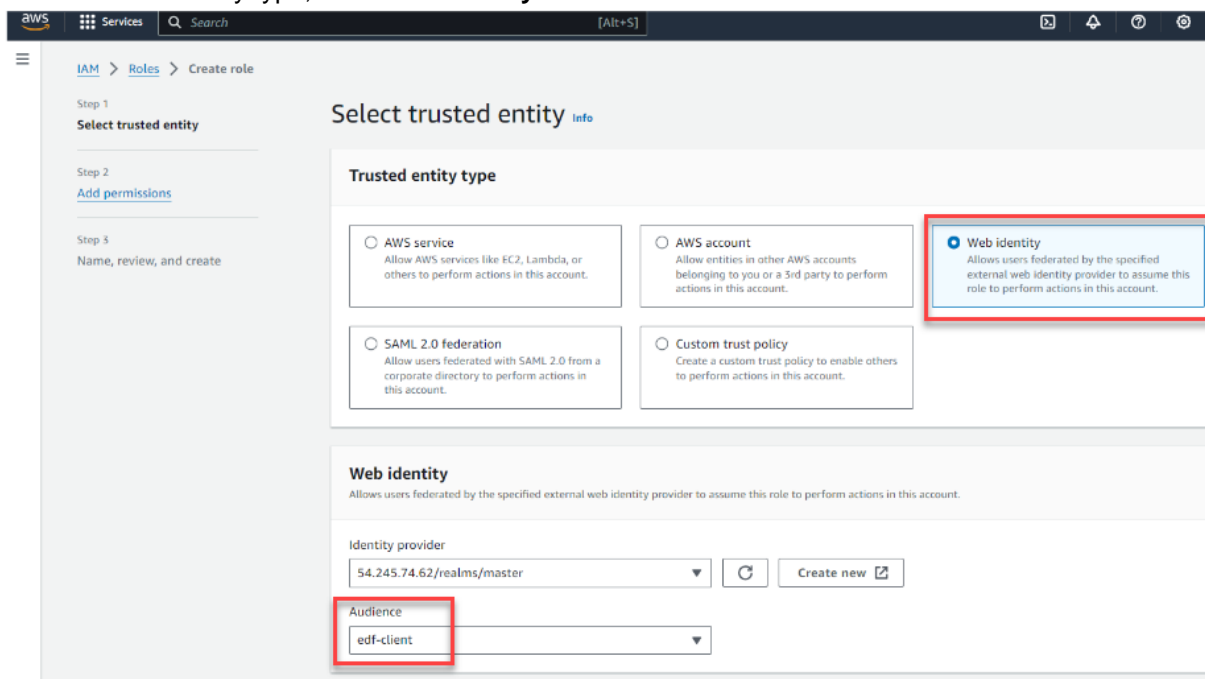


You can obtain this URL from the SSO setup card of the Data Fabric UI:

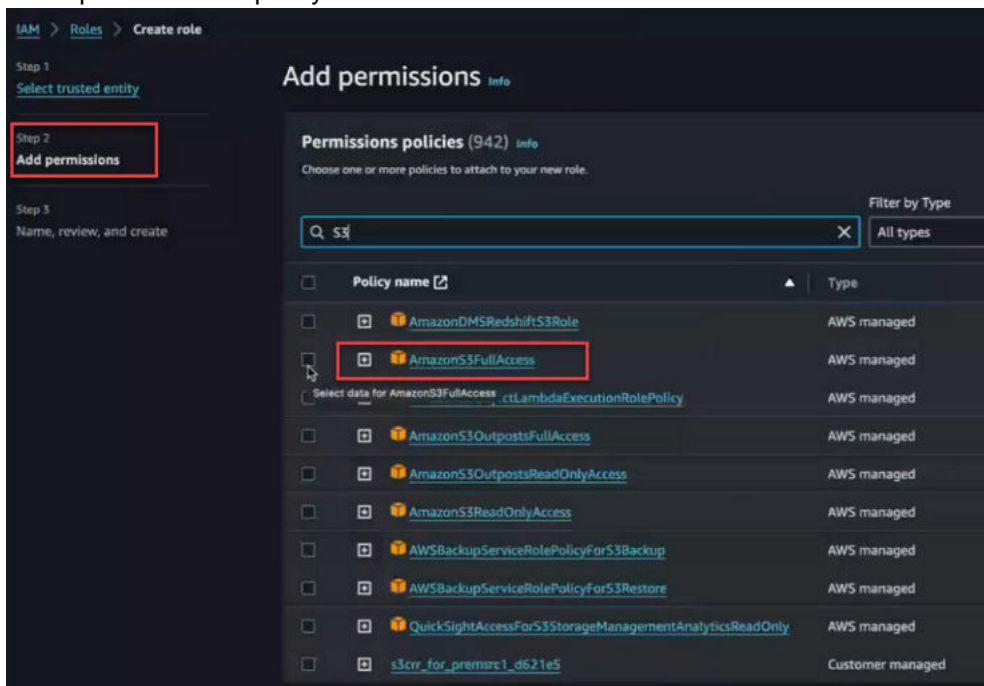


- Set the **Audience** field to **edf-client**, and click **Add provider**. The new provider is added to the list of providers on the **Identity providers** page.

7. In the left-navigation pane, click **Roles** to create a role.
8. For the trusted entity type, click **Web identity**:



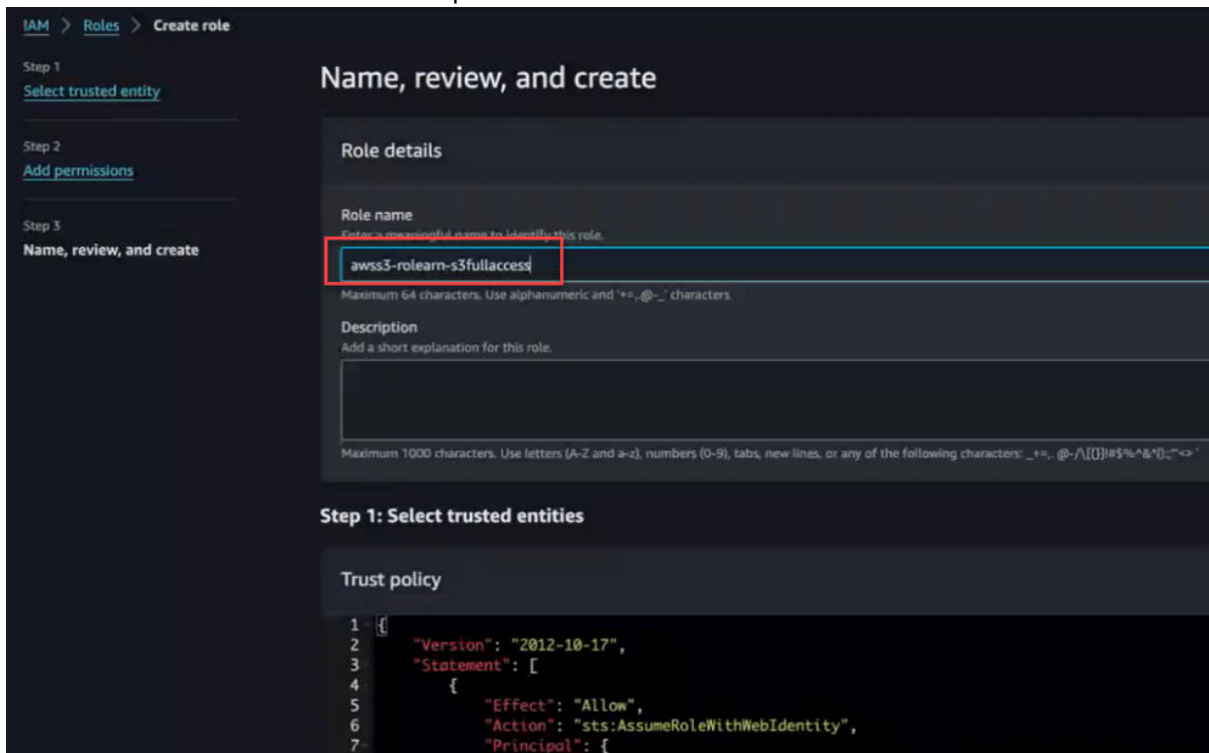
9. Select the identity provider that you created in the previous step, then verify that the **Audience** is **edf-client**.
10. Click **Next**.
11. Add the permission policies that are applicable for this role. Any entity assuming this role will have these permissions. Specify **AmazonS3FullAccess**:



12. Click **Next**.

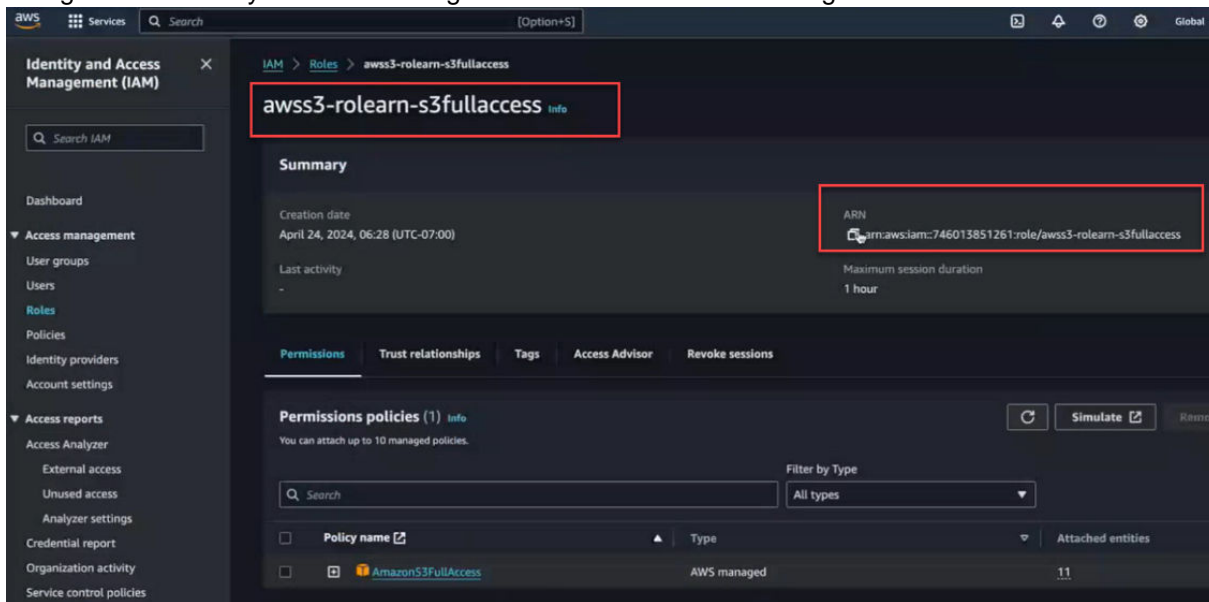


- Provide a name for the role. For example:



- Scroll down, and click **Create role**.

- Navigate to the newly created role to get its **ARN**. Note the **ARN** string:



You must provide the **ARN** in the next set of steps when you use the `maprcli` command to import AWS S3 into the global namespace by using STS.

### Importing the External AWS S3 Object Store by Using the `maprcli` Command

Starting with release 7.7.0, the `maprcli clustergroup addexternal` command is enhanced with a new option to support STS-based access. The command, which is used to import external S3 servers, includes a new `-awswebidrolearn` option. To enable STS when you run the command, you must:

- Specify the `-awswebidrolelearn` option
- Set the `-type` option to `s3`
- Set the `-s3vendor` option to `aws`

When you use these settings, Data Fabric ignores the provided access key and secret key and ensures that S3 access for the server is achieved through STS using the specified `-awswebidrolelearn`.

The following example command configures an external S3 object store to use STS access:

```
maprcli clustergroup addexternal -type s3 -s3vendor aws -awswebidrolelearn 'arn:aws:iam::74601xxxxxx:role/Keycloak-webid-s3-readonly'
```

### Related maprcli Commands

To implement the features described on this page, Data Fabric relies on the following `maprcli` command. A link to this command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `clustergroup addexternal`

## Administering Alarms

---

Manage alarms via the HPE Ezmeral Data Fabric UI.

### About this task

Alarms are the alerts or notifications generated by Data Fabric.

Alarms could be alerts related to errors, warnings, or information related to various fabric resources.

Following are the types of alarms are generated by Data Fabric.

- Fabric alarms
- User alarms
- Node alarms
- Volume alarms

Alarms raised by Data Fabric can be viewed, muted, or dismissed from the Data Fabric UI.

See [Viewing Alarms](#) on page 250 for the procedure to view alarms on the Data Fabric UI.

See [Muting/Dismissing Alarms](#) on page 251 to mute an alarm or dismiss an alarm via the Data Fabric UI.

### Viewing Alarms

View alarms on the [Data Fabric UI](#) on page 86.

#### Prerequisites

You must have access to the Data Fabric UI and the permission to view alarms.

#### About this task

Following are the types of alarms are generated by Data Fabric.

- Fabric alarms

- User alarms
- Node alarms
- Volume alarms

Use the following steps to view alarms generated by various events occurring on Data Fabric.

### Procedure

1. Log on to the Data Fabric UI.
2. Click the bell icon at the top right corner next to the help icon.

### Results

The list of alarms is displayed. You can mute or dismiss an alarm.

### Related `maprcli` Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `alarm list`

## Muting/Dismissing Alarms

Mute or dismiss an alarm via HPE Ezmeral Data Fabric UI.

### Prerequisites

You must have access to Data Fabric UI and the permission to mute or dismiss an alarm.

### About this task

You can mute or dismiss an alarm that is visible on the Data Fabric UI. An alarm can be muted for 24 hours, 6 hours or 1 hour.

Follow the steps given below to mute/dismiss an alarm.

### Procedure

1. Log on to the Data Fabric UI.
2. Click the bell icon at the top right corner next to the help icon.
3. Click **View all** to view all alarms.
4. To mute an alarm, click **Mute** and select the duration for which you wish to mute the alarm. Alternatively, click **Dismiss** to dismiss the alarm.

### Results

The alarm is muted for the specified duration or the alarm is dismissed, depending on the action you have selected.

### Related `maprcli` Commands

To implement the features described on this page, the Data Fabric UI relies on the following `maprcli` command. The command is provided for general reference. For more information, see [maprcli Commands in This Guide](#).

- `alarm mute`

## Monitoring

---

Describes monitoring with OpenTelemetry for HPE Ezmeral Data Fabric.

OpenTelemetry (OTel) is an observability framework that allows you to instrument, generate, collect, and export telemetry data. For more information on OTel, see [the official OpenTelemetry documentation](#).

The OTel endpoint provides centralized monitoring for your HPE Ezmeral Data Fabric deployments. Use OTel to generate metrics and logs for your fabrics, manage your OTel deployments through the Data Fabric UI, and view the generated telemetry data through EZ Central.

### Adding an OTel Endpoint

Describes how to add an OTel endpoint to a fabric.

#### Prerequisites

You must be a fabric manager to perform this operation.

#### About this task

Proceed as follows to add an OTel endpoint to a fabric:

#### Procedure

1. Log on to the Data Fabric UI.
2. Click the **Fabric administration** tab.
3. On the **OTEL endpoints** card, click **Add endpoint**. The **Add OTEL endpoint** side drawer opens.
4. Enter the **Name**.
5. Enter the **URL** of your OTel endpoint.
6. If your OTel endpoint contains a port, enter the port number.
7. To enable your OTel endpoint to return logs and/or metrics data, select **Logs** and/or **Metrics**.
8. Click **Select file** to select a key file to upload. Alternatively, drag and drop the key file to the **Upload files** area.
9. Click **Select file** to select a client certificate file to upload. Alternatively, drag and drop the client certificate file to the **Upload files** area.
10. Click **Add**.

#### Results

The OTel endpoint is created on the fabric.

If you selected **Logs** and/or **Metrics**, the OTel endpoint now returns the selected telemetry data for the fabric.

You can view the telemetry data generated for the fabric through EZ Central.

## Getting Started with Iceberg

---

Summarizes what you need to know to begin using Iceberg with HPE Ezmeral Data Fabric release 7.6.x.

## Version Support

HPE Ezmeral Data Fabric 7.6.x has been tested with:

- [Iceberg 1.4.2](#)
- [mapr-spark-3.3.3.0](#)
- [iceberg-spark-runtime-3.3\\_2.12-1.4.2.jar](#)

Other data-processing engines, such as open-source Spark, PrestoDB, Flink, and data-processing technologies, such as Snowflake, have not been tested.

## Catalog Support

Catalogs manage the metadata for datasets and tables in Iceberg. You must specify the catalog when interacting with Iceberg tables through Spark. The following built-in catalogs have been tested for use with Data Fabric 7.6.x:

- HiveCatalog
- HadoopCatalog

## Spark Setup for Iceberg

Setting up Spark to use Iceberg is a two-step process:

1. Add the `org.apache.iceberg:iceberg-spark-runtime-<spark.version>_<scala.version>:<iceberg.version>` jar file to your application classpath. Add the runtime to the `jars` folder in your `spark` directory. Add it directly to the application classpath by using the `--package` or `--jars` option.
2. Configure a catalog. For information about using catalogs with Iceberg, see [Catalogs](#).

For examples, see the [Spark and Iceberg Quickstart](#).

## Configuring Your Spark Application

Consider adding the following parameters to your Spark application:

```
spark.sql.catalog.<catalog_name>.type=hive
spark.sql.catalog.<catalog_name>.warehouse=<path_to_your_warehouse>
spark.sql.catalog.<catalog_name>=org.apache.iceberg.spark.SparkSessionCatalog
spark.sql.legacy.pathOptionBehavior.enabled=true
```

## Configuring Data Fabric to Track User Behavior

---

Describes how to configure Data Fabric to be able to track user behavior.

When auditing is enabled in Data Fabric, files, streams, S3 objects, and tables can be audited for cluster administration and/or data access operations. See [Enabling and Disabling Auditing of Cluster Administration](#) to enable auditing for cluster administration.



**NOTE:** Audit logs are generated into audit log files when auditing is enabled.

See [Streaming Audit Logs](#) for information on audit log streaming. See [Enabling and Disabling Audit Streaming Using the CLI](#) to enable audit streaming.

Data Fabric generates audit logs for the services that are related to various Data Fabric components. The services include CLDB, S3, MFS, auth. Auth logs are authentication audit logs.

Auditing is useful to record user behavior and assists in tracking anomalies or potential data security threats with respect to Data Fabric. See [Auditing in Data Fabric](#) for information on auditing.

See [Log files](#) on information about the various log files generated by Data Fabric.

See [Viewing Audit Logs](#) for information on how to view audit logs in Data Fabric.

Data Fabric audit logs provide insights into the activity that has taken place in relation to a cluster. The insight service is available as a distributed service on a cluster/fabric, when run in production mode.

The audit logs are stored on nodes on which the respective Data Fabric service runs. This makes it cumbersome to establish a correlation between the various logs.

Data Fabric stores audit logs in files and the audit logs can be directed to streams. However, it is not possible to run queries on streams data. Hence, the insight service picks the streams data and adds the audit data on to the respective Apache Iceberg tables.

Tools like Apache Spark can be used to run queries on the audit log data stored on Apache Iceberg, and tools like Apache Zeppelin can be used to provide graphical insights. Apache Zeppelin can make use of customizable queries to generate dashboards.



**NOTE:** Data Fabric creates the relevant Iceberg tables for storage of insight data on successful installation of the `mapr-hivemetastore` package.

When you wish to run the insight service in a trial mode, enable audit streams so that the audit logs from the audit log files are available on streams as well. The audit data can be transferred to Apache Iceberg tables and this data can be further analyzed by using Apache Spark or Apache Zeppelin.



**IMPORTANT:** The insight service must be installed on all nodes of your cluster/fabric to effectively gather information about activities and events recorded in the audit log for each node. All the data gathered by the insight service is stored in a single set of Iceberg tables.

You can write applications that consume the audit log data stored on Iceberg to detect anomalies in user behavior.

The insight service can be installed from `mapr-insight` package available on the HPE website that hosts the MEP packages.

The insight service uses Hive Metastore to store and manage the Iceberg catalog. Hive Metastore must be accessible to insight service for storing audit logs in Iceberg table. You must have the `mapr-hivemetastore` package installed and configured on your cluster to be able to use the insight service.

Hive Metastore requires a relational database management system like MySQL in production setups. See [Using MySQL for the Hive Metastore](#) to use MySQL with Hive Metastore.

To set up MySQL to work with the Hive Metastore and Data Fabric, see [Configuring a Remote MySQL Database for Hive Metastore](#).

Data Fabric supports other production grade databases like PostgreSQL. See [Configuring Data Fabric for Hive Metastore](#) for details.

### Types of Audit Logs Copied for Analysis

The following audit logs are expanded using a variant of `expand-audit` utility to include user-friendlier versions of uids, volids, etc. (usernames, volume names, etc.) before committing the logs to Iceberg table.

- CLDB logs
- MFS logs

- authentication logs
- S3 logs

There are four distinct Iceberg tables designated for each type of audit log stream. See [Enabling Insight Gathering in Trial Mode](#) and [Enabling Insight Gathering in Production Mode](#) for the Iceberg table names for each of the modes in which the insight gathering can be operated.

### Configure Data Fabric to Track User Behavior

The insight service can be enabled to gather insights at the cluster level, type level, and the node level.

If you do not want data from nodes to be copied to Iceberg, you can disable audit log/insights from node. By default, insights are disabled. When insight feature is enabled at the global level, audit logs for all types and all nodes are committed to Iceberg tables periodically. [insight](#) commands are available to enable/disable insights based on the type of logs, or at the node level on a node-by-node basis.

**!** **IMPORTANT:** Install and set up your fabric/cluster before installing the packages for the insight service and enabling the insight service.

You can use tools like Spark and Zeppelin to run queries on the Iceberg tables to generate various reports and charts required by you to detect any anomalies in user behavior related to the data access operations and cluster administration.

You can customize the insights with the [insight](#) CLI command.

For instance, you can turn off insight gathering on some nodes (although this is not recommended), or you can turn off insight gathering of certain audit components such as S3 due to heavy S3 traffic on your cluster/fabric.

### Enable Insight Gathering

See [insight cluster](#) to enable insight gathering.

## Reference

---

Provides reference information for the HPE Ezmeral Data Fabric.

## Release History

---

Describes the currently released versions of the HPE Ezmeral Data Fabric as-a-service platform.

Only core version 7.4.0 and later are currently supported for the as-a-service platform of the HPE Ezmeral Data Fabric.

Date	Build Version	Core Version	EEP Version	Docker Image
October 31, 2024	7.9.0.0.202410281 25132.GA	7.9.0.0	9.3.1	maprtech/ edf-seed-container:7.9.0_9.3.1_edf
July 31, 2024	7.8.0.0.202407240 93818.GA	7.8.0.0	9.3.0	maprtech/ edf-seed-container:7.8.0_9.3.0_edf
April 30, 2024	7.7.0.0.202404220 22544.GA	7.7.0.0	9.2.2	maprtech/ edf-seed-container:7.7.0_9.2.2_edf
February 16, 2024	7.6.1.0.202402071 05416.GA	7.6.1.0	9.2.1	maprtech/ edf-seed-container:7.6.1_9.2.1_edf

Date	Build Version	Core Version	EEP Version	Docker Image
October 30, 2023	7.5.0.0.202310262 22149.GA	7.5.0.0	9.2.0	maprtech/ edf-seed-container:7.5.0_9.2.0_edf
August 8, 2023	7.4.0.0.202307281 33744.GA	7.4.0.0	9.1.2	maprtech/ edf-seed-container:7.4.0_9.1.2_dfaas
May 12, 2023	7.3.0.0.202304250 02320.GA	7.3.0.0	9.1.1	maprtech/ dev-sandbox-container:7.3.0_9.1.1_dfaas

### Related concepts

[Viewing the Software Version](#) on page 115

Describes several ways to identify the core software version for a fabric.

## Cloud Instance Specifications

Compares different aspects of the supported cloud instances of the HPE Ezmeral Data Fabric.

### AWS Cloud Instance Specifications

The following table describes the AWS cloud instance for different storage tiers:

Specification	Storage Tier				
	100GB	1TB	10TB	100TB	1PB
Number of Instances	1	3	5	10	15
Instance Type	m6i.4xlarge	m6i.4xlarge	m6i.4xlarge	m6i.4xlarge	m6i.4xlarge
Number of DF Disks	1	3	4	10	14
DF Disk Type	st1	gp3	gp3	gp3	gp3
DF Disk Size	128	128	512	1024	5120
Swap Disk Type	gp3	gp3	gp3	gp3	gp3
RAM	64GB	64GB	64GB	64GB	64GB
CPU	16	16	16	16	16
Swap Disk Size	16	64	128	256	512

### Azure Cloud Instance Specifications

The following table describes the Azure cloud instance for different storage tiers:

Specification	Storage Tier				
	100GB	1TB	10TB	100TB	1PB
Number of Instances	1	3	5	10	15
Instance Type	Standard_B 12ms	Standard_B 16ms	Standard_B 16ms	Standard_B 20ms	Standard_B 20ms
Number of DF Disks	1	3	4	10	14
DF Disk Type	Standard_L RS	Standard_L RS	Standard_L RS	Standard_L RS	Standard_L RS
DF Disk Size	100	115	512	1024	4800



Specification	Storage Tier				
	100GB	1TB	10TB	100TB	1PB
RAM	48GB	64GB	64GB	80GB	80GB
CPU	12	16	16	20	20
Swap Disk Type	Standard_L RS	Standard_L RS	Standard_L RS	Standard_L RS	Standard_L RS
Swap Disk Size	32	32	32	32	32

### GCP Cloud Instance Specifications

The following table describes the GCP cloud instance for different storage tiers:

Specification	Storage Tier				
	100GB	1TB	10TB	100TB	1PB
Number of Instances	1	3	5	10	15
RAM	64GB	64GB	64GB	64GB	64GB
CPU	16	16	16	16	16
Instance Type	n2-standar d-16	n2-standar d-16	n2-standar d-16	n2-standar d-16	n2-standar d-16
Root Disk Type	pd-ssd	pd-ssd	pd-ssd	pd-ssd	pd-ssd
Root Disk Size	200	200	200	200	200
Number of DF Disks	2	3	4	10	14
DF Disk Type	pd-standard	pd-standard	pd-standard	pd-standard	pd-standard
DF Disk Size	100	115	512	1024	4880
Swap Disk Type	pd-standard	pd-standard	pd-standard	pd-standard	pd-standard
Swap Disk Size	32	32	32	32	32

### Third-Party Storage Solutions

Describes global-namespace support for HPE partner storage technologies, including Scality, WEKA, and VAST.

The HPE Ezmeral Data Fabric 7.6.x and later global namespace is compatible with the following third-party, object-storage solutions:

Storage Product	External NFS Integration with GNS		External S3 Integration with GNS	
	Using System Security(AD/LDAP)	Using Kerberos	Using Secret Key and Access Key	
			HTTPS	HTTP
WEKA	Supported	Not Supported	Supported	Supported
VAST Data on HPE Alletra	Supported	Supported	Supported	Supported
Scality ARTESCA	N/A*	N/A*	Supported	Supported
Scality RING	Supported	Supported	Supported	Supported

Storage Product	External NFS Integration with GNS		External S3 Integration with GNS	
	Using System Security(AD/LDAP)	Using Kerberos	Using Secret Key and Access Key	
			HTTPS	HTTP
Minio Server	N/A*	N/A*	Supported	Supported
NFS Ganesha	Supported	Supported	N/A*	N/A*

\*N/A means not supported by the storage vendor.

#### More information

[Scality Documentation](#)

[WEKA Documentation](#)

[VAST Data Documentation](#)

## Port Information

Describes the ports used by HPE Ezmeral Data Fabric services.

The following table lists the principal services, the ports they use, and the associated protocol. For traffic within a subnet, the port is not relevant. Hosts can communicate on any available port:

Service	Port	Protocol
CLDB	7222	TCP
Data Fabric Gateway	7660	TCP
Data Fabric Keycloak	443*	TCP
Data Fabric UI	8443	TCP
Fileserver	5660	TCP
Fileserver	5692	TCP
Installer	9443	TCP
MOSS (Multithreaded Object Store Server)	9000	TCP
NFS (into VPC UDP)	111	UDP
NFS (into VPC TCP)	2049	TCP
OpenTSDB	4242	TCP
SSH	22	TCP
Kafka Wire Protocol Service	9092	TCP

The port can be 6443 for upgraded fabrics. For more information, see [Special Considerations for Upgrading to Release 7.7.0 and Later](#) on page 81.

## maprcli Commands in This Guide

Describes how to use `maprcli` commands provided as reference links in this guide.

Some procedures in this guide include links to `maprcli` commands at the bottom of the page. Clicking a link to a `maprcli` command opens a page in the customer-managed documentation [website](#), where more detailed information for all `maprcli` commands is located.

These commands are listed for reference purposes only. The commands are not intended to replace the documented Data Fabric UI procedures. In addition, some as-a-service features are not supported on the customer-managed platform.

HPE Ezmeral Data Fabric users are encouraged to use the Data Fabric UI for all operations. Users of the as-a-service HPE Ezmeral Data Fabric generally are not encouraged to use `maprccli` commands when a UI control is available. However, users who are interested can read more about the commands in the customer-managed documentation.

## Operating System Support Matrix

The tables on this page show the Linux operating-system versions that are supported for HPE Ezmeral Data Fabric releases.

On-premises deployments can use any of the following operating-system versions. Cloud (AWS, Azure, GCP) deployments use Rocky Linux as the default OS.

### Red Hat Enterprise Linux (64-bit)

RHEL Version	Release 7.9.x	Release 7.8.x	Release 7.7.x	Release 7.6.x	Release 7.5.0
9.4	Yes	Yes	No	No	No
9.0	Yes	Yes	Yes	No	No
8.10	Yes	Yes	No	No	No
8.8	Yes	Yes	Yes	Yes	Yes
8.6	Yes	Yes	Yes	Yes	Yes
8.5	Yes	Yes	Yes	Yes	Yes
8.4	Yes	Yes	Yes	Yes	Yes
8.3	Yes	Yes	Yes	Yes	Yes
8.2	Yes	Yes	Yes	Yes	Yes
8.1	Yes	Yes	Yes	Yes	Yes

### Rocky Linux (64-bit)

Rocky Version	Release 7.9.x	Release 7.8.x	Release 7.7.x	Release 7.6.x	Release 7.5.0
9.4	Yes	Yes	No	No	No
8.5	Yes	Yes	Yes	Yes	Yes
8.4	Yes	Yes	Yes	Yes	Yes

### Ubuntu (64-bit)

Ubuntu Version	Release 7.9.x	Release 7.8.x	Release 7.7.x	Release 7.6.x	Release 7.5.0
22.04	Yes <sup>1</sup>	Yes <sup>1</sup>	Yes <sup>1</sup>	No	No
20.04	Yes	Yes	Yes	Yes	Yes
18.04	Yes	Yes	Yes	Yes	Yes

<sup>1</sup>Release 7.x has a dependency on the `libssl1.1` package, which is not included in Ubuntu 22.04. To resolve this issue, see the entry for MFS-18734 in [Known Issues \(Release 7.9.0\)](#) on page 10.

**SLES (64-bit)**

SLES Version	Release 7.9.x	Release 7.8.x	Release 7.7.x	Release 7.6.x	Release 7.5.0
15 SP5	Yes	Yes	No	No	No
15 SP3	Yes	Yes	Yes	Yes	Yes
15 SP2	Yes	Yes	Yes	Yes	Yes

**Oracle Enterprise Linux (OEL)**

OEL Version	Release 7.9.0	Release 7.8.0	Release 7.7.x	Release 7.6.x	Release 7.5.0
8.4	No	No	Yes	Yes	Yes
8.3	No	No	Yes	Yes	Yes
8.2	No	No	Yes	Yes	Yes

## Container Image Vulnerabilities and CVE Reports

---

Describes how HPE Ezmeral Engineering provides software updates to address container image vulnerabilities.

HPE Ezmeral Engineering takes security very seriously and makes every effort to ensure that the container images for HPE Ezmeral software products are free of known vulnerabilities at the time of release. However, because new vulnerabilities are always being discovered and reported, it is likely that scanning product images with tools such as Trivy will show lists of CVEs that affect packages inside the images.

The HPE Ezmeral Engineering team also regularly scans product images to identify new vulnerabilities and creates action plans to modify the product images. Please note that most vulnerabilities are present in open-source software leveraged by HPE Ezmeral Engineering. Therefore, HPE Ezmeral Engineering determines when it is best to update products with updated open-source content.

HPE Ezmeral Engineering typically updates vulnerable packages from one minor software product version to the next (for example, from 1.3 to 1.4). For critical vulnerabilities, HPE may provide security-patched container images outside of the established software release cycle, in accordance with the following table.

To keep your platform as secure as possible, please ensure that you upgrade or patch your HPE Ezmeral Software to the latest available software.

Severity (CVSS Base Score Range)	SLA of Response
Critical (9.0 – 10.0)	HPE Ezmeral Engineering will prioritize and begin working on a fix. The team will make the fix available as soon as possible. This might take the form of a special maintenance release of an HPE Ezmeral software product for the sole purpose of making the fix available. If it is possible to deploy the fix as a patch more quickly or conveniently, the patch will also be made available. In the meantime, the support team will work with the community to mitigate the issue.
High (7.0 – 8.9)	HPE Ezmeral Engineering will include a fix in the next planned release (major or minor) of the HPE Ezmeral software product. HPE Ezmeral software releases typically happen on a quarterly basis. The fix will be made available in patch form for customers who want to deploy it sooner, and the support team will assist with applying the patch.
Medium (4.0 – 6.9)	HPE Ezmeral Engineering will include a fix in the next planned release (major or minor) of the HPE Ezmeral product.
Low (0.1 – 3.9)	HPE Ezmeral Engineering will include a fix in the next major release of the HPE Ezmeral product, or the team will provide detailed steps that can be taken to mitigate the issue.

## Doc Site Available as a PDF

---


Provides a link to the downloadable PDF file containing all the information for the current release.

For a given release, you can access HPE Ezmeral Data Fabric documentation as a single, downloadable PDF file. A PDF file of each release is compiled several weeks after the release becomes public and is available for download from the [HPE Support Center](#).

Here is the PDF location for the current release:

To download the PDF from the [HPE Support Center](#):

1. Navigate to the Support Center home page for a Data Fabric release:
  - [Support Center HPE Ezmeral Data Fabric 7.7.0 Documentation](#)
  - [Support Center HPE Ezmeral Data Fabric 7.6.1 Documentation](#)
  - [Support Center HPE Ezmeral Data Fabric 7.5.0 Documentation](#)
  - [Support Center HPE Ezmeral Data Fabric 7.4.0 Documentation](#)
2. Above the right-navigation pane, click the **PDF** button, and select **Export all content**. A PDF file is downloaded to your workstation.

 **IMPORTANT:** PDF files are updated infrequently. They are a snapshot of the available information at the time the PDF was created. For the most current technical information, HPE recommends that you refer to the HTML pages at [this location](#). The HTML pages:

- Are updated continuously.
- Provide a **Feedback** button that enables you to submit comments or corrections.
- Can make it easier to access multimedia resources, such as product videos.

## Product Licensing

---

Provides information related to product licensing.

### Additional License Authorizations (ALA)

Provides Additional License Authorizations for HPE Ezmeral Software, including HPE Ezmeral Runtime Enterprise, HPE Ezmeral ML Ops, HPE Ezmeral Data Fabric, and Open Source Software.

[Additional License Authorizations for HPE Ezmeral Software](#)

### Open-Source Software Acknowledgements (Release 7.9.x)

Provides licensing information and acknowledges the use of open-source projects with HPE software.

#### About the NOTICE.txt File

The `NOTICE.txt` file provides licensing information and software acknowledgements for open-source software used by the HPE Ezmeral Data Fabric. On a release 7.9.x Data Fabric node, you can find the file in the `/opt/mapr` directory. The release 7.9.x file contains the following information:

## Open Source Notice

The Hewlett Packard Enterprise ("HPE") software accompanied by this notice is provided along with certain third party software licensed under various open source software licenses ("Open Source Components"). The below list of Open Source Components includes, as applicable, copyright notices, original source code URLs and license URLs, and indicates whether HPE has modified the original source code of the Open Source Components. With respect to licenses that require a particular language to be provided (such as the complete terms of the license itself), that language is included below under the first Open Source Component that is subject to such license.

With respect to Open Source Components licensed under the AGPL, CPL, GPL or LGPL, HPE hereby offers to provide upon request the source code thereof, including the HPE modifications, if any. Such modifications are documented by way of comments included in the source code files.

In addition to the warranty disclaimers contained in the open source licenses linked below and thus included herein by reference, HPE makes the following disclaimers regarding the Open Source Components on behalf of itself, the copyright holders, contributors, and licensors of such Open Source Components:

TO THE FULLEST EXTENT PERMITTED UNDER APPLICABLE LAW, THE OPEN SOURCE COMPONENTS ARE PROVIDED BY THE COPYRIGHT HOLDERS, CONTRIBUTORS, LICENSORS, AND HPE "AS IS" AND ANY REPRESENTATIONS OR WARRANTIES OF ANY KIND, WHETHER ORAL OR WRITTEN, WHETHER EXPRESS, IMPLIED, OR ARISING BY STATUTE, CUSTOM, COURSE OF DEALING, OR TRADE USAGE, INCLUDING WITHOUT LIMITATION THE IMPLIED WARRANTIES OF TITLE, MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT, ARE DISCLAIMED. IN NO EVENT WILL THE COPYRIGHT OWNER, CONTRIBUTORS, LICENSORS, OR HPE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION), HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THE OPEN SOURCE COMPONENTS, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## Project-Specific Copyright, Source Code, and License Information

-----  
Hadoop

Copyright (c) 2011 The Apache Software Foundation.

Source code: <http://hadoop.apache.org/>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----

Apache Hive

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----

Apache Zeppelin

Copyright (c) 2015 - 2016 The Apache Software Foundation

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----

Apache Tez

Copyright (c) 2016 The Apache Software Foundation

Source code: <git://git.apache.org/tez.git>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----

Apache HBase

Source code: <http://hbase.apache.org/>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----

Async HBase

Copyright (C) 2010-2012 The Async HBase Authors. All rights reserved.

New BSD License

<http://opensource.org/licenses/BSD-3-Clause>

-----

Apache Thrift

Copyright (c) 2006-2010 The Apache Software Foundation.

Source code: <http://incubator.apache.org/thrift/>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----

Apache RocksDB

Copyright (c) 2004 The Apache Software Foundation.

Source code: <http://incubator.apache.org/thrift/>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

---

#### Apache Kafka

Source code: <https://github.com/apache/kafka>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

---

#### Elasticsearch

Copyright 2009-2016 Elasticsearch

Source code: <https://github.com/elastic/elasticsearch>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0>

---

#### Grafana

Copyright 2012-2013 Elasticsearch BV

Source code: <https://github.com/grafana/grafana>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0>

---

#### Kibana

Copyright 2012-2016 Elasticsearch BV

Source code: <https://github.com/elastic/kibana>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0>

---

#### collectd

Copyright (C) 1989, 1991 Free Software Foundation

Source code: <https://github.com/collectd/collectd>

License: LGPL 2  
<https://github.com/collectd/collectd/blob/master/COPYING>

---

#### fluentd

Copyright (C) 2011 FURUHASHI Sadayuki



Source code: <https://github.com/fluent/fluentd>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0>

-----  
 MySQL Connector/J

Copyright (C) 1989, 1991 Free Software Foundation

Source code: <https://github.com/mysql/mysql-connector-j>

License: LGPL 2  
<https://github.com/mysql/mysql-connector-j/blob/release/5.1/COPYING>

-----  
 Ganesha

Copyright (C) 2007 Free Software Foundation, Inc.

Source code: <https://github.com/nfs-ganesha/nfs-ganesha>

License: LGPL 3  
<https://github.com/nfs-ganesha/nfs-ganesha/blob/next/src/LICENSE.txt>

-----  
 Minio

Copyright (c) 2004, The Apache Software Foundation

MinIO Client (C) 2014-2020 MinIO, Inc.

This product includes software developed at MinIO, Inc.  
 (<https://min.io/>).

The MinIO project contains unmodified/modified subcomponents too with separate copyright notices and license terms. Your use of the source code for the these subcomponents is subject to the terms and conditions of the following licenses.

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
 gRPC

Copyright (c) 2004, The Apache Software Foundation

Source code: <https://github.com/grpc/grpc>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
 Kafka-connect-jdbc  
 Copyright (c) 2015 Confluent Inc.

The following libraries are included in packaged versions of this project:

- \* SQLite JDBC Driver
  - \* COPYRIGHT: Copyright Taro L. Saito, David Crenshaw
  - \* LICENSE: licenses/LICENSE.apache2.txt
  - \* NOTICE: licenses/NOTICE.sqlite-jdbc.txt
  - \* HOMEPAGE: <https://github.com/xerial/sqlite-jdbc>
- \* PostgreSQL JDBC Driver
  - \* COPYRIGHT: Copyright 1997-2011, PostgreSQL Global Development Group
  - \* LICENSE: licenses/LICENSE.bsd.txt
  - \* HOMEPAGE: <https://jdbc.postgresql.org/>
- \* MariaDB JDBC Driver
  - \* COPYRIGHT: Copyright 2012 Monty Program Ab., 2009-2011, Marcus Eriksson
  - \* LICENSE: licenses/LICENSE.lgpl.txt
  - \* HOMEPAGE: <https://mariadb.com/kb/en/mariadb/about-mariadb-connector-j/>

-----  
 kafka-connect-hdfs  
 Copyright (c) 2015 Confluent Inc.  
 -----

kafka-rest  
 Confluent Community License Agreement Version 1.0  
 -----

schema-registry

The project is licensed under the Confluent Community License, except for client libs, which is under the Apache 2.0 license.

See LICENSE file in each subfolder for detailed license agreement.  
 -----

KSQL

Confluent Community License Agreement Version 1.0

The project is licensed under the Confluent Community License.

Apache, Apache Kafka, Kafka, and associated open source project names are trademarks of the Apache Software Foundation.  
 -----

rest-utils

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

The following libraries are included in packaged versions of this project:

- \* ClassMate
  - \* COPYRIGHT: Copyright 2010 The Apache Software Foundation
  - \* LICENSE: licenses/LICENSE.apache2.txt
  - \* HOMEPAGE: <https://github.com/cowtowncoder/java-classmate>

```

* Confluent Common
* COPYRIGHT: Confluent Inc.
* LICENSE: licenses/LICENSE.apache2.txt
* HOMEPAGE: https://github.com/confluentinc/common

* Hamcrest
* COPYRIGHT: Copyright (c) 2000-2006, www.hamcrest.org
* LICENSE: licenses/LICENSE.bsd.txt
* HOMEPAGE: http://hamcrest.org/

* Hibernate
* COPYRIGHT: licenses/COPYRIGHT.hibernate.txt
* LICENSE: licenses/LICENSE.apache2.txt
* HOMEPAGE: http://hibernate.org/validator/

* HK2
* COPYRIGHT: Copyright (c) 2010-2014 Oracle and/or its affiliates. All
rights reserved.
* LICENSE: licenses/LICENSE.cddl+gpl2.html
* HOMEPAGE: https://hk2.java.net

* Jackson annotations
* LICENSE: licenses/LICENSE.jackson-annotations.txt (Apache 2)
* HOMEPAGE: http://github.com/FasterXML/jackson

* Jackson core
* LICENSE: licenses/LICENSE.jackson-core.txt (Apache 2)
* NOTICE: licenses/NOTICE.jackson-core.txt
* HOMEPAGE: http://github.com/FasterXML/jackson

* Jackson databind
* LICENSE: licenses/LICENSE.jackson-databind.txt (Apache 2)
* NOTICE: licenses/NOTICE.jackson-databind.txt
* HOMEPAGE: http://github.com/FasterXML/jackson

* Jackson jaxrs-json-provider
* LICENSE: licenses/LICENSE.jackson-core.txt (Apache 2)
* NOTICE: licenses/NOTICE.jackson-core.txt
* HOMEPAGE: http://github.com/FasterXML/jackson

* Javassist
* COPYRIGHT: Copyright (C) 1999- by Shigeru Chiba, All rights reserved.
* LICENSE: licenses/LICENSE.javassist.txt (MPL, LGPL, Apache 2)
* HOMEPAGE: http://www.javassist.org

* javax.annotation-api, javax.el, javax.el-api, javax.inject,
javax.servlet, javax.ws.rs-api, javax.validation
* COPYRIGHT: Copyright Oracle
* LICENSE: licenses/LICENSE.cddl+gpl2.html

* JBoss Logging
* COPYRIGHT: Copyright 2014 Red Hat, Inc.
* LICENSE: licenses/LICENSE.apache2.txt
* HOMEPAGE: http://www.jboss.org

* Jersey
* LICENSE: licenses/LICENSE.cddl+gpl2.html
* HOMEPAGE: http://jersey.java.net

* Jetty
* COPYRIGHT: Copyright Mort Bay Consulting Pty Ltd unless otherwise noted
* LICENSE: licenses/LICENSE.apache2.txt, licenses/LICENSE.epl.html
* NOTICE: licenses/NOTICE.jetty.txt
* HOMEPAGE: http://eclipse.org/jetty/

```

```
* JUnit
* LICENSE: licenses/LICENSE.epl.txt
* NOTICE: licenses/NOTICE.junit.txt
* HOMEPAGE: http://junit.org/
```

---

#### KStreams

Copyright (c) 2004, The Apache Software Foundation

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

---

#### HttpComponents

Copyright (c) 2004, The Apache Software Foundation

Source code: <http://hc.apache.org>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

---

#### Quartz-Scheduler Hazelcast Job Store

Copyright (c) 2004, The Apache Software Foundation

Source code: <https://github.com/FlavioF/quartz-scheduler-hazelcast-jobstore>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

---

#### Quartz

Copyright (c) 2004, The Apache Software Foundation

Source code: <https://github.com/quartz-scheduler/quartz>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

---

#### AWS JAVA-SDK

Copyright (c) 2004, The Apache Software Foundation

Source code: <https://aws.amazon.com/sdk-for-java>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

---

## ZIP4J

Copyright (c) 2004, The Apache Software Foundation

Source code: <http://www.lingala.net/zip4j/>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

---

## Args4j

Copyright (c) 2013, Kohsuke Kawaguchi and other contributors

Source code: <https://github.com/kohsuke/args4j>

License: MIT  
<http://www.opensource.org/licenses/mit-license.php>

---

## Curator

Copyright (c) 2004, The Apache Software Foundation

Source code: <https://curator.apache.org/>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

---

## Hazelcast Discovery Plugin for Apache ZooKeeper

Copyright (c) 2004, The Apache Software Foundation

Source code: <https://github.com/hazelcast/hazelcast-zookeeper>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

---

## Intel(R) Intelligent Storage Acceleration Library

Copyright (c) 2004, The Apache Software Foundation

Source code: <https://github.com/01org/isa-l>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

---

## Intel(R) Intelligent Storage Acceleration Library Crypto Version

Copyright (c) 2004, The Apache Software Foundation

Source code: [https://github.com/01org/isa-l\\_crypto](https://github.com/01org/isa-l_crypto)

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
MapR-DB Client Driver for Python Application

Copyright (c) 2004, The Apache Software Foundation

Source code: <https://github.com/mapr/maprdb-python-client>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
MapR-DB Client Driver for Node.JS Application

Copyright (c) 2004, The Apache Software Foundation

Source code: <https://github.com/mapr/maprdb-node-client>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
Mesosphere Mesos-DNS

Copyright (c) 2015, The Apache Software Foundation

Source code: <https://github.com/mesosphere/mesos-dns>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
Java Library for Processing JSON

Copyright (c) 2015, The Apache Software Foundation

Source Code: Source: <https://github.com/FasterXML>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>  
<http://wiki.fasterxml.com/JacksonLicensing>

-----  
Spring Framework

Source code: <https://github.com/spring-projects>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
Spring Shell

Source code: <https://github.com/spring-projects>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
 TCMalloc

New BSD License

<http://opensource.org/licenses/BSD-3-Clause>  
 -----

Antlr4 Runtime

Source Code: <https://github.com/antlr/antlr4/>

License: BSD License

<http://wwwantlr.org/license.html>  
 -----

AOP Alliance

Source Code: <http://sourceforge.net/p/aopalliance/code/>

License: Public Domain  
 -----

ASM Java Bytecode Manipulation and Analysis Framework

Source Code: [http://forge.ow2.org/plugins/scmsvn/index.php?group\\_id=23](http://forge.ow2.org/plugins/scmsvn/index.php?group_id=23)

License: BSD License

<http://forge.ow2.org/projects/asm/>  
 -----

JLine (Java Library for Handling Console Input v. 2)

Source Code: <https://github.com/jline/jline2>

License: BSD License

<https://github.com/jline/jline2/blob/master/LICENSE.txt>  
 -----

OpenTSDB

LGPL v2.1

<https://github.com/OpenTSDB/opentsdb/blob/master/COPYING.LESSER>  
 -----

Apache Spark

License: Apache License, Version 2.0

<http://www.apache.org/licenses/LICENSE-2.0.html>  
 -----

Snappy 1.0.5

New BSD License

<http://opensource.org/licenses/BSD-3-Clause>

-----  
Hue

Copyright (c) Cloudera

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
Ansible

GPL  
<https://github.com/ansible/ansible/blob/devel/COPYING>

-----  
Apache Drill

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
gperftools 2.0  
New BSD License

<http://opensource.org/licenses/BSD-3-Clause>

-----  
Apache ZooKeeper

Copyright (c) 2009 The Apache Software Foundation.

Source code: <http://zookeeper.apache.org>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
Open

Application Interface (OJAI)

Copyright (c) 2015 The Apache Software Foundation.

Source code: <https://github.com/ojai/ojai>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
Apache Commons

Copyright (c) 2003-2007 The Apache Software Foundation.

Source code and additional copyright: <http://commons.apache.org/>

License: Apache License, Version 2.0



<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
 Google Collections (Guava)

Copyright (c) 2007 Google Inc.

Source code: <http://code.google.com/p/guava-libraries/>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
 Apache Tomcat

Copyright (c) 1999-2011 The Apache Software Foundation.

Source code: <http://tomcat.apache.org>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
 Jetty Web Container

Copyright (c) 1995-2009 Mort Bay Consulting Pty Ltd.

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
 Open Json

Android JSON library  
 Copyright (C) 2010 The Android Open Source Project

Source code: <https://github.com/tdunning/open-json>  
 License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
 JUnit

License: Common Public License - v 1.0  
<http://www.junit.org/license>

-----  
 log4j

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
 JavaMail

Copyright (c) 1997-2011, Oracle and/or its affiliates.

Source code: <http://www.oracle.com/technetwork/java/index-138643.html>

License: Oracle Corporation ("ORACLE") ENTITLEMENT for SOFTWARE  
See below.

-----  
Protocol Buffers

Copyright (c) 2008 Google Inc.

Source code: <http://protobuf.googlecode.com>

License: New BSD License  
<http://www.opensource.org/licenses/bsd-license.php>

-----  
uuid - DCE compatible Universally Unique Identifier library

Copyright (C) 1996, 1997, 1998 Theodore Ts'o.

License: below.

-----  
MurmurHash

Source code: <http://code.google.com/p/smhasher/>

License: MIT License  
<http://www.opensource.org/licenses/mit-license.php>

-----  
Eval - A Simple Expression Evaluator for Java

Source code: <http://java.net/projects/eval/pages/Home>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
Guava Release 11.0.1

Source code: <http://code.google.com/p/guava-libraries/>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
suEXEC - Apache HTTP Server Version 2.0

Source code: <http://httpd.apache.org/docs/2.0/suexec.html>

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
LZ4 compression

Copyright (C) 2011-2012, Yann Collet.

Source code: <http://code.google.com/p/lz4/>

License: New BSD License

<http://www.opensource.org/licenses/bsd-license.php>

-----  
ZLIB compression

Copyright (C) 1995-2012 Jean-loup Gailly and Mark Adler

Source code: <http://www.zlib.net/>

License: below.

-----  
D3.js

Copyright (c) 2012, Michael Bostock

License: New BSD License (below)

<http://opensource.org/licenses/BSD-3-Clause>

-----  
c3p0 - JDBC3 Connection and Statement Pooling

Copyright (c) 2012 Machinery For Change, Inc.

Source code: <http://www.mchange.com/projects/c3p0/index.html>

License: Lesser GNU Public License (LGPL)

<http://www.gnu.org/licenses/lgpl.html>

-----  
Hibernate

Source code: <http://www.hibernate.org/>

License: Lesser GNU Public License (LGPL) v2.1

<http://www.gnu.org/licenses/old-licenses/lgpl-2.1.html>

-----  
Trove

Source code: <https://bitbucket.org/trove4j/trove>

License: Lesser GNU Public License (LGPL) v2.1

<http://www.gnu.org/licenses/old-licenses/lgpl-2.1.html>

-----  
SOI

Source code: <http://soci.sourceforge.net/>

License: Boost Software License

[http://www.boost.org/LICENSE\\_1\\_0.txt](http://www.boost.org/LICENSE_1_0.txt)

-----

PCRE

Copyright (c) 2007-2012 Google Inc.  
 Copyright (c) 2009-2012 Zoltan Herczeg  
 Copyright (c) 1997-2012 University of Cambridge

Source code: <http://www.pcre.org/>

License: New BSD License  
<http://www.opensource.org/licenses/bsd-license.php>

-----

"react@16.14.0"

"licenses": "MIT",  
 "repository": "https://github.com/facebook/react",  
 "licenseUrl": "https://github.com/facebook/react/raw/main/LICENSE",

-----

"react@17.0.2"

"licenses": "MIT",  
 "repository": "https://github.com/facebook/react",  
 "licenseUrl": "https://github.com/facebook/react/raw/main/LICENSE",

-----

"react@18.3.1"

"licenses": "MIT",  
 "repository": "https://github.com/facebook/react",  
 "licenseUrl": "https://github.com/facebook/react/raw/main/LICENSE",

-----

"prop-types@15.7.2"

"licenses": "MIT",  
 "repository": "https://github.com/facebook/prop-types",  
 "licenseUrl": "https://github.com/facebook/prop-types/raw/main/LICENSE",

-----

"prop-types@15.8.1"

"licenses": "MIT",  
 "repository": "https://github.com/facebook/prop-types",  
 "licenseUrl": "https://github.com/facebook/prop-types/raw/main/LICENSE",

-----

"react-bootstrap@0.32.4"

"licenses": "MIT",  
 "repository": "https://github.com/react-bootstrap/react-bootstrap",  
 "licenseUrl": "https://github.com/react-bootstrap/react-bootstrap/raw/master/LICENSE",

-----

```

"rxjs@5.5.12"

  "licenses": "Apache-2.0",
  "repository": "https://github.com/reactivex/rxjs",
  "licenseUrl": "https://github.com/ReactiveX/rxjs/raw/master/
LICENSE.txt",
-----

"classnames@2.2.5"

  "licenses": "MIT",
  "repository": "https://github.com/JedWatson/classnames",
  "licenseUrl": "https://github.com/JedWatson/classnames/raw/master/
LICENSE",
-----

"classnames@2.3.1"

  "licenses": "MIT",
  "repository": "https://github.com/JedWatson/classnames",
  "licenseUrl": "https://github.com/JedWatson/classnames/raw/master/
LICENSE",
-----

"react-redux@7.2.4"

  "licenses": "MIT",
  "repository": "https://github.com/reduxjs/react-redux",
  "licenseUrl": "https://github.com/reduxjs/react-redux/raw/master/
LICENSE.md",
-----

"react-redux@7.2.8"

  "licenses": "MIT",
  "repository": "https://github.com/reduxjs/react-redux",
  "licenseUrl": "https://github.com/reduxjs/react-redux/raw/master/
LICENSE.md",
-----

"redux-form@8.3.8"

  "licenses": "MIT",
  "repository": "https://github.com/redux-form/redux-form",
  "licenseUrl": "https://github.com/redux-form/redux-form/raw/master/
LICENSE",
-----

"immutable@3.8.1"

  "licenses": "BSD-3-Clause",
  "repository": "https://github.com/facebook/immutable-js",
  "licenseUrl": "https://raw.githubusercontent.com/immutable-js/
immutable-js/e96d73f7e1fbef00d03b09aa4352e04de61abb3/LICENSE",
-----

"moment@2.29.4"

```

```

    "licenses": "MIT",
    "repository": "https://github.com/moment/moment",
    "licenseUrl": "https://github.com/moment/moment/raw/develop/LICENSE",
    -----

"graphql@14.7.0"

    "licenses": "MIT",
    "repository": "https://github.com/graphql/graphql-js",
    "licenseUrl": "https://github.com/graphql/graphql-js/raw/main/LICENSE",
    -----

"graphql@16.6.0"

    "licenses": "MIT",
    "repository": "https://github.com/graphql/graphql-js",
    "licenseUrl": "https://github.com/graphql/graphql-js/raw/main/LICENSE",
    -----

"lodash-es@14.7.0"

    "licenses": "MIT",
    "repository": "https://github.com/lodash/lodash",
    "licenseUrl": "https://github.com/lodash/lodash/raw/master/LICENSE",
    -----

"react-dom@16.14.0"

    "licenses": "MIT",
    "repository": "https://github.com/facebook/react",
    "licenseUrl": "https://github.com/facebook/react/raw/main/LICENSE",
    -----

"react-dom@17.0.2"

    "licenses": "MIT",
    "repository": "https://github.com/facebook/react",
    "licenseUrl": "https://github.com/facebook/react/raw/main/LICENSE",
    -----

"react-dom@18.3.1"

    "licenses": "MIT",
    "repository": "https://github.com/facebook/react",
    "licenseUrl": "https://github.com/facebook/react/raw/main/LICENSE",
    -----

"antlr4@4.8.0"

    "licenses": "BSD-3-Clause",
    "repository": "https://github.com/antlr/antlr4",
    "licenseUrl": "https://github.com/antlr/antlr4/raw/master/LICENSE.txt",
    -----

"react-router@6.15.0"

```

```

    "licenses": "MIT",
    "repository": "https://github.com/remix-run/react-router",
    "licenseUrl": "https://github.com/remix-run/react-router/raw/main/LICENSE.md",
-----

"react-router-dom@4.2.2"

    "licenses": "MIT",
    "repository": "https://github.com/remix-run/react-router",
    "licenseUrl": "https://github.com/remix-run/react-router/raw/main/LICENSE.md",
-----

"react-router-dom@5.2.0"

    "licenses": "MIT",
    "repository": "https://github.com/remix-run/react-router",
    "licenseUrl": "https://github.com/remix-run/react-router/raw/main/LICENSE.md",
-----

"react-router-dom@6.15.0"

    "licenses": "MIT",
    "repository": "https://github.com/remix-run/react-router",
    "licenseUrl": "https://github.com/remix-run/react-router/raw/main/LICENSE.md",
-----

"fbjs@0.8.16"

    "licenses": "MIT",
    "repository": "https://github.com/facebook/fbjs",
    "licenseUrl": "https://github.com/facebook/fbjs/raw/main/LICENSE",
-----

"redux@4.2.0"

    "licenses": "MIT",
    "repository": "https://github.com/reduxjs/redux",
    "licenseUrl": "https://github.com/reduxjs/redux/raw/master/LICENSE.md",
-----

"react-highcharts@16.1.0"

    "licenses": "MIT",
    "repository": "https://github.com/kirjs/react-highcharts",
    "licenseUrl": "https://github.com/kirjs/react-highcharts/raw/master/LICENSE",
-----

"lodash@4.17.21"

    "licenses": "MIT",
    "repository": "https://github.com/lodash/lodash",

```

```

    "licenseUrl": "https://github.com/lodash/lodash/raw/master/LICENSE" ,
-----

"highcharts@7.2.2"

    "licenses": "https://www.highcharts.com/license" ,
    "repository": "https://github.com/highcharts/highcharts-dist" ,
-----

"highcharts@9.1.0"

    "licenses": "https://www.highcharts.com/license" ,
    "repository": "https://github.com/highcharts/highcharts-dist" ,
-----

"pegjs@0.10.0"

    "licenses": "MIT" ,
    "repository": "https://github.com/pegjs/pegjs" ,
    "licenseUrl": "https://github.com/pegjs/pegjs/raw/master/LICENSE" ,
-----

"react-overlays@0.7.3"

    "licenses": "MIT" ,
    "repository": "https://github.com/react-bootstrap/react-overlays" ,
    "licenseUrl": "https://github.com/react-bootstrap/react-overlays/raw/
master/LICENSE" ,
-----

"jquery@3.6.1"

    "licenses": "MIT" ,
    "repository": "https://github.com/jquery/jquery" ,
    "licenseUrl": "https://github.com/jquery/jquery/raw/main/LICENSE.txt" ,
-----

"react-bootstrap-typeahead@1.4.2"

    "licenses": "MIT" ,
    "repository": "https://github.com/ericgio/react-bootstrap-typeahead" ,
    "licenseUrl": "https://github.com/ericgio/react-bootstrap-typeahead/raw/
main/LICENSE.md" ,
-----

"@reduxjs/toolkit@1.5.1"

    "licenses": "MIT" ,
    "repository": "https://github.com/reduxjs/redux-toolkit" ,
    "licenseUrl": "https://github.com/reduxjs/redux-toolkit/raw/master/
LICENSE" ,
-----

"@reduxjs/toolkit@1.8.2"

    "licenses": "MIT" ,
    "repository": "https://github.com/reduxjs/redux-toolkit" ,

```



```

    "licenseUrl": "https://github.com/reduxjs/redux-toolkit/raw/master/
LICENSE",
-----

"react-table@6.11.5"

  "licenses": "MIT",
  "repository": "https://github.com/TanStack/table",
  "licenseUrl": "https://github.com/TanStack/table/raw/main/LICENSE",
-----

"json-structure-validator@1.2.1"

  "licenses": "none",
  "repository": "https://github.com/AntJanus/JSON-structure-validator",
-----

"keycode@2.2.1"

  "licenses": "MIT",
  "repository": "https://github.com/timoxley/keycode",
  "licenseUrl": "https://github.com/timoxley/keycode/raw/master/LICENSE",
-----

"react-intl@2.4.0"

  "licenses": "BSD-3-Clause",
  "repository": "https://github.com/formatjs/formatjs",
-----

"intl-messageformat@2.1.0"

  "licenses": "BSD-3-Clause",
  "repository": "https://github.com/formatjs/formatjs",
-----

"intl-messageformat@9.6.16"

  "licenses": "BSD-3-Clause",
  "repository": "https://github.com/formatjs/formatjs",
-----

"rc-slider@8.3.1"

  "licenses": "MIT",
  "repository": "https://github.com/react-component/slider",
  "licenseUrl": "https://github.com/react-component/slider/raw/master/
LICENSE",
-----

"graphql-tag@2.12.6"

  "licenses": "MIT",
  "repository": "https://github.com/apollographql/graphql-tag",
  "licenseUrl": "https://github.com/apollographql/graphql-tag/raw/main/
LICENSE",

```

```

-----
"react-notification-system@0.2.15"

  "licenses": "MIT",
  "repository": "https://github.com/igorprado/react-notification-system",
  "licenseUrl": "https://github.com/igorprado/
react-notification-system/raw/master/LICENSE",
-----

"history@4.10.1"

  "licenses": "MIT",
  "repository": "https://github.com/remix-run/history",
  "licenseUrl": "https://github.com/remix-run/history/raw/dev/LICENSE",
-----

"rc-datetime-picker@4.10.1"

  "licenses": "MIT",
  "repository": "https://github.com/AllenWoouooo/rc-datetime-picker",
  "licenseUrl": "https://github.com/AllenWoouooo/rc-datetime-picker/raw/
master/LICENSE",
-----

"rc-tooltip@3.4.9"

  "licenses": "MIT",
  "repository": "https://github.com/react-component/tooltip",
  "licenseUrl": "https://github.com/react-component/tooltip/raw/master/
LICENSE",
-----

"react-addons-shallow-compare@15.6.3"

  "licenses": "MIT",
  "repository": "https://github.com/facebook/react",
  "licenseUrl": "https://github.com/facebook/react/raw/main/LICENSE",
-----

"react-router-bootstrap@0.25.0"

  "licenses": "Apache-2.0",
  "repository": "https://github.com/react-bootstrap/
react-router-bootstrap",
  "licenseUrl": "https://github.com/react-bootstrap/
react-router-bootstrap/raw/master/LICENSE",
-----

"redux-thunk@2.2.0"

  "licenses": "MIT",
  "repository": "https://github.com/reduxjs/redux-thunk",
  "licenseUrl": "https://github.com/reduxjs/redux-thunk/raw/master/
LICENSE.md",
-----

```

```

"apollo-boost@0.1.28"

  "licenses": "MIT",
  "repository": "https://github.com/apollographql/apollo-client",
  "licenseUrl": "https://github.com/apollographql/apollo-client/raw/main/
LICENSE",
-----

"redux-observable@0.18.0"

  "licenses": "MIT",
  "repository": "https://github.com/redux-observable/redux-observable",
  "licenseUrl": "https://github.com/redux-observable/redux-observable/raw/
master/LICENSE",
-----

"react-router-redux@4.0.8"

  "licenses": "MIT",
  "repository": "https://github.com/reactjs/react-router-redux",
  "licenseUrl": "https://github.com/reactjs/react-router-redux/raw/master/
LICENSE",
-----

"intl@1.2.5"

  "licenses": "MIT",
  "repository": "https://github.com/andyearnshaw/Intl.js",
  "licenseUrl": "https://github.com/andyearnshaw/Intl.js/raw/master/
LICENSE.txt",
-----

"babel-polyfill@6.26.0"

  "licenses": "MIT",
  "repository": "https://github.com/babel/babel/tree/master/packages/
babel-polyfill",
  "licenseUrl": "https://github.com/babel/babel/raw/main/LICENSE",
-----

"@babel-runtime@7.21.0"

  "licenses": "MIT",
  "repository": "https://github.com/babel/babel/tree/main/packages/
babel-runtime",
  "licenseUrl": "https://raw.githubusercontent.com/babel/babel/main/
LICENSE",
-----

"whatwg-fetch@2.0.3"

  "licenses": "MIT",
  "repository": "https://github.com/github/fetch",
  "licenseUrl": "https://github.com/github/fetch/raw/master/LICENSE",
-----

```

```

"react-text-mask@5.0.2"

  "licenses": "Unlicense",
  "repository": "https://github.com/text-mask/text-mask",
  "licenseUrl": "https://github.com/text-mask/text-mask/raw/master/
LICENSE",
-----

"react-select@1.3.0"

  "licenses": "MIT",
  "repository": "https://github.com/JedWatson/react-select/tree/master/
packages/react-select",
  "licenseUrl": "https://github.com/JedWatson/react-select/raw/master/
LICENSE",
-----

"react-dock@0.2.4"

  "licenses": "MIT",
  "repository": "https://github.com/reduxjs/redux-devtools",
  "licenseUrl": "https://github.com/reduxjs/redux-devtools/raw/main/
LICENSE.md",
-----

"css-toggle-switch@4.1.0"

  "licenses": "MIT",
  "repository": "https://github.com/ghinda/css-toggle-switch",
  "licenseUrl": "https://github.com/ghinda/css-toggle-switch/raw/master/
LICENSE",
-----

"dompurify@2.3.8"

  "licenses": "MPL-2.0 OR Apache-2.0",
  "repository": "https://github.com/cure53/DOMPurify",
  "licenseUrl": "https://github.com/cure53/DOMPurify/raw/main/LICENSE",
-----

"react-copy-to-clipboard@5.0.0"

  "licenses": "MIT",
  "repository": "https://github.com/nkbt/react-copy-to-clipboard",
  "licenseUrl": "https://github.com/nkbt/react-copy-to-clipboard/raw/
master/LICENSE",
-----

"react-duallist@1.1.6"

  "licenses": "MIT",
  "repository": "https://github.com/jyotirmaybanerjee/react-duallist",
  "licenseUrl": "https://github.com/jyotirmaybanerjee/react-duallist/raw/
master/LICENSE",
-----

"redux-devtools-extension@2.13.2"

```

```

    "licenses": "MIT",
    "repository": "https://github.com/zalmoxisus/redux-devtools-extension",
    "licenseUrl": "https://github.com/zalmoxisus/
redux-devtools-extension/raw/master/LICENSE",
-----

"axios-mock-adapter@1.19.0"

    "licenses": "MIT",
    "repository": "https://github.com/ctimmerm/axios-mock-adapter",
    "licenseUrl": "https://github.com/ctimmerm/axios-mock-adapter/raw/
master/LICENSE",
-----

"axios@0.21.1"

    "licenses": "MIT",
    "repository": "https://github.com/axios/axios",
    "licenseUrl": "https://github.com/axios/axios/raw/v0.x/LICENSE",
-----

"axios@0.27.2"

    "licenses": "MIT",
    "repository": "https://github.com/axios/axios",
    "licenseUrl": "https://github.com/axios/axios/raw/v0.x/LICENSE",
-----

"codemirror@5.62.2"

    "licenses": "MIT",
    "repository": "https://github.com/codemirror/basic-setup",
    "licenseUrl": "https://github.com/codemirror/basic-setup/raw/main/
LICENSE",
-----

"codemirror@5.65.12"

    "licenses": "MIT",
    "repository": "https://github.com/codemirror/basic-setup",
    "licenseUrl": "https://github.com/codemirror/basic-setup/raw/main/
LICENSE",
-----

"grommet-icons@4.9.0"

    "licenses": "Apache-2.0",
    "repository": "https://github.com/grommet/grommet-icons",
    "licenseUrl": "https://github.com/grommet/grommet-icons/raw/master/
LICENSE",
-----

"grommet-icons@4.12.1"

    "licenses": "Apache-2.0",
    "repository": "https://github.com/grommet/grommet-icons",

```

```

    "licenseUrl": "https://github.com/grommet/grommet-icons/raw/master/
LICENSE",
-----

"grommet@2.25.1"

    "licenses": "Apache-2.0",
    "repository": "https://github.com/grommet/grommet",
    "licenseUrl": "https://github.com/grommet/grommet/raw/master/LICENSE",
-----

"grommet@2.39.0"

    "licenses": "Apache-2.0",
    "repository": "https://github.com/grommet/grommet",
    "licenseUrl": "https://github.com/grommet/grommet/raw/master/LICENSE",
-----

"highcharts-react-official@3.0.0"

    "licenses": "https://github.com/highcharts/highcharts-react/raw/master/
LICENSE",
    "repository": "https://github.com/highcharts/highcharts-react",
-----

"react-codemirror2@7.2.1"

    "licenses": "MIT",
    "repository": "https://github.com/scniro/react-codemirror2",
    "licenseUrl": "https://github.com/scniro/react-codemirror2/raw/master/
LICENSE",
-----

"styled-components@5.3.0"

    "licenses": "MIT",
    "repository": "https://github.com/styled-components/styled-components",
    "licenseUrl": "https://github.com/styled-components/
styled-components/raw/main/LICENSE",
-----

"styled-components@5.3.9"

    "licenses": "MIT",
    "repository": "https://github.com/styled-components/styled-components",
    "licenseUrl": "https://github.com/styled-components/
styled-components/raw/main/LICENSE",
-----

"grommet-theme-hpe@3.2.1"

    "licenses": "Apache-2.0",
    "repository": "https://github.com/grommet/grommet-theme-hpe",
    "licenseUrl": "https://github.com/grommet/grommet-theme-hpe/raw/master/
LICENSE",
-----

```

```

"uuid@8.3.2"

  "licenses": "MIT",
  "repository": "https://github.com/uuidjs/uuid",
  "licenseUrl": "https://github.com/uuidjs/uuid/raw/main/LICENSE.md",
-----

"uuid@9.0.0"

  "licenses": "MIT",
  "repository": "https://github.com/uuidjs/uuid",
  "licenseUrl": "https://github.com/uuidjs/uuid/raw/main/LICENSE.md",
-----

"use-debounce@7.0.1"

  "licenses": "MIT",
  "repository": "https://github.com/xnimorz/use-debounce",
  "licenseUrl": "https://github.com/xnimorz/use-debounce/raw/master/
LICENSE",
-----

"deep-equal@1.0.1"

  "licenses": "MIT",
  "repository": "https://github.com/inspect-js/node-deep-equal",
  "licenseUrl": "https://github.com/inspect-js/node-deep-equal/raw/master/
LICENSE",
-----

"react-d3@0.4.0"

  "licenses": "MIT",
  "repository": "https://github.com/esbullington/react-d3",
  "licenseUrl": "https://github.com/esbullington/react-d3/raw/master/
LICENSE.md",
-----

"react-immutable-proptypes@2.1.0"

  "licenses": "MIT",
  "repository": "https://github.com/HurricaneJames/
react-immutable-proptypes",
  "licenseUrl": "https://github.com/HurricaneJames/
react-immutable-proptypes/raw/master/LICENSE",
-----

"swagger-ui-dist@3.23.11"

  "licenses": "Apache-2.0",
  "repository": "https://github.com/swagger-api/swagger-ui",
  "licenseUrl": "https://github.com/swagger-api/swagger-ui/raw/master/
LICENSE",
-----

"swagger-ui-themes@3.0.0"

```

```

    "licenses": "MIT",
    "repository": "https://github.com/ostranme/swagger-ui-themes",
    -----

"deepmerge@4.3.1"

    "licenses": "MIT",
    "repository": "https://github.com/TehShrike/deepmerge",
    "licenseUrl": "https://raw.githubusercontent.com/TehShrike/deepmerge/
master/license.txt",
    -----

"exenv@1.2.2"

    "licenses": "BSD",
    "repository": "https://github.com/JedWatson/exenv",
    "licenseUrl": "https://raw.githubusercontent.com/JedWatson/exenv/master/
LICENSE",
    -----

"grommet-styles@0.2.0"

    "licenses": "Apache-2.0",
    "repository": "https://github.com/grommet/grommet-styles",
    "licenseUrl": "https://raw.githubusercontent.com/grommet/grommet-styles/
master/LICENSE",
    -----

"hoist-non-react-statics@3.3.2"

    "licenses": "BSD",
    "repository": "https://github.com/mridgway/hoist-non-react-statics",
    "licenseUrl": "https://raw.githubusercontent.com/mridgway/
hoist-non-react-statics/master/LICENSE.md",
    -----

"object-assign@4.1.1"

    "licenses": "MIT",
    "repository": "https://github.com/sindresorhus/object-assign",
    "licenseUrl": "https://raw.githubusercontent.com/sindresorhus/
object-assign/main/license",
    -----

"react-fast-compare@3.2.1"

    "licenses": "MIT",
    "repository": "https://github.com/FormidableLabs/react-fast-compare",
    "licenseUrl": "https://raw.githubusercontent.com/FormidableLabs/
react-fast-compare/master/LICENSE",
    -----

"react-is@18.2.0"

    "licenses": "MIT",
    "repository": "https://github.com/facebook/react",
    "licenseUrl": "https://raw.githubusercontent.com/facebook/react/main/

```



```

LICENSE" ,

-----

"react-joyride@2.5.3"

  "licenses": "MIT",
  "repository": "https://github.com/gilbarbara/react-joyride",
  "licenseUrl": "https://raw.githubusercontent.com/gilbarbara/
react-joyride/main/LICENSE" ,

-----

"react-proptype-conditional-require@1.0.4"

  "licenses": "MIT",
  "repository": "https://github.com/beefancohen/
react-proptype-conditional-require" ,
  "licenseUrl": "https://raw.githubusercontent.com/beefancohen/
react-proptype-conditional-require/master/LICENSE" ,

-----

"@tanstack/react-query@5.32.0"

  "licenses": "MIT",
  "repository": "https://github.com/TanStack/query",
  "licenseUrl": "https://raw.githubusercontent.com/TanStack/query/main/
LICENSE" ,

-----

"react-side-effect@2.1.2"

  "licenses": "MIT",
  "repository": "https://github.com/gaearon/react-side-effect",
  "licenseUrl": "https://raw.githubusercontent.com/gaearon/
react-side-effect/master/LICENSE" ,

-----

"scheduler@0.23.2"

  "licenses": "MIT",
  "repository": "https://github.com/facebook/react",
  "licenseUrl": "https://github.com/facebook/react/raw/main/LICENSE" ,

-----

"scroll@3.0.1"

  "licenses": "MIT",
  "repository": "https://github.com/michaelrhodes/scroll",
  "licenseUrl": "https://raw.githubusercontent.com/michaelrhodes/scroll/
master/LICENSE" ,

-----

"scrollparent@2.0.1"

  "licenses": "MIT",
  "repository": "https://github.com/olahol/scrollparent.js",
  "licenseUrl": "https://raw.githubusercontent.com/olahol/scrollparent.js/
master/LICENSE" ,

```

```
-----  
"shallowequal@1.1.0"
```

```
  "licenses": "MIT",  
  "repository": "https://github.com/dashed/shallowequal",  
  "licenseUrl": "https://raw.githubusercontent.com/dashed/shallowequal/  
master/LICENSE",
```

```
-----  
"@mswjs/cookies@0.2.2"
```

```
  "licenses": "MIT",  
  "repository": "https://github.com/mswjs/cookies",  
  "licenseUrl": "https://raw.githubusercontent.com/mswjs/cookies/main/  
LICENSE.md",
```

```
-----  
"@open-draft/until@1.0.3"
```

```
  "licenses": "MIT",  
  "repository": "https://github.com/open-draft/until",  
  "licenseUrl": "https://raw.githubusercontent.com/open-draft/until/main/  
LICENSE",
```

```
-----  
"@xmldom/xmldom@0.8.7"
```

```
  "licenses": "MIT",  
  "repository": "https://github.com/xmldom/xmldom",  
  "licenseUrl": "https://raw.githubusercontent.com/xmldom/xmldom/master/  
LICENSE",
```

```
-----  
"available-typed-arrays@1.0.5"
```

```
  "licenses": "MIT",  
  "repository": "https://github.com/inspect-js/available-typed-arrays",  
  "licenseUrl": "https://raw.githubusercontent.com/inspect-js/  
available-typed-arrays/main/LICENSE",
```

```
-----  
"base64-js@1.5.1"
```

```
  "licenses": "MIT",  
  "repository": "https://github.com/beatgammit/base64-js",  
  "licenseUrl": "https://raw.githubusercontent.com/beatgammit/base64-js/  
master/LICENSE",
```

```
-----  
"buffer@6.0.3"
```

```
  "licenses": "MIT",  
  "repository": "https://github.com/feross/buffer",  
  "licenseUrl": "https://raw.githubusercontent.com/feross/buffer/master/  
LICENSE",
```

```

-----
"call-bind@1.0.2"
  "licenses": "MIT",
  "repository": "https://github.com/ljharb/call-bind",
  "licenseUrl": "https://raw.githubusercontent.com/ljharb/call-bind/main/LICENSE",
-----

"cookie@0.4.2"
  "licenses": "MIT",
  "repository": "https://github.com/jshttp/cookie",
  "licenseUrl": "https://raw.githubusercontent.com/jshttp/cookie/master/LICENSE",
-----

"debug@4.3.4"
  "licenses": "MIT",
  "repository": "https://github.com/debug-js/debug",
  "licenseUrl": "https://raw.githubusercontent.com/debug-js/debug/master/LICENSE",
-----

"esprima@4.0.1"
  "licenses": "BSD-2-Clause",
  "repository": "https://github.com/jquery/esprima",
  "licenseUrl": "https://raw.githubusercontent.com/jquery/esprima/main/LICENSE.BSD",
-----

"events@3.3.0"
  "licenses": "MIT",
  "repository": "https://github.com/browserify/events",
  "licenseUrl": "https://raw.githubusercontent.com/browserify/events/main/LICENSE",
-----

"for-each@0.3.3"
  "licenses": "MIT",
  "repository": "https://github.com/Raynos/for-each",
  "licenseUrl": "https://raw.githubusercontent.com/Raynos/for-each/master/LICENSE",
-----

"function-bind@1.1.1"
  "licenses": "MIT",
  "repository": "https://github.com/Raynos/function-bind",
  "licenseUrl": "https://raw.githubusercontent.com/Raynos/function-bind/master/LICENSE",
-----

```

```

"get-intrinsic@1.2.0"

  "licenses": "MIT",
  "repository": "https://github.com/ljharb/get-intrinsic",
  "licenseUrl": "https://raw.githubusercontent.com/ljharb/get-intrinsic/
main/LICENSE",
-----

"lodash@4.17.21"

  "licenses": "MIT",
  "repository": "https://github.com/lodash/lodash",
  "licenseUrl": "https://raw.githubusercontent.com/lodash/lodash/main/
LICENSE",
-----

"has-symbols@1.0.3"

  "licenses": "MIT",
  "repository": "https://github.com/inspect-js/has-symbols",
  "licenseUrl": "https://raw.githubusercontent.com/inspect-js/has-symbols/
main/LICENSE",
-----

"has-tostringtag@1.0.0"

  "licenses": "MIT",
  "repository": "https://github.com/inspect-js/has-tostringtag",
  "licenseUrl": "https://raw.githubusercontent.com/inspect-js/
has-tostringtag/main/LICENSE",
-----

"has@1.0.3"

  "licenses": "MIT",
  "repository": "https://github.com/tarruda/has",
  "licenseUrl": "https://raw.githubusercontent.com/tarruda/has/master/
LICENSE-MIT",
-----

"headers-polyfill@3.1.2"

  "licenses": "MIT",
  "repository": "https://github.com/mswjs/headers-polyfill",
  "licenseUrl": "https://raw.githubusercontent.com/mswjs/headers-polyfill/
main/LICENSE",
-----

"ieee754@1.2.1"

  "licenses": "BSD-3-Clause",
  "repository": "https://github.com/feross/ieee754",
  "licenseUrl": "https://raw.githubusercontent.com/feross/ieee754/master/
LICENSE",
-----

```

```

"inherits@2.0.4"

  "licenses": "ISC",
  "repository": "https://github.com/isaacs/inherits",
  "licenseUrl": "https://raw.githubusercontent.com/isaacs/inherits/main/LICENSE",
  -----

"is-arguments@1.1.1"

  "licenses": "MIT",
  "repository": "https://github.com/inspect-js/is-arguments",
  "licenseUrl": "https://raw.githubusercontent.com/inspect-js/is-arguments/main/LICENSE",
  -----

"is-callable@1.2.7"

  "licenses": "MIT",
  "repository": "https://github.com/inspect-js/is-callable",
  "licenseUrl": "https://raw.githubusercontent.com/inspect-js/is-callable/main/LICENSE",
  -----

"is-generator-function@1.0.10"

  "licenses": "MIT",
  "repository": "https://github.com/inspect-js/is-generator-function",
  "licenseUrl": "https://raw.githubusercontent.com/inspect-js/is-generator-function/main/LICENSE",
  -----

"is-node-process@1.2.0"

  "licenses": "MIT",
  "repository": "https://github.com/mswjs/is-node-process",
  -----

"is-typed-array@1.1.10"

  "licenses": "MIT",
  "repository": "https://github.com/inspect-js/is-typed-array",
  "licenseUrl": "https://raw.githubusercontent.com/inspect-js/is-typed-array/main/LICENSE",
  -----

"js-levenshtein@1.1.6"

  "licenses": "MIT",
  "repository": "https://github.com/gustf/js-levenshtein",
  "licenseUrl": "https://github.com/gustf/js-levenshtein/blob/master/LICENSE",
  -----

"js-yaml@3.14.1"

  "licenses": "MIT",

```

```

    "repository": "https://github.com/nodeca/js-yaml",
    "licenseUrl": "https://raw.githubusercontent.com/nodeca/js-yaml/master/
LICENSE",
-----

"ms@2.1.2"

    "licenses": "MIT",
    "repository": "https://github.com/vercel/ms",
    "licenseUrl": "https://raw.githubusercontent.com/vercel/ms/master/
license.md",
-----

"msw@1.2.1"

    "licenses": "MIT",
    "repository": "https://github.com/mswjs/msw",
    "licenseUrl": "https://raw.githubusercontent.com/mswjs/msw/main/
LICENSE.md",
-----

"node-fetch@2.6.9"

    "licenses": "MIT",
    "repository": "https://github.com/node-fetch/node-fetch",
    "licenseUrl": "https://raw.githubusercontent.com/node-fetch/node-fetch/
main/LICENSE.md",
-----

"outvariant@1.4.0"

    "licenses": "MIT",
    "repository": "https://github.com/open-draft/outvariant",
    "licenseUrl": "https://raw.githubusercontent.com/open-draft/outvariant/
main/LICENSE",
-----

"path-to-regexp@6.2.1"

    "licenses": "MIT",
    "repository": "https://github.com/pillarjs/path-to-regexp",
    "licenseUrl": "https://raw.githubusercontent.com/pillarjs/
path-to-regexp/master/LICENSE",
-----

"set-cookie-parser@2.6.0"

    "licenses": "MIT",
    "repository": "https://github.com/nfriedly/set-cookie-parser",
    "licenseUrl": "https://raw.githubusercontent.com/nfriedly/
set-cookie-parser/master/LICENSE",
-----

"strict-event-emitter@0.4.6"

    "licenses": "MIT",
    "repository": "https://github.com/open-draft/strict-event-emitter",

```

```

-----
"util@0.12.5"
  "licenses": "MIT",
  "repository": "https://github.com/browserify/node-util",
  "licenseUrl": "https://raw.githubusercontent.com/browserify/node-util/
master/LICENSE",
-----

"web-encoding@1.1.5"
  "licenses": "MIT",
  "repository": "https://github.com/gozala/web-encoding",
-----

"which-typed-array@1.1.9"
  "licenses": "MIT",
  "repository": "https://github.com/inspect-js/which-typed-array",
  "licenseUrl": "https://raw.githubusercontent.com/inspect-js/
which-typed-array/main/LICENSE",
-----

"react-toastify@9.1.2"
  "licenses": "MIT",
  "repository": "https://github.com/fkhadra/react-toastify",
  "licenseUrl": "https://raw.githubusercontent.com/fkhadra/react-toastify/
main/LICENSE",
-----

"react-syntax-highlighter@15.5.0"
  "licenses": "MIT",
  "repository": "https://github.com/react-syntax-highlighter/
react-syntax-highlighter",
  "licenseUrl": "https://github.com/react-syntax-highlighter/
react-syntax-highlighter/blob/master/LICENSE",
-----

"character-entities@1.2.4"
  "licenses": "MIT",
  "repository": "https://github.com/wooorm/character-entities",
  "licenseUrl": "https://github.com/wooorm/character-entities/blob/main/
license",
-----

"character-entities-legacy@1.1.4"
  "licenses": "MIT",
  "repository": "https://github.com/wooorm/character-entities-legacy",
  "licenseUrl": "https://github.com/wooorm/character-entities-legacy/blob/
main/license",
-----

```

```
"character-reference-invalid@1.1.4"
```

```
  "licenses": "MIT",
  "repository": "https://github.com/woorm/character-reference-invalid",
  "licenseUrl": "https://github.com/woorm/character-reference-invalid/
blob/main/license",
```

```
-----
```

```
"comma-separated-tokens@1.0.8"
```

```
  "licenses": "MIT",
  "repository": "https://github.com/woorm/comma-separated-tokens",
  "licenseUrl": "https://github.com/woorm/comma-separated-tokens/blob/
main/license",
```

```
-----
```

```
"hast-util-parse-selector@2.2.5"
```

```
  "licenses": "MIT",
  "repository": "https://github.com/syntax-tree/hast-util-parse-selector",
  "licenseUrl": "https://github.com/syntax-tree/hast-util-parse-selector/
blob/main/license",
```

```
-----
```

```
"hastscript@6.0.0"
```

```
  "licenses": "MIT",
  "repository": "https://github.com/syntax-tree/hastscript",
  "licenseUrl": "https://github.com/syntax-tree/hastscript/blob/main/
license",
```

```
-----
```

```
"is-alphabetical@1.0.4"
```

```
  "licenses": "MIT",
  "repository": "https://github.com/woorm/is-alphabetical",
  "licenseUrl": "https://github.com/woorm/is-alphabetical/blob/main/
license",
```

```
-----
```

```
"is-alphanumeric@1.0.4"
```

```
  "licenses": "MIT",
  "repository": "https://github.com/woorm/is-alphanumeric",
  "licenseUrl": "https://github.com/woorm/is-alphanumeric/blob/main/
license",
```

```
-----
```

```
"is-decimal@1.0.4"
```

```
  "licenses": "MIT",
  "repository": "https://github.com/woorm/is-decimal",
  "licenseUrl": "https://github.com/woorm/is-decimal/blob/main/license",
```

```
-----
```

```
"is-hexadecimal@1.0.4"
```



```

    "licenses": "MIT",
    "repository": "https://github.com/woorm/is-hexadecimal",
    "licenseUrl": "https://github.com/woorm/is-hexadecimal/blob/main/
license",
-----

"parse-entities@2.0.0"

    "licenses": "MIT",
    "repository": "https://github.com/woorm/parse-entities",
    "licenseUrl": "https://github.com/woorm/parse-entities/blob/main/
license",
-----

"prismjs@1.29.0"

    "licenses": "MIT",
    "repository": "https://github.com/PrismJS/prism",
    "licenseUrl": "https://github.com/PrismJS/prism/blob/master/LICENSE",
-----

"property-information@5.6.0"

    "licenses": "MIT",
    "repository": "https://github.com/woorm/property-information",
    "licenseUrl": "https://github.com/woorm/property-information/blob/main/
license",
-----

"refractor@3.6.0"

    "licenses": "MIT",
    "repository": "https://github.com/woorm/refractor",
    "licenseUrl": "https://github.com/woorm/refractor/blob/main/license",
-----

"space-separated-tokens@1.1.5"

    "licenses": "MIT",
    "repository": "https://github.com/woorm/space-separated-tokens",
    "licenseUrl": "https://github.com/woorm/space-separated-tokens/blob/
main/license",
-----

"xtend@4.0.2"

    "licenses": "MIT",
    "repository": "https://github.com/Raynos/xtend",
    "licenseUrl": "https://github.com/Raynos/xtend/blob/master/LICENSE",
-----

"dayjs@1.11.10"

    "licenses": "MIT",
    "repository": "https://github.com/iamkun/dayjs",
    "licenseUrl": "https://github.com/iamkun/dayjs#MIT-1-ov-file",

```

-----  
commons-beanutils

Copyright (c) 2009 The Apache Software Foundation.

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
commons-configuration

Copyright (c) 2009 The Apache Software Foundation.

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
joda-time

Copyright (c) 2009 The Apache Software Foundation.

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
jna

Copyright (c) 2009 The Apache Software Foundation.

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
commons-lang

Copyright (c) 2009 The Apache Software Foundation.

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
ehcache-core

Copyright (c) 2009 The Apache Software Foundation.

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

-----  
annotations

License: GNU Lesser Public License  
<http://www.gnu.org/licenses/lgpl.html>

`hazelcast`

Copyright (c) 2009 The Apache Software Foundation.

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

---

`jersey-server`

License: CDDL+GPL License  
[http://glassfish.java.net/public/CDDL+GPL\\_1\\_1.html](http://glassfish.java.net/public/CDDL+GPL_1_1.html)

---

`libpam4j`

License: The MIT license  
<http://www.opensource.org/licenses/mit-license.php>

---

`lombok`

License: The MIT License  
<https://projectlombok.org/LICENSE>

---

`spring-security-core`

Copyright (c) 2009 The Apache Software Foundation.

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

---

`spring-security-kerberos-core`

Copyright (c) 2009 The Apache Software Foundation.

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

---

`swagger-annotations`

Copyright (c) 2009 The Apache Software Foundation.

License: Apache License, Version 2.0  
<http://www.apache.org/licenses/LICENSE-2.0.html>

---

`Apache Ranger`

Copyright 2014-2022 The Apache Software Foundation

License: Apache License Version 2.0, January 2004  
<http://www.apache.org/licenses/LICENSE-2.0>

-----  
Apache NiFi

Copyright 2014-2022 The Apache Software Foundation

License: Apache License Version 2.0, January 2004  
<http://www.apache.org/licenses/LICENSE-2.0>

-----  
Apache Airflow

Copyright 2016-2021 The Apache Software Foundation

License: Apache License Version 2.0, January 2004  
<http://www.apache.org/licenses/LICENSE-2.0>

=====

## Apache License

Apache License  
Version 2.0, January 2004  
<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

### 1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but

not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:
  - (a) You must give any other recipients of the Work or Derivative Works a copy of this License; and
  - (b) You must cause any modified files to carry prominent notices stating that You changed the files; and
  - (c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
  - (d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.
8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.
9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[ ]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License");  
you may not use this file except in compliance with the License.  
You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software

distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

=====

### MIT License

The MIT License

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

=====

### License for uuid

License for uuid:

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, and the entire permission notice in its entirety, including the disclaimer of warranties.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, ALL OF WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE



LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF NOT ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## License for JavaMail

License for JavaMail:

Oracle Corporation ("ORACLE") ENTITLEMENT for SOFTWARE

Licensee/Company: Entity receiving Software.

Effective Date: Date of delivery of the Software to You.

Software: JavaMail 1.4.4

License Term: Perpetual (subject to termination under the SLA).

Licensed Unit: Software Copy.

Licensed unit Count: Unlimited.

Permitted Uses:

1. You may reproduce and use the Software for Your own Individual, Commercial and Research and Instructional Use only for the purposes of designing, developing, testing, and running Your applets and applications ("Programs").

2. Subject to the terms and conditions of this Agreement and restrictions and exceptions set forth in the Software's documentation, You may reproduce and distribute portions of Software identified as a redistributable in the documentation (each a "Redistributable"), provided that You comply with the following (note that You may be entitled to reproduce and distribute other portions of the Software not defined in the documentation as a Redistributable under certain other licenses as described in the THIRDPARTYLICENSEREADME, if applicable):

(a) You distribute Redistributable complete and unmodified and only bundled as part of Your Programs,

(b) Your Programs add significant and primary functionality to the Redistributable,

(c) You distribute Redistributable for the sole purpose of running Your Programs,

(d) You do not distribute additional software intended to replace any component(s) of the Redistributable,

(e) You do not remove or alter any proprietary legends or notices contained in or on the Redistributable.

(f) You only distribute the Redistributable subject to a license agreement that protects Oracle's interests consistent with the terms contained in this Agreement, and

(g) You agree to defend and indemnify Oracle and its licensors from and against any damages, costs, liabilities, settlement amounts and/or expenses (including attorneys' fees) incurred in connection with any claim, lawsuit or action by any third party that arises or results from the use or distribution of any and all Programs and/or Redistributable.

3. Java Technology Restrictions. You may not create, modify, or change the behavior of, or authorize Your licensees to create, modify, or change the behavior of, classes, interfaces, or subpackages that are in any way identified as "java", "javax", "sun" or similar convention as specified by Oracle in any naming convention designation.

4. No Diagnostic, Maintenance, Repair or Technical Support Services. The scope of Your license does not include any right, express or implied, (i) to access, copy, distribute, display or use the Software to provide diagnostic, maintenance, repair or technical support services for Oracle software or Oracle hardware on behalf of any third party for Your direct or indirect commercial gain or advantage, without Oracle's prior written authorization, or (ii) for any third party to access, copy, distribute, display or use the Software to provide diagnostic, maintenance, repair or technical support services for Oracle software or Oracle hardware on Your behalf for such party's direct or indirect commercial gain or advantage, without Oracle's prior written authorization. The limitations set forth in this paragraph apply to any and all error corrections, patches, updates, and upgrades to the Software You may receive, access, download or otherwise obtain from Oracle.

5. Records and Documentation. During the term of the SLA and Entitlement, and for a period of three (3) years thereafter, You agree to keep proper records and documentation of Your compliance with the SLA and Entitlement. Upon Oracle's reasonable request, You will provide copies of such records and documentation to Oracle for the purpose of confirming Your compliance with the terms and conditions of the SLA and Entitlement. This section will survive any termination of the SLA and Entitlement. You may terminate this SLA and Entitlement at any time by destroying all copies of the Software in which case the obligations set forth in Section 7 of the SLA shall apply.

Oracle Corporation ("ORACLE")

## SOFTWARE LICENSE AGREEMENT

READ THE TERMS OF THIS AGREEMENT ("AGREEMENT") CAREFULLY BEFORE OPENING SOFTWARE MEDIA PACKAGE. BY OPENING SOFTWARE MEDIA PACKAGE, YOU AGREE TO THE TERMS OF THIS AGREEMENT. IF YOU ARE ACCESSING SOFTWARE ELECTRONICALLY, INDICATE YOUR ACCEPTANCE OF THESE TERMS BY SELECTING THE "ACCEPT" BUTTON AT THE END OF THIS AGREEMENT. IF YOU DO NOT AGREE TO ALL OF THE TERMS, PROMPTLY RETURN THE UNUSED SOFTWARE TO YOUR PLACE OF PURCHASE FOR A REFUND OR, IF SOFTWARE IS ACCESSED ELECTRONICALLY, SELECT THE "DECLINE" (OR "EXIT") BUTTON AT THE END OF THIS AGREEMENT. IF YOU HAVE SEPARATELY AGREED TO LICENSE TERMS ("MASTER TERMS") FOR YOUR LICENSE TO THIS SOFTWARE, THEN SECTIONS 1-6 OF THIS AGREEMENT ("SUPPLEMENTAL LICENSE TERMS") SHALL SUPPLEMENT AND SUPERSEDE THE MASTER TERMS IN RELATION TO THIS SOFTWARE.

## 1. Definitions.

(a) "Entitlement" means the collective set of applicable documents authorized by Oracle evidencing your obligation to pay associated fees (if any) for the license, associated Services, and the authorized scope of use of Software under this Agreement.

(b) "Licensed Unit" means the unit of measure by which your use of Software and/or Service is licensed, as described in your Entitlement.

(c) "Permitted Use" means the licensed Software use(s) authorized in this Agreement as specified in your Entitlement. The Permitted Use for any bundled Oracle software not specified in your Entitlement will be evaluation use as provided in Section 3.

(d) "Service" means the service(s) that Oracle or its delegate will provide, if any, as selected in your Entitlement and as further described in the applicable service listings at [www.sun.com/service/servicelist](http://www.sun.com/service/servicelist).

(e) "Software" means the Oracle software described in your Entitlement. Also, certain software may be included for evaluation use under Section 3.

(f) "You" and "Your" means the individual or legal entity specified in the Entitlement, or for evaluation purposes, the entity performing the evaluation.

## 2. License Grant and Entitlement.

Subject to the terms of your Entitlement, Oracle grants you a nonexclusive, nontransferable limited license to use Software for its Permitted Use for the license term. Your Entitlement will specify (a) Software licensed, (b) the Permitted Use, (c) the license term, and (d) the Licensed Units.

Additionally, if your Entitlement includes Services, then it will also specify the (e) Service and (f) service term.

If your rights to Software or Services are limited in duration and the date such rights begin is other than the purchase date, your Entitlement will provide that beginning date(s).

The Entitlement may be delivered to you in various ways depending on the manner in which you obtain Software and Services, for example, the Entitlement may be provided in your receipt, invoice or your contract with Oracle or authorized Oracle reseller. It may also be in electronic format if you download Software.

### 3. Permitted Use.

As selected in your Entitlement, one or more of the following Permitted Uses will apply to your use of Software. Unless you have an Entitlement that expressly permits it, you may not use Software for any of the other Permitted Uses. If you don't have an Entitlement, or if your Entitlement doesn't cover additional software delivered to you, then such software is for your Evaluation Use.

(a) Evaluation Use. You may evaluate Software internally for a period of 90 days from your first use.

(b) Research and Instructional Use. You may use Software internally to design, develop and test, and also to provide instruction on such uses.

(c) Individual Use. You may use Software internally for personal, individual use.

(d) Commercial Use. You may use Software internally for your own commercial purposes.

(e) Service Provider Use. You may make Software functionality accessible (but not by providing Software itself or through outsourcing services) to your end users in an extranet deployment, but not to your affiliated companies or to government agencies.

### 4. Licensed Units.

Your Permitted Use is limited to the number of Licensed Units stated in your Entitlement. If you require additional Licensed Units, you will need additional Entitlement(s).

### 5. Restrictions.

(a) The copies of Software provided to you under this Agreement are licensed, not sold, to you by Oracle. Oracle reserves all rights not expressly granted. (b) You may make a single archival copy of Software, but otherwise may not copy, modify, or distribute Software. However if the Oracle documentation accompanying Software lists specific portions of

Software, such as header files, class libraries, reference source code, and/or redistributable files, that may be handled differently, you may do so only as provided in the Oracle documentation. (c) You may not rent, lease, lend or encumber Software. (d) Unless enforcement is prohibited by applicable law, you may not decompile, or reverse engineer Software. (e) The terms and conditions of this Agreement will apply to any Software updates, provided to you at Oracle's discretion, that replace and/or supplement the original Software, unless such update contains a separate license. (f) You may not publish or provide the results of any benchmark or comparison tests run on Software to any third party without the prior written consent of Oracle. (g) Software is confidential and copyrighted. (h) Unless otherwise specified, if Software is delivered with embedded or bundled software that enables functionality of Software, you may not use such software on a stand-alone basis or use any portion of such software to interoperate with any program(s) other than Software. (i) Software may contain programs that perform automated collection of system data and/or automated software updating services. System data collected through such programs may be used by Oracle, its subcontractors, and its service delivery partners for the purpose of providing you with remote system services and/or improving Oracle's software and systems. (j) Software is not designed, licensed or intended for use in the design, construction, operation or maintenance of any nuclear facility and Oracle and its licensors disclaim any express or implied warranty of fitness for such uses. (k) No right, title or interest in or to any trademark, service mark, logo or trade name of Oracle or its licensors is granted under this Agreement.

## 6. Java Compatibility and Open Source.

Software may contain Java technology. You may not create additional classes to, or modifications of, the Java technology, except under compatibility requirements available under a separate agreement available at [www.java.net](http://www.java.net).

Oracle supports and benefits from the global community of open source developers, and thanks the community for its important contributions and open standards-based technology, which Oracle has adopted into many of its products.

Please note that portions of Software may be provided with notices and open source licenses from such communities and third parties that govern the use of those portions, and any licenses granted hereunder do not alter any rights and obligations you may have under such open source licenses, however, the disclaimer of warranty and limitation of liability provisions in this Agreement will apply to all Software in this distribution.

## 7. Term and Termination.

The license and service term are set forth in your Entitlement(s). Your rights under this Agreement will terminate immediately without notice from Oracle if you materially breach it or take any action in derogation of Oracle's and/or its licensors' rights to Software. Oracle may terminate this Agreement should any Software become, or in Oracle's reasonable opinion likely to become, the subject of a claim of intellectual property infringement or trade secret misappropriation. Upon termination, you will cease use of, and destroy, Software and confirm compliance in writing to Oracle. Sections 1, 5, 6, 7, and 9-15 will

survive termination of the Agreement.

8. Limited Warranty.

Oracle warrants to you that for a period of 90 days from the date of purchase, as evidenced by a copy of the receipt, the media on which Software is furnished (if any) will be free of defects in materials and workmanship under normal use. Except for the foregoing, Software is provided "AS IS". Your exclusive remedy and Oracle's entire liability under this limited warranty will be at Oracle's option to replace Software media or refund the fee paid for Software. Some states do not allow limitations on certain implied warranties, so the above may not apply to you. This limited warranty gives you specific legal rights. You may have others, which vary from state to state.

9. Disclaimer of Warranty.

UNLESS SPECIFIED IN THIS AGREEMENT, ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT THESE DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

10. Limitation of Liability.

TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL ORACLE OR ITS LICENSORS BE LIABLE FOR ANY LOST REVENUE, PROFIT OR DATA, OR FOR SPECIAL, INDIRECT, CONSEQUENTIAL, INCIDENTAL OR PUNITIVE DAMAGES, HOWEVER CAUSED REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF OR RELATED TO THE USE OF OR INABILITY TO USE SOFTWARE, EVEN IF ORACLE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. In no event will Oracle's liability to you, whether in contract, tort (including negligence), or otherwise, exceed the amount paid by you for Software under this Agreement. The foregoing limitations will apply even if the above stated warranty fails of its essential purpose. Some states do not allow the exclusion of incidental or consequential damages, so some of the terms above may not be applicable to you.

11. Export Regulations.

All Software, documents, technical data, and any other materials delivered under this Agreement are subject to U.S. export control laws and may be subject to export or import regulations in other countries. You agree to comply strictly with these laws and regulations and acknowledge that you have the responsibility to obtain any licenses to export, re-export, or import as may be required after delivery to you.

12. U.S. Government Restricted Rights.

If Software is being acquired by or on behalf of the U.S. Government or by a U.S. Government prime contractor or subcontractor (at any tier), then the Government's rights in Software and accompanying documentation will be only as set forth in this Agreement; this is in accordance with 48 CFR 227.7201 through 227.7202-4 (for Department of Defense (DOD))

acquisitions) and with 48 CFR 2.101 and 12.212 (for non-DOD acquisitions).

#### 13. Governing Law.

Any action related to this Agreement will be governed by California law and controlling U.S. federal law. No choice of law rules of any jurisdiction will apply.

#### 14. Severability.

If any provision of this Agreement is held to be unenforceable, this Agreement will remain in effect with the provision omitted, unless omission would frustrate the intent of the parties, in which case this Agreement will immediately terminate.

#### 15. Integration.

This Agreement, including any terms contained in your Entitlement, is the entire agreement between you and Oracle relating to its subject matter. It supersedes all prior or contemporaneous oral or written communications, proposals, representations and warranties and prevails over any conflicting or additional terms of any quote, order, acknowledgment, or other communication between the parties relating to its subject matter during the term of this Agreement. No modification of this Agreement will be binding, unless in writing and signed by an authorized representative of each party.

For inquiries please contact: Oracle Corporation, 500 Oracle Parkway, Redwood Shores, California 94065, USA.

## ZLIB License

ZLIB license

zlib.h -- interface of the 'zlib' general purpose compression library  
version 1.2.7, May 2nd, 2012

Copyright (C) 1995-2012 Jean-loup Gailly and Mark Adler

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software

- in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
  3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly  
jloup@gzip.org

Mark Adler  
madler@alumni.caltech.edu

### D3.js license (New BSD License)

D3.js license (New BSD License)

Copyright (c) 2012, Michael Bostock  
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- \* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

- \* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

- \* The name Michael Bostock may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL MICHAEL BOSTOCK BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

### Lesser GNU Public License (LGPL)

GNU LESSER GENERAL PUBLIC LICENSE  
Version 3, 29 June 2007

Copyright (C) 2007 Free Software Foundation, Inc. <<http://fsf.org/>>  
Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.



This version of the GNU Lesser General Public License incorporates the terms and conditions of version 3 of the GNU General Public License, supplemented by the additional permissions listed below.

## 0. Additional Definitions.

As used herein, "this License" refers to version 3 of the GNU Lesser General Public License, and the "GNU GPL" refers to version 3 of the GNU General Public License.

"The Library" refers to a covered work governed by this License, other than an Application or a Combined Work as defined below.

An "Application" is any work that makes use of an interface provided by the Library, but which is not otherwise based on the Library. Defining a subclass of a class defined by the Library is deemed a mode of using an interface provided by the Library.

A "Combined Work" is a work produced by combining or linking an Application with the Library. The particular version of the Library with which the Combined Work was made is also called the "Linked Version".

The "Minimal Corresponding Source" for a Combined Work means the Corresponding Source for the Combined Work, excluding any source code for portions of the Combined Work that, considered in isolation, are based on the Application, and not on the Linked Version.

The "Corresponding Application Code" for a Combined Work means the object code and/or source code for the Application, including any data and utility programs needed for reproducing the Combined Work from the Application, but excluding the System Libraries of the Combined Work.

## 1. Exception to Section 3 of the GNU GPL.

You may convey a covered work under sections 3 and 4 of this License without being bound by section 3 of the GNU GPL.

## 2. Conveying Modified Versions.

If you modify a copy of the Library, and, in your modifications, a facility refers to a function or data to be supplied by an Application that uses the facility (other than as an argument passed when the facility is invoked), then you may convey a copy of the modified version:

- a) under this License, provided that you make a good faith effort to ensure that, in the event an Application does not supply the function or data, the facility still operates, and performs

whatever part of its purpose remains meaningful, or

b) under the GNU GPL, with none of the additional permissions of this License applicable to that copy.

### 3. Object Code Incorporating Material from Library Header Files.

The object code form of an Application may incorporate material from a header file that is part of the Library. You may convey such object code under terms of your choice, provided that, if the incorporated material is not limited to numerical parameters, data structure layouts and accessors, or small macros, inline functions and templates (ten or fewer lines in length), you do both of the following:

a) Give prominent notice with each copy of the object code that the Library is used in it and that the Library and its use are covered by this License.

b) Accompany the object code with a copy of the GNU GPL and this license document.

### 4. Combined Works.

You may convey a Combined Work under terms of your choice that, taken together, effectively do not restrict modification of the portions of the Library contained in the Combined Work and reverse engineering for debugging such modifications, if you also do each of the following:

a) Give prominent notice with each copy of the Combined Work that the Library is used in it and that the Library and its use are covered by this License.

b) Accompany the Combined Work with a copy of the GNU GPL and this license document.

c) For a Combined Work that displays copyright notices during execution, include the copyright notice for the Library among these notices, as well as a reference directing the user to the copies of the GNU GPL and this license document.

d) Do one of the following:

0) Convey the Minimal Corresponding Source under the terms of this License, and the Corresponding Application Code in a form suitable for, and under terms that permit, the user to recombine or relink the Application with a modified version of the Linked Version to produce a modified Combined Work, in the manner specified by section 6 of the GNU GPL for conveying Corresponding Source.

1) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (a) uses at run time a copy of the Library already present on the user's computer system, and (b) will operate properly with a modified version of the Library that is interface-compatible with the Linked Version.

e) Provide Installation Information, but only if you would otherwise be required to provide such information under section 6 of the GNU GPL, and only to the extent that such information is necessary to install and execute a modified version of the Combined Work produced by recombining or relinking the Application with a modified version of the Linked Version. (If you use option 4d0, the Installation Information must accompany the Minimal Corresponding Source and Corresponding Application Code. If you use option 4d1, you must provide the Installation Information in the manner specified by section 6 of the GNU GPL for conveying Corresponding Source.)

## 5. Combined Libraries.

You may place library facilities that are a work based on the Library side by side in a single library together with other library facilities that are not Applications and are not covered by this License, and convey such a combined library under terms of your choice, if you do both of the following:

a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities, conveyed under the terms of this License.

b) Give prominent notice with the combined library that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

## 6. Revised Versions of the GNU Lesser General Public License.

The Free Software Foundation may publish revised and/or new versions of the GNU Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library as you received it specifies that a certain numbered version of the GNU Lesser General Public License "or any later version" applies to it, you have the option of following the terms and conditions either of that published version or of any later version published by the Free Software Foundation. If the Library as you received it does not specify a version number of the GNU Lesser General Public License, you may choose any version of the GNU Lesser General Public License ever published by the Free Software Foundation.

If the Library as you received it specifies that a proxy can decide whether future versions of the GNU Lesser General Public License shall

apply, that proxy's public statement of acceptance of any version is permanent authorization for you to choose that version for the Library.

## Lesser GNU Public License (LGPL) v2.1

Lesser GNU Public License (LGPL) v2.1

GNU LESSER GENERAL PUBLIC LICENSE  
Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc.  
51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA  
Everyone is permitted to copy and distribute verbatim copies  
of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts  
as the successor of the GNU Library Public License, version 2, hence  
the version number 2.1.]

### Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages--typically libraries--of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling

it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating

system.

Although the Lesser General Public License is Less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

GNU LESSER GENERAL PUBLIC LICENSE  
TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) The modified work must itself be a software library.
- b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
- c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
- d) If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)



Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)

b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.

c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.

d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.

e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception,

the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.

b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or

otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

## NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

## END OF TERMS AND CONDITIONS

## How to Apply These Terms to Your New Libraries

If you develop a new library, and you want it to be of the greatest possible use to the public, we recommend making it free software that everyone can redistribute and change. You can do so by permitting redistribution under these terms (or, alternatively, under the terms of the ordinary General Public License).

To apply these terms, attach the following notices to the library. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the library's name and a brief idea of what it does.>
Copyright (C) <year> <name of author>
```

```
This library is free software; you can redistribute it and/or
modify it under the terms of the GNU Lesser General Public
License as published by the Free Software Foundation; either
version 2.1 of the License, or (at your option) any later version.
```

```
This library is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU
Lesser General Public License for more details.
```

```
You should have received a copy of the GNU Lesser General Public
```

```
License along with this library; if not, write to the Free Software
Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA
02110-1301 USA
```

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the library, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the
library `Frob' (a library for tweaking knobs) written by James Random
Hacker.
```

```
<signature of Ty Coon>, 1 April 1990
Ty Coon, President of Vice
```

That's all there is to it!

### Boost Software License - Version 1.0 - August 17th, 2003

```
Boost Software License - Version 1.0 - August 17th, 2003
```

Permission is hereby granted, free of charge, to any person or organization obtaining a copy of the software and accompanying documentation covered by this license (the "Software") to use, reproduce, display, distribute, execute, and transmit the Software, and to prepare derivative works of the Software, and to permit third-parties to whom the Software is furnished to do so, all subject to the following:

The copyright notices in the Software and this entire statement, including the above license grant, this restriction and the following disclaimer, must be included in all copies of the Software, in whole or in part, and all derivative works of the Software, unless such copies or derivative works are solely in the form of machine-executable object code generated by a source language processor.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR ANYONE DISTRIBUTING THE SOFTWARE BE LIABLE FOR ANY DAMAGES OR OTHER LIABILITY, WHETHER IN CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

### GNU GENERAL PUBLIC LICENSE Version 3, 29 June 2007

```
GNU GENERAL PUBLIC LICENSE
Version 3, 29 June 2007
```

```
Copyright (C) 2007 Free Software Foundation, Inc. <http://fsf.org/>
Everyone is permitted to copy and distribute verbatim copies
```

of this license document, but changing it is not allowed.

### Preamble

The GNU General Public License is a free, copyleft license for software and other kinds of works.

The licenses for most software and other practical works are designed to take away your freedom to share and change the works. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change all versions of a program--to make sure it remains free software for all its users. We, the Free Software Foundation, use the GNU General Public License for most of our software; it applies also to any other work released this way by its authors. You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for them if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs, and that you know you can do these things.

To protect your rights, we need to prevent others from denying you these rights or asking you to surrender the rights. Therefore, you have certain responsibilities if you distribute copies of the software, or if you modify it: responsibilities to respect the freedom of others.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must pass on to the recipients the same freedoms that you received. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

Developers that use the GNU GPL protect your rights with two steps: (1) assert copyright on the software, and (2) offer you this License giving you legal permission to copy, distribute and/or modify it.

For the developers' and authors' protection, the GPL clearly explains that there is no warranty for this free software. For both users' and authors' sake, the GPL requires that modified versions be marked as changed, so that their problems will not be attributed erroneously to authors of previous versions.

Some devices are designed to deny users access to install or run modified versions of the software inside them, although the manufacturer can do so. This is fundamentally incompatible with the aim of protecting users' freedom to change the software. The systematic pattern of such abuse occurs in the area of products for individuals to use, which is precisely where it is most unacceptable. Therefore, we have designed this version of the GPL to prohibit the practice for those products. If such problems arise substantially in other domains, we stand ready to extend this provision to those domains in future versions of the GPL, as needed to protect the freedom of users.

Finally, every program is threatened constantly by software patents. States should not allow patents to restrict development and use of software on general-purpose computers, but in those that do, we wish to avoid the special danger that patents applied to a free program could make it effectively proprietary. To prevent this, the GPL assures that patents cannot be used to render the program non-free.

The precise terms and conditions for copying, distribution and modification follow.

## TERMS AND CONDITIONS

### 0. Definitions.

"This License" refers to version 3 of the GNU General Public License.

"Copyright" also means copyright-like laws that apply to other kinds of works, such as semiconductor masks.

"The Program" refers to any copyrightable work licensed under this License. Each licensee is addressed as "you". "Licensees" and "recipients" may be individuals or organizations.

To "modify" a work means to copy from or adapt all or part of the work in a fashion requiring copyright permission, other than the making of an exact copy. The resulting work is called a "modified version" of the earlier work or a work "based on" the earlier work.

A "covered work" means either the unmodified Program or a work based on the Program.

To "propagate" a work means to do anything with it that, without permission, would make you directly or secondarily liable for infringement under applicable copyright law, except executing it on a computer or modifying a private copy. Propagation includes copying, distribution (with or without modification), making available to the public, and in some countries other activities as well.

To "convey" a work means any kind of propagation that enables other parties to make or receive copies. Mere interaction with a user through a computer network, with no transfer of a copy, is not conveying.

An interactive user interface displays "Appropriate Legal Notices" to the extent that it includes a convenient and prominently visible feature that (1) displays an appropriate copyright notice, and (2) tells the user that there is no warranty for the work (except to the extent that warranties are provided), that licensees may convey the work under this License, and how to view a copy of this License. If the interface presents a list of user commands or options, such as a menu, a prominent item in the list meets this criterion.

## 1. Source Code.

The "source code" for a work means the preferred form of the work for making modifications to it. "Object code" means any non-source form of a work.

A "Standard Interface" means an interface that either is an official standard defined by a recognized standards body, or, in the case of interfaces specified for a particular programming language, one that is widely used among developers working in that language.

The "System Libraries" of an executable work include anything, other than the work as a whole, that (a) is included in the normal form of packaging a Major Component, but which is not part of that Major Component, and (b) serves only to enable use of the work with that Major Component, or to implement a Standard Interface for which an implementation is available to the public in source code form. A "Major Component", in this context, means a major essential component (kernel, window system, and so on) of the specific operating system (if any) on which the executable work runs, or a compiler used to produce the work, or an object code interpreter used to run it.

The "Corresponding Source" for a work in object code form means all the source code needed to generate, install, and (for an executable work) run the object code and to modify the work, including scripts to control those activities. However, it does not include the work's System Libraries, or general-purpose tools or generally available free programs which are used unmodified in performing those activities but which are not part of the work. For example, Corresponding Source includes interface definition files associated with source files for the work, and the source code for shared libraries and dynamically linked subprograms that the work is specifically designed to require, such as by intimate data communication or control flow between those subprograms and other parts of the work.

The Corresponding Source need not include anything that users can regenerate automatically from other parts of the Corresponding Source.

The Corresponding Source for a work in source code form is that same work.

## 2. Basic Permissions.

All rights granted under this License are granted for the term of copyright on the Program, and are irrevocable provided the stated conditions are met. This License explicitly affirms your unlimited permission to run the unmodified Program. The output from running a covered work is covered by this License only if the output, given its content, constitutes a covered work. This License acknowledges your rights of fair use or other equivalent, as provided by copyright law.

You may make, run and propagate covered works that you do not convey, without conditions so long as your license otherwise remains in force. You may convey covered works to others for the sole purpose



of having them make modifications exclusively for you, or provide you with facilities for running those works, provided that you comply with the terms of this License in conveying all material for which you do not control copyright. Those thus making or running the covered works for you must do so exclusively on your behalf, under your direction and control, on terms that prohibit them from making any copies of your copyrighted material outside their relationship with you.

Conveying under any other circumstances is permitted solely under the conditions stated below. Sublicensing is not allowed; section 10 makes it unnecessary.

### 3. Protecting Users' Legal Rights From Anti-Circumvention Law.

No covered work shall be deemed part of an effective technological measure under any applicable law fulfilling obligations under article 11 of the WIPO copyright treaty adopted on 20 December 1996, or similar laws prohibiting or restricting circumvention of such measures.

When you convey a covered work, you waive any legal power to forbid circumvention of technological measures to the extent such circumvention is effected by exercising rights under this License with respect to the covered work, and you disclaim any intention to limit operation or modification of the work as a means of enforcing, against the work's users, your or third parties' legal rights to forbid circumvention of technological measures.

### 4. Conveying Verbatim Copies.

You may convey verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice; keep intact all notices stating that this License and any non-permissive terms added in accord with section 7 apply to the code; keep intact all notices of the absence of any warranty; and give all recipients a copy of this License along with the Program.

You may charge any price or no price for each copy that you convey, and you may offer support or warranty protection for a fee.

### 5. Conveying Modified Source Versions.

You may convey a work based on the Program, or the modifications to produce it from the Program, in the form of source code under the terms of section 4, provided that you also meet all of these conditions:

- a) The work must carry prominent notices stating that you modified it, and giving a relevant date.
- b) The work must carry prominent notices stating that it is released under this License and any conditions added under section 7. This requirement modifies the requirement in section 4 to

"keep intact all notices".

c) You must license the entire work, as a whole, under this License to anyone who comes into possession of a copy. This License will therefore apply, along with any applicable section 7 additional terms, to the whole of the work, and all its parts, regardless of how they are packaged. This License gives no permission to license the work in any other way, but it does not invalidate such permission if you have separately received it.

d) If the work has interactive user interfaces, each must display Appropriate Legal Notices; however, if the Program has interactive interfaces that do not display Appropriate Legal Notices, your work need not make them do so.

A compilation of a covered work with other separate and independent works, which are not by their nature extensions of the covered work, and which are not combined with it such as to form a larger program, in or on a volume of a storage or distribution medium, is called an "aggregate" if the compilation and its resulting copyright are not used to limit the access or legal rights of the compilation's users beyond what the individual works permit. Inclusion of a covered work in an aggregate does not cause this License to apply to the other parts of the aggregate.

## 6. Conveying Non-Source Forms.

You may convey a covered work in object code form under the terms of sections 4 and 5, provided that you also convey the machine-readable Corresponding Source under the terms of this License, in one of these ways:

a) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by the Corresponding Source fixed on a durable physical medium customarily used for software interchange.

b) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by a written offer, valid for at least three years and valid for as long as you offer spare parts or customer support for that product model, to give anyone who possesses the object code either (1) a copy of the Corresponding Source for all the software in the product that is covered by this License, on a durable physical medium customarily used for software interchange, for a price no more than your reasonable cost of physically performing this conveying of source, or (2) access to copy the Corresponding Source from a network server at no charge.

c) Convey individual copies of the object code with a copy of the written offer to provide the Corresponding Source. This alternative is allowed only occasionally and noncommercially, and only if you received the object code with such an offer, in accord with subsection 6b.

d) Convey the object code by offering access from a designated place (gratis or for a charge), and offer equivalent access to the Corresponding Source in the same way through the same place at no further charge. You need not require recipients to copy the Corresponding Source along with the object code. If the place to copy the object code is a network server, the Corresponding Source may be on a different server (operated by you or a third party) that supports equivalent copying facilities, provided you maintain clear directions next to the object code saying where to find the Corresponding Source. Regardless of what server hosts the Corresponding Source, you remain obligated to ensure that it is available for as long as needed to satisfy these requirements.

e) Convey the object code using peer-to-peer transmission, provided you inform other peers where the object code and Corresponding Source of the work are being offered to the general public at no charge under subsection 6d.

A separable portion of the object code, whose source code is excluded from the Corresponding Source as a System Library, need not be included in conveying the object code work.

A "User Product" is either (1) a "consumer product", which means any tangible personal property which is normally used for personal, family, or household purposes, or (2) anything designed or sold for incorporation into a dwelling. In determining whether a product is a consumer product, doubtful cases shall be resolved in favor of coverage. For a particular product received by a particular user, "normally used" refers to a typical or common use of that class of product, regardless of the status of the particular user or of the way in which the particular user actually uses, or expects or is expected to use, the product. A product is a consumer product regardless of whether the product has substantial commercial, industrial or non-consumer uses, unless such uses represent the only significant mode of use of the product.

"Installation Information" for a User Product means any methods, procedures, authorization keys, or other information required to install and execute modified versions of a covered work in that User Product from a modified version of its Corresponding Source. The information must suffice to ensure that the continued functioning of the modified object code is in no case prevented or interfered with solely because modification has been made.

If you convey an object code work under this section in, or with, or specifically for use in, a User Product, and the conveying occurs as part of a transaction in which the right of possession and use of the User Product is transferred to the recipient in perpetuity or for a fixed term (regardless of how the transaction is characterized), the Corresponding Source conveyed under this section must be accompanied by the Installation Information. But this requirement does not apply if neither you nor any third party retains the ability to install modified object code on the User Product (for example, the work has been installed in ROM).

The requirement to provide Installation Information does not include a requirement to continue to provide support service, warranty, or updates for a work that has been modified or installed by the recipient, or for the User Product in which it has been modified or installed. Access to a

network may be denied when the modification itself materially and adversely affects the operation of the network or violates the rules and protocols for communication across the network.

Corresponding Source conveyed, and Installation Information provided, in accord with this section must be in a format that is publicly documented (and with an implementation available to the public in source code form), and must require no special password or key for unpacking, reading or copying.

## 7. Additional Terms.

"Additional permissions" are terms that supplement the terms of this License by making exceptions from one or more of its conditions. Additional permissions that are applicable to the entire Program shall be treated as though they were included in this License, to the extent that they are valid under applicable law. If additional permissions apply only to part of the Program, that part may be used separately under those permissions, but the entire Program remains governed by this License without regard to the additional permissions.

When you convey a copy of a covered work, you may at your option remove any additional permissions from that copy, or from any part of it. (Additional permissions may be written to require their own removal in certain cases when you modify the work.) You may place additional permissions on material, added by you to a covered work, for which you have or can give appropriate copyright permission.

Notwithstanding any other provision of this License, for material you add to a covered work, you may (if authorized by the copyright holders of that material) supplement the terms of this License with terms:

- a) Disclaiming warranty or limiting liability differently from the terms of sections 15 and 16 of this License; or
- b) Requiring preservation of specified reasonable legal notices or author attributions in that material or in the Appropriate Legal Notices displayed by works containing it; or
- c) Prohibiting misrepresentation of the origin of that material, or requiring that modified versions of such material be marked in reasonable ways as different from the original version; or
- d) Limiting the use for publicity purposes of names of licensors or authors of the material; or
- e) Declining to grant rights under trademark law for use of some trade names, trademarks, or service marks; or
- f) Requiring indemnification of licensors and authors of that material by anyone who conveys the material (or modified versions of it) with contractual assumptions of liability to the recipient, for any liability that these contractual assumptions directly impose on

those licensors and authors.

All other non-permissive additional terms are considered "further restrictions" within the meaning of section 10. If the Program as you received it, or any part of it, contains a notice stating that it is governed by this License along with a term that is a further restriction, you may remove that term. If a license document contains a further restriction but permits relicensing or conveying under this License, you may add to a covered work material governed by the terms of that license document, provided that the further restriction does not survive such relicensing or conveying.

If you add terms to a covered work in accord with this section, you must place, in the relevant source files, a statement of the additional terms that apply to those files, or a notice indicating where to find the applicable terms.

Additional terms, permissive or non-permissive, may be stated in the form of a separately written license, or stated as exceptions; the above requirements apply either way.

## 8. Termination.

You may not propagate or modify a covered work except as expressly provided under this License. Any attempt otherwise to propagate or modify it is void, and will automatically terminate your rights under this License (including any patent licenses granted under the third paragraph of section 11).

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, you do not qualify to receive new licenses for the same material under section 10.

## 9. Acceptance Not Required for Having Copies.

You are not required to accept this License in order to receive or run a copy of the Program. Ancillary propagation of a covered work occurring solely as a consequence of using peer-to-peer transmission

to receive a copy likewise does not require acceptance. However, nothing other than this License grants you permission to propagate or modify any covered work. These actions infringe copyright if you do not accept this License. Therefore, by modifying or propagating a covered work, you indicate your acceptance of this License to do so.

#### 10. Automatic Licensing of Downstream Recipients.

Each time you convey a covered work, the recipient automatically receives a license from the original licensors, to run, modify and propagate that work, subject to this License. You are not responsible for enforcing compliance by third parties with this License.

An "entity transaction" is a transaction transferring control of an organization, or substantially all assets of one, or subdividing an organization, or merging organizations. If propagation of a covered work results from an entity transaction, each party to that transaction who receives a copy of the work also receives whatever licenses to the work the party's predecessor in interest had or could give under the previous paragraph, plus a right to possession of the Corresponding Source of the work from the predecessor in interest, if the predecessor has it or can get it with reasonable efforts.

You may not impose any further restrictions on the exercise of the rights granted or affirmed under this License. For example, you may not impose a license fee, royalty, or other charge for exercise of rights granted under this License, and you may not initiate litigation (including a cross-claim or counterclaim in a lawsuit) alleging that any patent claim is infringed by making, using, selling, offering for sale, or importing the Program or any portion of it.

#### 11. Patents.

A "contributor" is a copyright holder who authorizes use under this License of the Program or a work on which the Program is based. The work thus licensed is called the contributor's "contributor version".

A contributor's "essential patent claims" are all patent claims owned or controlled by the contributor, whether already acquired or hereafter acquired, that would be infringed by some manner, permitted by this License, of making, using, or selling its contributor version, but do not include claims that would be infringed only as a consequence of further modification of the contributor version. For purposes of this definition, "control" includes the right to grant patent sublicenses in a manner consistent with the requirements of this License.

Each contributor grants you a non-exclusive, worldwide, royalty-free patent license under the contributor's essential patent claims, to make, use, sell, offer for sale, import and otherwise run, modify and propagate the contents of its contributor version.

In the following three paragraphs, a "patent license" is any express agreement or commitment, however denominated, not to enforce a patent (such as an express permission to practice a patent or covenant not to

sue for patent infringement). To "grant" such a patent license to a party means to make such an agreement or commitment not to enforce a patent against the party.

If you convey a covered work, knowingly relying on a patent license, and the Corresponding Source of the work is not available for anyone to copy, free of charge and under the terms of this License, through a publicly available network server or other readily accessible means, then you must either (1) cause the Corresponding Source to be so available, or (2) arrange to deprive yourself of the benefit of the patent license for this particular work, or (3) arrange, in a manner consistent with the requirements of this License, to extend the patent license to downstream recipients. "Knowingly relying" means you have actual knowledge that, but for the patent license, your conveying the covered work in a country, or your recipient's use of the covered work in a country, would infringe one or more identifiable patents in that country that you have reason to believe are valid.

If, pursuant to or in connection with a single transaction or arrangement, you convey, or propagate by procuring conveyance of, a covered work, and grant a patent license to some of the parties receiving the covered work authorizing them to use, propagate, modify or convey a specific copy of the covered work, then the patent license you grant is automatically extended to all recipients of the covered work and works based on it.

A patent license is "discriminatory" if it does not include within the scope of its coverage, prohibits the exercise of, or is conditioned on the non-exercise of one or more of the rights that are specifically granted under this License. You may not convey a covered work if you are a party to an arrangement with a third party that is in the business of distributing software, under which you make payment to the third party based on the extent of your activity of conveying the work, and under which the third party grants, to any of the parties who would receive the covered work from you, a discriminatory patent license (a) in connection with copies of the covered work conveyed by you (or copies made from those copies), or (b) primarily for and in connection with specific products or compilations that contain the covered work, unless you entered into that arrangement, or that patent license was granted, prior to 28 March 2007.

Nothing in this License shall be construed as excluding or limiting any implied license or other defenses to infringement that may otherwise be available to you under applicable patent law.

## 12. No Surrender of Others' Freedom.

If conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot convey a covered work so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not convey it at all. For example, if you agree to terms that obligate you to collect a royalty for further conveying from those to whom you convey the Program, the only way you could satisfy both those terms and this License would be to refrain entirely from conveying the Program.

### 13. Use with the GNU Affero General Public License.

Notwithstanding any other provision of this License, you have permission to link or combine any covered work with a work licensed under version 3 of the GNU Affero General Public License into a single combined work, and to convey the resulting work. The terms of this License will continue to apply to the part which is the covered work, but the special requirements of the GNU Affero General Public License, section 13, concerning interaction through a network will apply to the combination as such.

### 14. Revised Versions of this License.

The Free Software Foundation may publish revised and/or new versions of the GNU General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies that a certain numbered version of the GNU General Public License "or any later version" applies to it, you have the option of following the terms and conditions either of that numbered version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of the GNU General Public License, you may choose any version ever published by the Free Software Foundation.

If the Program specifies that a proxy can decide which future versions of the GNU General Public License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Program.

Later license versions may give you additional or different permissions. However, no additional obligations are imposed on any author or copyright holder as a result of your choosing to follow a later version.

### 15. Disclaimer of Warranty.

THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

### 16. Limitation of Liability.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE



USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

#### 17. Interpretation of Sections 15 and 16.

If the disclaimer of warranty and limitation of liability provided above cannot be given local legal effect according to their terms, reviewing courts shall apply local law that most closely approximates an absolute waiver of all civil liability in connection with the Program, unless a warranty or assumption of liability accompanies a copy of the Program in return for a fee.

#### END OF TERMS AND CONDITIONS

#### How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively state the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the program's name and a brief idea of what it does.>
Copyright (C) <year> <name of author>
```

```
This program is free software: you can redistribute it and/or modify
it under the terms of the GNU General Public License as published by
the Free Software Foundation, either version 3 of the License, or
(at your option) any later version.
```

```
This program is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
GNU General Public License for more details.
```

```
You should have received a copy of the GNU General Public License
along with this program. If not, see <http://www.gnu.org/licenses/>.
```

Also add information on how to contact you by electronic and paper mail.

If the program does terminal interaction, make it output a short notice like this when it starts in an interactive mode:

```
<program> Copyright (C) <year> <name of author>
This program comes with ABSOLUTELY NO WARRANTY; for details type `show
```

w' .

This is free software, and you are welcome to redistribute it under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, your program's commands might be different; for a GUI interface, you would use an "about box".

You should also get your employer (if you work as a programmer) or school, if any, to sign a "copyright disclaimer" for the program, if necessary. For more information on this, and how to apply and follow the GNU GPL, see <http://www.gnu.org/licenses/>.

The GNU General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License. But first, please read <http://www.gnu.org/philosophy/why-not-lgpl.html>.

## Other Resources

---

Provides links to additional resources such as on-demand training, videos, blogs, and the HPE Ezmeral Data Fabric community.

In addition to the product documentation, you may be interested in the following resources:

<b>Training</b>	<a href="https://learn.software.hpe.com/">https://learn.software.hpe.com/</a>
<b>Blogs and Videos</b>	<a href="https://community.hpe.com/t5/hpe-ezmeral-uncut/bg-p/software#.XzXV2-hKg2w">https://community.hpe.com/t5/hpe-ezmeral-uncut/bg-p/software#.XzXV2-hKg2w</a>
<b>HPE Ezmeral Software Community</b>	<a href="#">HPE Ezmeral Software Community</a>
<b>HPE Developer Community</b>	<a href="https://developer.hpe.com/">https://developer.hpe.com/</a>
<b>Slack Community for Developers</b>	<a href="https://slack.hpedev.io/">https://slack.hpedev.io/</a>
<b>HPE Support Center</b>	<a href="https://support.hpe.com/">https://support.hpe.com/</a>
<b>Contact HPE</b>	<a href="https://www.hpe.com/us/en/contact-hpe.html">https://www.hpe.com/us/en/contact-hpe.html</a>
<b>Videos, Reports, and Case Studies</b>	<a href="https://www.hpe.com/us/en/resource-library.html">https://www.hpe.com/us/en/resource-library.html</a>
<b>HPE GreenLake Marketplace</b>	<a href="https://www.hpe.com/us/en/software/marketplace.html/platform/ezmeraldata">https://www.hpe.com/us/en/software/marketplace.html/platform/ezmeraldata</a>

## Contact HPE

---

Provides a link to contact HPE Sales or Support.

[Contact HPE](#)

## Glossary

---

List of terms (with description) used in HPE Ezmeral Data Fabric documentation.

## **.snapshot**

---

A special directory in the top level of each volume that contains all the snapshots created or preserved for the volume.

## **access control expression (ACE)**

---

A Boolean expression that defines a combination of users, groups, or roles that have access to an object stored natively such as a directory, file, or HPE Ezmeral Data Fabric Database table.

Access Control Expression (ACE)



**NOTE:** An ACE (up to 64KB in length) is a combination of users, groups, and/or roles for whom access (to volume data) is defined using boolean expressions and sub expressions within single quotes. When you pass in an access type that has already been set, the new value replaces the existing value for that access type. There is no change to access types that are not passed in with the command, whether or not they were set.

## **access control list (ACL)**

---

A list of permissions attached to an object. An ACL specifies users or system processes that can perform specific actions on an object.

Access Control List (ACL)



**NOTE:** An Access Control List (ACL) is a list of users or groups. Each user or group in the list is paired with a defined set of permissions that limit the actions that the user or group can perform on the object secured by the ACL. In the HPE Ezmeral Data Fabric, the objects secured by ACLs are the job queue, volumes, and the cluster itself.

## **access policy**

---

An ACL or policy in JSON format that describes user access. Grants accounts and IAM users permissions to perform resource operations, such as putting objects in a bucket. You associate access policies with accounts, users, buckets, and objects.

## **administrator**

---

A user or users with special privileges to administer the cluster or cluster resources. Administrative functions can include managing hardware resources, users, data, services, security, and availability.

## **advisory quota**

---

An advisory disk capacity limit that can be set for a volume, user, or group. When disk usage exceeds the advisory quota, an alert is sent.

## air gap

---

Physical isolation between a computer system and unsecured networks. To enhance security, air-gapped computer systems are disconnected from other systems and networks.

## chunk

---

Files in the file system are split into chunks (similar to Hadoop blocks) that are normally 256 MB by default. Any multiple of 65,536 bytes is a valid chunk size, but tuning the size correctly is important. Files inherit the chunk size settings of the directory that contains them, as do subdirectories on which chunk size has not been explicitly set. Any files written by a Hadoop application, whether via the file APIs or over NFS, use chunk size specified by the settings for the directory where the file is written.

## client node

---

A node that runs the `mapr-client` that can access every cluster node and is used to access the cluster. Also referred to as an "edge node." Client nodes and edge nodes are NOT part of a Data Fabric cluster.

See also [node](#) and [edge node](#).

## cluster admin

---

The Data Fabric user.

For more information, see [Data Fabric user](#) on page 342.

## compute node

---

A compute node is used to process data using a compute engine (for example, YARN, Hive, Spark, or Drill). A compute node is by definition a Data Fabric cluster node.

See also [node](#).

Compare with [data node](#).

## container

---

The unit of shared storage in a Data Fabric cluster. Every container is either a name container or a data container.

### More information

[Docker containers](#) on page 342

[YARN resource containers](#) on page 348

## container location database (CLDB)

---

A service, running on one or more Data Fabric nodes, that maintains the locations of services, containers, and other cluster information.

**NOTE:**

The Container Location Database (CLDB) service tracks the following information about every container in the file system:

- The node where the container is located
- Size of the container
- The volume to which the container belongs
- The policies, quotas, and usage for that volume

For more information about the CLDB, see the following HPE Ezmeral Data Fabric – Customer Managed documentation topic: [CLDB](#).

## core

---

The minimum complement of software packages required to construct Data Fabric cluster. These packages include `mapr-core`, `mapr-core-internal`, `mapr-cldb`, `mapr-apiserver`, `mapr-fileserver`, `mapr-zookeeper`, and others. Note that ecosystem components are not part of core.

## data-access gateway

---

A service that acts as a proxy and gateway for translating requests between lightweight client applications and the Data Fabric cluster.

## data compaction

---

A process that enables users to remove empty or deleted space in the database and to compact the database to occupy contiguous space.

**More information**

[log compaction](#) on page 344

## data container

---

One of the two types of containers in a Data Fabric cluster. Data containers typically have a cascaded configuration (master replicates to replica1, replica1 replicates to replica2, and so on). Every data container is either a master container, an intermediate container, or a tail container depending on its replication role.

## Data Fabric

---

A collection of nodes that work together under a unified architecture, along with the services or technologies running on that architecture. A fabric is similar to a Linux cluster. Fabrics help you manage your data, making it possible to access, integrate, model, analyze, and provision your data seamlessly.

## Data Fabric administrator

---

The "Data Fabric user." The user that cluster services run as (typically named `mapr` or `hadoop`) on each node.

See [Data Fabric user](#) on page 342.

## Data Fabric gateway

---

A gateway that supports table and stream replication. The Data Fabric gateway mediates one-way communication between a source Data Fabric cluster and a destination cluster. The Data Fabric gateway also applies updates from JSON tables to their secondary indexes and propagates Change Data Capture (CDC) logs.

## Data Fabric user

---

The user that cluster services run as (typically named `mapr` or `hadoop`) on each node. The Data Fabric user, also known as the "Data Fabric admin," has full privileges to administer the cluster. The administrative privilege, with varying levels of control, can be assigned to other users as well.

## data node

---

A data node has the function of storing data and always runs FileServer. A data node is by definition a Data Fabric cluster node.

See also [node](#).

Compare with [compute node](#).

## desired replication factor

---

The number of copies of a volume that should be maintained by the Data Fabric cluster for normal operation.

When the number of copies falls below the desired replication factor, but remains equal to or above the [minimum replication factor](#), re-replication occurs after the timeout specified in the `cldb.fs.mark.rereplicate.sec` parameter.

## developer preview

---

A label for a feature or collection of features that have usage restrictions. Developer previews are not tested for production environments, and should be used with caution.

## Docker containers

---

The application containers used by Docker software. Docker is a leading proponent of OS virtualization using application containers ("containerization").

## Domain

---

Relates to Object Store. A domain is a management entity for accounts and users. The number of users, the amount of disk space, number of buckets in each of the accounts, total number of accounts, and the number of disabled accounts are all tracked within a domain. Currently, Object Store only supports the primary domain; you cannot create additional domains. Administrators can create multiple accounts in the primary domain.

## domain user

---

Relates to Object Store. A domain user is a cluster security principal authenticated through AD/LDAP. Domain users only exist in the default account. Domain users can log in to the Object Store UI with their domain username and password.

## Ecosystem Pack (EEP)

---

A selected set of stable, interoperable, and widely used components from the Hadoop ecosystem that are fully supported on the Data Fabric platform.

## edge cluster

---

A small-footprint edition of the HPE Ezmeral Data Fabric designed to capture, process, and analyze IoT data close to the source of the data.

## edge node

---

A node that runs the `mapr-client` that can access every cluster node and is used to access the cluster. Also referred to as a "client node." Client nodes and edge nodes are NOT part of a Data Fabric cluster.

See also [node](#) and [client node](#).

## fabric

---

A collection of nodes that work together under a unified architecture, along with the services or technologies running on that architecture. A fabric is similar to a Linux cluster. Fabrics help you manage your data, making it possible to access, integrate, model, analyze, and provision your data seamlessly.

## filelet

---

A filelet, also called an fid, is a 256MB shard of a file. A 1 GB file for instance is comprised of the following filelets: 64K (primary fid)+(256MB-64KB)+256MB+256MB+256MB.

## file system

---

The NFS-mountable, distributed, high-performance HPE Ezmeral Data Fabric data-storage system.

## gateway node

---

A node on which a `mapr-gateway` is installed. A gateway node is by definition a Data Fabric cluster node.

See also [node](#)

## global namespace (GNS)

---

The data plane that connects HPE Ezmeral Data Fabric deployments. The global namespace is a mechanism that aggregates disparate and remote data sources and provides a namespace that

encompasses all of your infrastructure and deployments. Global namespace technology lets you manage globally deployed data as a single resource. Because of the global namespace, you can view and run multiple fabrics as a single, logical, and local fabric. The global namespace is designed to span multiple edge nodes, on-prem data centers, and clouds. See [Global Namespace \(GNS\)](#) on page 88.

## heartbeat

---

A signal sent by each FileServer and NFS node every second to provide information to the CLDB about the node's health and resource usage.

## IAM users

---

Relates to Object Store. An IAM (Identity and Access Management) user represents an actual user or an application. An administrator creates IAM users in an Object Store account and assigns access policies to them to control user and application access to resources in the account.

## Installer

---

A program that simplifies installation of the HPE Ezmeral Data Fabric. The Installer guides you through the process of installing a cluster with Data Fabric services and ecosystem components. You can also use the Installer to update a previously installed cluster with additional nodes, services, and ecosystem components. And you can use the Installer to upgrade a cluster to a newer core version if the cluster was installed using the Installer or an Installer Stanza.

## log compaction

---

A process that purges messages previously published to a topic partition, retaining the latest version.

### More information

[data compaction](#) on page 341

## MAST Gateway

---

A gateway that serves as a centralized entry point for all the operations that need to be performed on tiered storage.

## minimum replication factor

---

The minimum number of copies of a volume that should be maintained by the Data Fabric cluster for normal operation. When the replication factor falls below this minimum, re-replication occurs as aggressively as possible to restore the replication level. If any containers in the CLDB volume fall below the minimum replication factor, writes are disabled until aggressive re-replication restores the minimum level of replication.

## mirror

---

A replica of a volume.



## MOSS

---

MOSS is the acronym for Multithreaded Object Store Server.

## name container

---

A container in a Data Fabric cluster that holds a volume's namespace information and file chunk locations, and the first 64 KB of each file in the volume.

## Network File System (NFS)

---

A protocol that allows a user on a client computer to access files over a network as though they were stored locally.

## node

---

An individual physical or virtual machine in a fabric.

## NodeManager (NM)

---

A data service that works with the ResourceManager to host the YARN resource containers that run on each data node.

## object

---

File and metadata that describes the file. You upload an object into a bucket. You can then download, open, move, or delete the object.

## Object Store

---

Object and metadata storage solution built into the HPE Ezmeral Data Fabric. Object Store efficiently stores data for fast access and leverages the capabilities of the patented HPE Ezmeral Data Fabric file system for performance, reliability, and scalability.

## policy server

---

The service that manages security policies and composite IDs.

## quota

---

A disk capacity limit that can be set for a volume, user, or group. When disk usage exceeds the quota, no more data can be written.

## replication factor

---

The number of copies of a volume.

## replication role

---

The replication role of a container determines how that container is replicated to other storage pools in the cluster.

A *name container* may have one of two replication roles: master or replica. A *data container* may have one of three replication roles: master, intermediate, or tail.

## replication role balancer

---

The replication role balancer is a tool that switches the replication roles of containers to ensure that every node has an equal share of of master and replica containers (for name containers) and an equal share of master, intermediate, and tail containers (for data containers).

## re-replication

---

Re-replication occurs whenever the number of available replica containers drops below the number prescribed by that volume's replication factor. Re-replication may occur for a variety of reasons including replica container corruption, node unavailability, hard disk failure, or an increase in replication factor.

## ResourceManager (RM)

---

A YARN service that manages cluster resources and schedules applications.

## role

---

The service that the node runs in a cluster. You can use a node for one, or a combination of the following roles: CLDB, JobTracker, WebServer, ResourceManager, Zookeeper, FileServer, TaskTracker, NFS, and HBase.

## secret

---

A Kubernetes object that holds sensitive information, such as passwords, tokens, and keys. Pods that require this sensitive information reference the secret in their pod definition. Secrets are the method Kubernetes uses to move sensitive data into pods.

## secure by default

---

The HPE Ezmeral Data Fabric platform and supported ecosystem components are designed to implement security unless the user takes specific steps to turn off security options.

## schedule

---

A group of rules that specify recurring points in time at which certain actions are determined to occur.

## snapshot

---

A read-only logical image of a volume at a specific point in time.

## storage pool

---

A unit of storage made up of one or more disks. By default, Data Fabric storage pools contain two or three disks. For high-volume reads and writes, you can create larger storage pools when initially formatting storage during cluster creation.



**NOTE:** Storage pool refers to the combined storage capacity that is obtained by combining one or more storage devices. Storage devices can be anything from a very small disk drive to large arrays of disk drives (each containing 20-30 drives).

A storage pool is created to get a very large capacity of GBs/TBs/PBs available, from which users are provided needed amounts of storage

For example, one can combine 10 hard disk drives of 4TB each, totaling to 40TBs. Now, one can either directly use the 40TB as a single device or partition the space out to many smaller storage capacities such as 100GB, 1TB and so on from this 40TB and provide that access to different users.

## stripe width

---

The number of disks in a storage pool. See [storage pool](#).

## super group

---

The group that has administrative access to the Data Fabric cluster.

## super user

---

The user that has administrative access to the Data Fabric cluster.

## tagging

---

Operation of applying a security policy to a resource.

## ticket

---

In the Data Fabric platform, a file that contains keys used to authenticate users and cluster servers. Tickets are created using the `maprlogin` or `configure.sh` utilities and are encrypted to protect their contents.

Different types of tickets are provided for users and services. For example, every user who wants to access a cluster must have a user ticket, and every node in a cluster must have a server ticket.

## volume

---

A tree of files and directories grouped for the purpose of applying a policy or set of policies to all of them at once.

## Warden

---

A Data Fabric process that coordinates the starting and stopping of configured services on a node.

## **YARN resource containers**

---

A unit of memory allocated for use by YARN to process each map or reduce task.

## **ZooKeeper**

---

A coordination service for distributed applications. It provides a shared hierarchical namespace that is organized like a standard file system.