

Elastic Cloud Storage (ECS)

Version 2.2

Administrator's Guide

302-002-523

02

Copyright © 2013-2016 EMC Corporation. All rights reserved. Published in the USA.

Published February, 2016

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>).

EMC Corporation
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.EMC.com

CONTENTS

Figures		7
Tables		9
Part 1	Use the ECS Portal	11
Chapter 1	Introduction	13
	Introduction to the ECS Portal.....	14
	Log in to the ECS Portal.....	14
	Change Password.....	16
	Access to portal areas.....	16
	Ordering and searching tables in the portal.....	19
	About this VDC.....	20
Chapter 2	Use the Getting Started Checklist	23
	Using the Getting Started Checklist.....	24
Chapter 3	Use the Dashboard	27
	Use the Portal Dashboard.....	28
Part 2	Configure and Manage	33
Chapter 4	Configure One or More Sites	35
	Configure storage pools, VDCs, and replication groups.....	36
	Storage pools.....	36
	Create storage pools.....	37
	Virtual data centers (VDCs).....	39
	Create a VDC for a single site.....	40
	Add a VDC to a federation.....	40
	Fail over a site/Delete a VDC.....	41
	Replication groups.....	41
	Create replication groups.....	42
	Configuring ESRS.....	43
Chapter 5	Configure a namespace	45
	Introduction.....	46
	Understanding tenants.....	46
	Understanding namespace settings.....	47
	Working with namespaces at the ECS portal.....	50
	Create and configure a namespace.....	50
Chapter 6	Configure authentication and manage users	53

	Introduction.....	54
	Understanding users and roles in ECS.....	54
	Users in ECS.....	54
	User roles.....	55
	Domain and local users.....	57
	User scope: global or namespace.....	57
	Working with the users at the ECS Portal.....	58
	Add a new object user.....	60
	Add a domain user as an object user.....	61
	Create a local management user or assign a domain user to a management role.....	61
	Assign an Active Directory group name to the system admin or system monitor role.....	62
	Create a namespace administrator.....	63
	Assign an Active Directory group name to the namespace admin role	64
	Working with the authentication providers at the ECS Portal.....	64
	Add an authentication provider.....	65
	Authentication provider settings.....	66
	Considerations when adding authentication providers.....	69
	Understanding the mapping of users into a namespace.....	70
	Map domain users into a namespace.....	71
Chapter 7	Manage tenants	73
	Introduction.....	74
	Quotas.....	74
	Retention periods and policies.....	75
	Lock buckets and users.....	76
	Metering.....	77
	Audit buckets.....	78
Chapter 8	Remove a site	81
	Fail over a site/Delete a VDC.....	82
Chapter 9	Manage licenses	83
	Licensing.....	84
	Obtain the EMC ECS license file.....	84
Chapter 10	Create and manage buckets	85
	Introduction.....	86
	Bucket concepts.....	86
	Bucket access.....	86
	Bucket ownership.....	86
	Access to a bucket during temporary site outage.....	87
	Bucket attributes.....	87
	Default Group.....	89
	Metadata index keys.....	90
	Bucket tagging.....	91
	Bucket ACLs.....	92
	Create a bucket using the ECS Portal.....	93
	Edit a bucket.....	94

	Set the bucket ACL permissions for a user.....	95
	Set the bucket ACL permissions for a pre-defined group.....	96
	Set custom group bucket ACLs.....	97
	Create bucket using the object APIs.....	98
	Create a bucket using the S3 API (with s3curl).....	99
	Bucket and key naming conventions.....	102
	S3 bucket and object naming in ECS.....	102
	OpenStack Swift container and object naming in ECS.....	103
	Atmos bucket and object naming in ECS.....	103
	CAS pool and object naming in ECS.....	103
Chapter 11	Configure NFS file access	105
	Configure file access.....	106
Chapter 12	Set the Base URL	107
	Introduction.....	108
	Bucket addressing.....	108
	DNS Configuration.....	109
	Base URL.....	109
	Add a Base URL.....	110
Part 3	Monitor	113
Chapter 13	Monitoring basics	115
	Using monitoring pages.....	116
Chapter 14	Monitor metering	119
	Monitor metering data.....	120
	Metering data.....	121
Chapter 15	Monitor events	123
	About event monitoring.....	124
	Monitor audit data.....	124
	Monitor alerts.....	125
Chapter 16	Monitor capacity utilization	127
	Monitor capacity.....	128
	Storage capacity data.....	128
Chapter 17	Monitor traffic metrics	131
	Monitor network traffic.....	132
Chapter 18	Monitor hardware health	135
	Monitor hardware.....	136
Chapter 19	Monitor node and process health	137

	Monitor node and process health.....	138
Chapter 20	Monitor chunk summary	141
	Monitor chunks.....	142
Chapter 21	Monitor erasure coding	145
	Monitor erasure coding.....	146
Chapter 22	Monitor recovery status	149
	Monitor recovery status.....	150
Chapter 23	Monitor disk bandwidth	151
	Monitor disk bandwidth.....	152
Chapter 24	Monitor geo-replication	155
	Introduction to Geo-replication monitoring.....	156
	Monitor geo-replication: Rate and Chunks.....	156
	Monitor geo-replication: Recovery Point Objective (RPO).....	157
	Monitor geo-replication: Failover Processing.....	158
	Monitor geo replication: Bootstrap Processing.....	159
Chapter 25	Service logs	161
	Service logs.....	162
	ECS service log locations.....	162

FIGURES

1	Table Column with Sort Control Available.....	20
2	Starred Rows Indicate Nodes with a Different Software Version than the Current Node...	21
3	A Completed Checklist.....	25
4	Storage Pool Management page.....	36
5	VDCManagement page.....	39
6	Manage Replication Groups page.....	41
7	Namespace management page.....	50
8	User mappings for a tenant using AD attributes.....	70
9	Using multiple mapping criteria.....	71
10	Refresh.....	116
11	Open Filter panel with criteria selected.....	116
12	Closed Fiter panel showing summary of applied filter.....	116
13	Navigating with breadcrumbs.....	117
14	History chart with active cursor.....	118
15	Metering page with criteria selected.....	121
16	Network traffic charts for a VDC.....	133
17	Hardware Health.....	136
18	Node & Process Health.....	140
19	Chunk Summary.....	143
20	Disk Bandwidth.....	153
21	Geo replication: Rate and Chunks	157
22	RPO.....	157
23	Failover	159
24	Bootstrap processing.....	160

FIGURES

TABLES

1	Storage pool properties.....	36
2	VDC properties.....	39
3	Replication Group properties.....	41
4	Authentication provider settings.....	66
5	Bucket and namespace metering.....	77
6	Bucket attributes.....	87
7	Bucket ACLs.....	92
8	Bucket and namespace metering.....	121
9	Alert types.....	125
10	Capacity Utilization: Storage Pool.....	128
11	Capacity Utilization: Node.....	129
12	Capacity Utilization: Disk.....	130
13	Network traffic metrics.....	132
14	VDC, node, and process health metrics.....	138
15	Chunk tables.....	142
16	Chunk metrics.....	142
17	Erasur coding metrics.....	146
18	Recovery metrics.....	150
19	Disk bandwidth metrics.....	152
20	Rate and Chunk columns.....	156
21	RPO columns.....	157
22	Failover columns.....	158
23	Bootstrap Processing columns.....	159

PART 1

Use the ECS Portal

[Chapter 1, "Introduction"](#)

[Chapter 2, "Use the Getting Started Checklist"](#)

[Chapter 3, "Use the Dashboard"](#)

Use the ECS Portal

CHAPTER 1

Introduction

- [Introduction to the ECS Portal](#)..... 14
- [Log in to the ECS Portal](#)..... 14
- [Change Password](#)..... 16
- [Access to portal areas](#)..... 16
- [Ordering and searching tables in the portal](#)..... 19
- [About this VDC](#)..... 20

Introduction to the ECS Portal

The ECS Portal enables you to configure, manage, and monitor ECS.

The portal also allows tenants to manage and monitor their namespace and to create and configure buckets within their namespace.

The portal provides access for ECS management users: System Administrators, System Monitors, and Namespace/Tenant Administrators. Object storage users access ECS using the supported object protocols with clients that support those protocols. You can read more about ECS users and roles in [Add users and assign roles on page 54](#).

The portal uses the public ECS Management REST API. You can develop custom ECS clients using this API.

Log in to the ECS Portal

You can log into the ECS Portal from your browser by specifying the IP address or fully qualified domain name (FQDN) of any node, or the load balancer that acts as the front end to ECS.

Before you begin

- You can log into the ECS Portal if you are assigned to the System Admin, System Monitor, or Namespace Admin roles.
- A "root" user account, assigned to the System Admin role, is provided for initial access.

If you log in using the initial root credentials (root/ChangeMe), you must change the password for the root account immediately.

Your session ends when you close the browser, or log out. Logging out always closes the session. If you are unable to log in, contact your administrator.

You are automatically logged out after 2 hours of inactivity.

Procedure

1. Type the Public IP address of the first node in the system, or the address of the load balancer that has been configured as the front-end to ECS, in the following form:
`http://<node1_public_ip>`.

If this is the first time you are logging into a node with the root account, you will be prompted to change the password for the root account.

WELCOME TO EMC ELASTIC CLOUD STORAGE

⚠ Before continuing you must change your password
 You will be logged out after password change. Log back in to continue the checklist.

New Password

Confirm Password

Save

2. Click **Save**.

You are logged out and the standard login screen appears.

Elastic Cloud Storage

User Name *

Password *


Login

© 2015 EMC Corporation. All Rights Reserved.

EMC

3. Enter the **User Name** and **Password**.

4. To log out of the portal, locate the user menu in the upper-right corner of portal pages, select it, and choose **Logout**.



Change Password

When you are logged in at the ECS Portal, you can change your password.

System Admin, Namespace Admin, and System Monitor users have access to the **Change Password** page.

Procedure

1. At the ECS Portal, select **Settings > Change Password**
2. Enter a new password in the **Password** field and enter it again in the **Confirm Password** field.
3. Click **Save**.

Access to portal areas

The portal provides a left navigation menu and a right page area.

The System Admin can use all pages, a Namespace Admin can use a limited number of pages and perform only tenant-specific operations. A System Monitor can view all portal pages, but cannot create, edit, or delete portal settings. A system message displays if a user accesses a page or tries to perform an operation for which they do not have permissions.

The followings sections detail the permissions provided for different management users.

- [System Admin and System Monitor on page 16](#)
- [Namespace Admin on page 18](#)

System Admin and System Monitor

The following table lists the menu items that can be accessed and provides a link to documentation articles that provide more information on their use. System Monitors can view all that a System Admin can, but cannot make any changes.

Area	Menu	Operations Supported
Monitor	Metering	View object metering for namespace or bucket. For more information, see: Monitor metering data on page 120 .
	Events	View audit and alert events. For more information, see: About event monitoring on page 124 .
	Capacity Utilization	Monitor storage pool, node, and disk capacity. For more information, see: Monitor capacity on page 128 .

Area	Menu	Operations Supported
	Traffic Metrics	Monitor read and write bandwidth and latency. For more information, see: Monitor network traffic on page 132 .
	Hardware Health	Monitor storage node and disk status for each storage pool. For more information, see: Monitor hardware on page 136 .
	Node and Process Health	Monitor health of nodes and processes by memory and CPU utilization. For more information, see: Monitor node and process health on page 138 .
	Chunk Summary	Monitor chunks and chunks status. For more information, see: Monitor chunks on page 142 .
	Erasur Coding	Monitor erasure coding status. For more information, see: Monitor erasure coding on page 146 .
	Recovery Status	Monitor recovery status. For more information, see: Monitor recovery status on page 150 .
	Disk Bandwidth	Monitor disk bandwidth usage. For more information, see: Monitor disk bandwidth on page 152 .
	Geo Replication	Monitor geo-replication activity. For more information, see: Introduction to Geo-replication monitoring on page 156 .
Manage	Storage Pools	Enables the following operations: <ul style="list-style-type: none"> • Add a storage pool and specify the nodes that it comprises. • Add a VDC and define its connection details. • Configure a replication group by adding storage pools belonging to a VDC. For more information, see: Configure storage pools, VDCs, and replication groups on page 36 .
	Virtual Data Center	
	Replication Group	
	Authentication	Add an authentication provider that can authenticate domain users. See Add users and assign roles on page 54 .
	Namespace	Enables the following operations: <ul style="list-style-type: none"> • Create a new namespace. • Set quota for namespace. • Map object users into a namespace.

Area	Menu	Operations Supported
		For more information, see: Configure a namespace for a tenant on page 46
	Users	Enables the following operations: <ul style="list-style-type: none"> • Create object users for the namespace. • Edit object users. • Create secret keys. For more information, see: Add users and assign roles on page 54
	Buckets	Enables the following operations: <ul style="list-style-type: none"> • Create bucket. • Assign ACLs to bucket owner and object users. For more information, see: Create and configure buckets on page 86
	File	NFS File Access (not enabled for ECS 2.2)
Settings	Object Base URL	Set the Base URL to determine which part of object address is the bucket and namespace. For more information, see: Address ECS object storage and use the Base URL on page 108
	Change Password	Change own password. For more information, see: Change Password on page 16
	ESRS	Configure sending of alerts to EMC. For more information, see: Configuring ESRS on page 43 .
	Licensing	View license status and upload a license. For more information, see: Obtain and upload a license file to the ECS Portal on page 84
	About this VDC	View information about the VDC's nodes: node names, rack IDs, and software versions. For more information, see About this VDC on page 20 .

Namespace Admin

The following table lists the menu items that the Namespace Admin has permission to use and provides a link to documentation articles that provide more information on their use.


Area	Menu	Operations Supported
Monitor	Metering	View object metering for namespace or bucket. For more information, see Monitor metering data on page 120 .

Area	Menu	Operations Supported
Manage	Namespace	Edit the namespace. For more information, see Configure a namespace for a tenant on page 46
	Users	Enables the following operations: <ul style="list-style-type: none"> • Create object users for the namespace. • Edit object users • Create secret keys for object users For more information, see Add users and assign roles on page 54
	Bucket	Enables the following operations: <ul style="list-style-type: none"> • Create bucket. • Assign ACLs to bucket owner and object users. For more information, see Create and configure buckets on page 86
Settings	Change Password	Change own password. For more information, see Change Password on page 16

Ordering and searching tables in the portal

When a data set presented at the portal is large, and especially when it runs onto multiple pages, it is useful to reorder a table and to search for information in the table.

An example of a portal table is shown below.



Capacity Utilization

Storage Pools > sp1

Filter Current Clear Filter

Nodes	Disks	Usable Capacity	Used Capacity	Available Capacity	Node Status	Actions
10.241.51.201	60	326.95 TB	10.61 TB	316.34 TB	✓	History
10.241.51.205	59	321.5 TB	10.58 TB	310.93 TB	✓	History
10.241.51.204	60	326.95 TB	10.66 TB	316.29 TB	✓	History
10.241.51.203	60	326.95 TB	10.66 TB	316.29 TB	✓	History

Reordering Table Columns

You can reorder the rows in some tables based on the ordering of a selected column. A table column can be ordered by clicking on the table header.

Columns that contain textual data are sorted alphabetically. For example, if you select the Namespace field in the users table, that column will be ordered alphabetically and will

drive the ordering of rows. When you reenter the page, the default ordering will be applied. Similarly, refreshing the page will return the page to the default ordering.

Figure 1 Table Column with Sort Control Available

The screenshot shows a table with the following data:

Nodes
10.249.248.191
10.249.248.192
10.249.248.193
10.249.248.194

Other tables provide filter options to reduce table size. See [Monitoring basics on page 116](#).

Using Search

The Search facility enables some table rows to be filtered based on matching text strings.

As you type text in the Search box, rows that contain strings that match the search string are displayed. The order in which the rows that match the search criteria are displayed depends on the ordering applied by the table column ordering.

Refreshing a Page

A refresh control is provided on pages that contain table data. Using refresh will return the table to its default ordering.

About this VDC

Check software version numbers for the current node or any node in the VDC.

The **About this VDC** dialog lets you check the node names, rack IDs, and software version of the nodes in the VDC. The About page will give you information related to the node you are currently connected to. The Nodes page will give you information for all the nodes available in the VDC. The Nodes page will also identify any nodes that are not at the same software version as the node you are connected to.

Procedure

1. Select **Settings** > **About this VDC**.
2. Select **Nodes**.

Figure 2 Starred Rows Indicate Nodes with a Different Software Version than the Current Node

Current Node		Version	Rack ID
10.249.248.45		2.2.0.0	plum
Node	Node IP Address	ESX Object Version	Rack ID
<input checked="" type="radio"/> 10/249-248-45-100.vsan-010e.ses.lab.emc.com	10.249.248.45	2.2.0.0.7303180c8bc290e	plum
<input type="radio"/> 10/249-250-135.vsan-plum.ses.lab.emc.com	10.249.250.135	★ 2.2.0.0.731740c49a43ac	plum
<input type="radio"/> 10/249-248-47.vsan-010e.ses.lab.emc.com	10.249.248.47	2.2.0.0.7303180c8bc290e	plum
<input type="radio"/> 10/249-250-137.vsan-plum.ses.lab.emc.com	10.249.250.137	2.2.0.0.7303180c8bc290e	plum
<input type="radio"/> 10/249-248-46.vsan-010e.ses.lab.emc.com	10.249.248.46	2.2.0.0.7303180c8bc290e	plum
<input type="radio"/> 10/249-250-136.vsan-plum.ses.lab.emc.com	10.249.250.136	★ 2.2.0.0.731740c49a43ac	plum
<input type="radio"/> 10/249-248-48.vsan-010e.ses.lab.emc.com	10.249.248.48	2.2.0.0.7303180c8bc290e	plum
<input type="radio"/> 10/249-250-138.vsan-plum.ses.lab.emc.com	10.249.250.138	2.2.0.0.7303180c8bc290e	plum

CHAPTER 2

Use the Getting Started Checklist

- [Using the Getting Started Checklist](#)..... 24

Using the Getting Started Checklist


Use the Getting Started Checklist to guide you through the initial configuration of your ECS site.

The **Getting Started Checklist** is an app that overlays the portal and guides you through your initial configuration. The checklist appears when the portal detects that initial configuration is not complete. The checklist will automatically appear until you dismiss it. You can redisplay the checklist by selecting the Guide icon from the global menu at the top-right corner of all portal pages.



Note


Some parts of the initial configuration may be completed as part of the installation service.



1. The current step in the checklist.
2. A completed step.
3. An optional step. This step won't show a checkmark even if you have configured it.
4. Information about the current step.
5. Available actions will show here.
6. Dismiss the checklist.

A completed checklist will give you the option to browse the list again or recheck your configuration.

Figure 3 A Completed Checklist



Use the Getting Started Checklist

CHAPTER 3

Use the Dashboard


- [Use the Portal Dashboard](#).....28

Use the Portal Dashboard

The ECS Portal Dashboard provides critical information about the ECS processes on your local VDC.

The Dashboard

The Dashboard is the first page you encounter after login. To return to the Dashboard, select **Dashboard** in the left-hand menu.



Each panel title links to the portal monitoring page that shows deeper detail for the topic.

Global User menu

The Global User menu appears on each portal page.




Menu items include:

1. The Alert menu shows the most recent five alerts for the current VDC. The number indicates how many unacknowledged alerts are pending for the current VDC. The number displays "99+" if there are more than 99.
2. The Global Help icon brings up the online documentation for the current portal page.
3. The VDC menu shows the names of the current VDC. If your AD or LDAP credentials allow you to access more than one VDC, then you'll be able to switch your portal view to your other VDCs from here without re-entering your credentials.

4. The Guide icon brings up the Getting Started Checklist app.
5. The User menu shows the current user and allows you to log out.


Capacity

The Capacity panel displays the used and available storage on the local VDC as well as the percent used. Capacity takes into account ingested data, replicas, and system data.




Performance

The Performance panel shows you how network read and write operations are performing now and the average over the last 24 hours.




Data

The Data panel breaks down local VDC storage by user data and system data. Keep in mind that user data is the amount of your data ingested by ECS. The capacity used by your data will be affected by copies of your data and current system activities processing those copies.




Storage Efficiency

The Storage Efficiency panel shows how efficient the erasure coding (EC) process is currently working. The graph shows the progress of the current EC process, and the other values show the amount of EC data waiting for the EC process as well as the current rate of the EC process.



Geo Monitoring

The Geo Monitoring panel shows how much data from the local VDC is waiting for geo-replication as well as the rate of the replication. Recovery Point Objective (RPO) refers to the point in time in the past to which you can recover. The value here is the oldest data at risk of being lost if a local VDC fails before replication is complete. Failover Progress shows the progress of any active failover occurring in the federation involving the local VDC. Bootstrap Progress shows the progress of any active process to add a new VDC to the federation.



Nodes and Disks

The Nodes and Disks panel shows the health status (Good, Bad, or Suspect) of disks and nodes.

Node & Disks			
Disks		Nodes	
	479		7
	0		0
	0		1

Alerts

The Alert panel displays a count of critical alerts and errors. Click **Alerts** to see the full list of current events.

Alerts	
Critical Alerts 2	Errors 11

Use the Dashboard

PART 2

Configure and Manage

[Chapter 4, "Configure One or More Sites"](#)

[Chapter 5, "Configure a namespace"](#)

[Chapter 6, "Configure authentication and manage users"](#)

[Chapter 7, "Manage tenants"](#)

[Chapter 8, "Remove a site"](#)

[Chapter 9, "Manage licenses"](#)

[Chapter 10, "Create and manage buckets"](#)

[Chapter 11, "Configure NFS file access"](#)

[Chapter 12, "Set the Base URL"](#)

CHAPTER 4

Configure One or More Sites

- [Configure storage pools, VDCs, and replication groups](#).....36
- [Storage pools](#).....36
- [Virtual data centers \(VDCs\)](#).....39
- [Replication groups](#)..... 41
- [Configuring ESRS](#).....43

Configure storage pools, VDCs, and replication groups

Learn how to use the portal to create, modify, and delete storage pools, VDCs, and replication groups for single or federated deployments, and how to configure ConnectEMC for the object service.

Users must be assigned to the System Admin role to perform these procedures.

Storage pools

Storage pools let you organize storage resources based on business requirements. For example, if you require physical separation of data, you can partition the storage into multiple different storage pools.

Use the **Storage Pool Management** page available from **Manage > Storage Pools** to view the details of existing storage pools, to create new storage pools, to modify existing storage pools, and to delete storage pools.

Figure 4 Storage Pool Management page




Table 1 Storage pool properties

Field	Description
Name	The name of the storage pool.
# Nodes	The number of nodes assigned to the storage pool.
Status	The current state of the storage pool and of the nodes. Storage pool states are: <ul style="list-style-type: none"> Ready: At least four nodes are installed and all nodes are in the ready to use state. Not Ready: A node in the storage pool is not in the ready to use. Partially Ready: There are less than four nodes and all nodes are in the ready to use state.
Host Name	The fully qualified host name assigned to the node.
Node IP address	The public IP address assigned to the node.
Rack ID	The name assigned to the rack that contains the nodes.
Actions	Actions are:

Table 1 Storage pool properties (continued)

Field	Description
	<ul style="list-style-type: none"> • Edit: Use to change storage pool's name and the set of nodes included in the storage pool. • Delete: Use to delete storage pools. All nodes in storage pool must be removed before you can delete a storage pool. You cannot delete the system storage pool which is the first storage pool created. If the system storage pool has empty nodes, the empty nodes can be deleted if the number of nodes is greater than four.
Cold Storage	A storage pool with the Cold Storage property set uses an erasure coding (EC) scheme more efficient for infrequently accessed objects. Cold Storage is also known as a Cold Archive. Once a storage pool has been created, this setting cannot be changed.

Create storage pools

Use this procedure to assign nodes to storage pools. Storage pools must contain a minimum of four nodes. The first storage pool that is created is known as the system storage pool because it stores system metadata. The system storage pool cannot be deleted.

Procedure

1. From the portal, select **Manage** › **Storage Pools**.
2. Click **New Storage Pool**.

New Storage Pool

Name *

Cold Storage

Disabled Enabled

Note: Once Cold Storage has been enabled for a storage pool, it cannot be disabled.
Cold Storage functionality is applicable when the storage pool contains at least 6 nodes.

Available Nodes

search +

IP	Host

Selected Nodes * A minimum of 4 nodes is required

search ✕

IP	Host

3. Type the storage pool name. For example: **StoragePool11**.
4. Decide if this storage pool will be Cold Storage (also known as a Cold Archive). Cold storage contains infrequently accessed data. The ECS data protection scheme for cold storage is optimized to increase storage efficiency. Once a storage pool has been created, this setting cannot be changed.

Note

Cold storage requires a minimum hardware configuration of 6 nodes. See the [Data Protection](#) section of the ECS Planning Guide for more details.

5. Select the nodes to add to the storage pool from the Available Nodes list.
 - a. To select nodes one-by-one, click the + icon next for each node.
 - b. To select all available nodes, click the + icon at the top of the **Available Nodes** list.
 - c. To narrow the list of available nodes, type the node's public IP address or host name in the **search** field.
6. When you have completed the node selection, click **Save**.
7. Wait 10 minutes after the storage pool is in the **Ready** state before you perform other configuration tasks. This allows the storage pool time to initialize.

If you do not wait long enough, you receive the following error message: `Error 7000 (http: 500): An error occurred in the API Service. An`

error occurred in the API service.Cause: error insertVdcInfo. Virtual Data Center creation failure may occur when Data Services has not completed initialization.

If you receive this error, wait a few more minutes before attempting any further configuration.

Virtual data centers (VDCs)

VDCs are logical constructs. They are the top-level resource that represents the collection of ECS infrastructure to manage as a unit.

Use the **Virtual Data Center Management** page available from **Manage > Virtual Data Centers** to view VDC details, to create a new VDC, to modify existing VDCs, to delete VDCs and to federate multiple VDCs for a multi-site deployment. The following example shows the **Manage Virtual Data Center** page for a multi-site, federated deployment. It is configured with three sites. The VDCs are named vdc1, vdc2, and vdc3.

Figure 5 VDCManagement page

Name	Endpoints	Status	Actions
vdc1	10.249.249.151,10.249.249.152,10.249.249.153,10.249.249.154	Online	Edit
vdc2	10.249.249.155,10.249.249.156,10.249.249.157,10.249.249.158	Online	Edit
vdc3	10.249.248.155,10.249.248.156,10.249.248.157,10.249.248.158	Online	Edit

Table 2 VDC properties

Field	Description
Name	The VDC's name.
Endpoints	The public IP addresses of the nodes in the storage pools that comprise the VDC.
Status	States are: <ul style="list-style-type: none"> Online Permanently Failed: The VDC was deleted.
Actions	Actions are: <ul style="list-style-type: none"> Edit: Use to modify the VDC's name, the access key, and the public IP addresses of the nodes in the VDC's storage pools. Delete: Use to delete a VDC. The delete operation triggers permanent fail over of the VDC so you cannot add it back using the same name. You cannot delete a VDC that is part of a replication group until you first remove it from the replication group. You cannot delete a VDC when you are logged in to the VDC you are trying to delete.

Create a VDC for a single site

Use this procedure when you are creating a VDC for a single site deployment, or when you are creating the first VDC in a multi-site federation.

Before you begin

One or more storage pools are available and in the `Ready` state.

Procedure

1. From the ECS Portal, select **Manage** › **Virtual Data Center**.

2. Click **New Virtual Data Center**.

3. Type a name. For example: `vdc1`.

The name cannot have includes spaces or underscores.

4. Click **Get VDC Access Key**.

The VDC Access Key is used as a symmetric key for encrypting replication traffic between VDCs in a multi-site federation.

5. In the **Endpoints** text box, enter the public IP addresses of each node in the VDC's storage pools. Supply them as a comma-separated list.

6. Click **Save**.

Add a VDC to a federation

Use this procedure when you are adding a VDC (for example, VDC2) to an existing VDC (for example, VDC1) to create a federation.

Before you begin

Obtain the **ECS Portal** credentials for the root user, or for a user with system administrator credentials, to log in to both sites.

Ensure you have the list of public IP addresses for the nodes from the site you are adding (VDC2).

Ensure the site you are adding (VDC2) has a valid license uploaded and has at least one storage pool in the `Ready` state.

Procedure

1. Log in to the **ECS Portal** at the site you are adding (VDC2).

The default credentials are `root/ChangeMe`.

2. Select **Manage** › **Virtual Data Center**.

3. Click **Get VDC Access Key**.

4. Select the access key, and copy it using Ctrl-C to save it in the buffer.

5. Log out of the **ECS Portal** at the site you are adding (VDC2).

6. Log in to the **ECS Portal** of the first VDC (VDC1).

7. Select **Manage** › **Virtual Data Center**.

8. Click **New Virtual Data Center**.

9. Enter the VDC's name. For example: VDC2.

10. Click into the **Key** field and paste (CTRL-V) the Key you copied from the site you are adding (VDC2) from Steps 3 and 4 above.

11. Enter the public IP addresses of the site you are adding. Enter them as a comma-separated list.
12. Click **Save**.

Fail over a site/Delete a VDC

Use this procedure to delete a VDC. Deleting a VDC initiates site fail over when the VDC you are deleting is part of a multi-site federation.

If a disaster occurs, an entire VDC can become unrecoverable. ECS initially treats the unrecoverable VDC as a temporary site failure. If the failure is permanent, you must remove the VDC from the federation to initiate fail over processing which reconstructs and reprotects the objects stored on the failed VDC. The recovery tasks run as a background process. Review the recovery process by using the **Monitor > Geo Replication > Failover Processing**.

Procedure

1. Log in to one of the operational VDCs in the federation.
2. Go to **Manage > Replication Group**.
3. Click **Edit** for the replication group that contains the VDC to delete.
4. Click **Delete** in the row that contains the VDC and storage pool to remove.
5. Click **Save**.
6. Go to **Manage > VDC**. The status for the permanently removed VDC changes to `Permanently failed`.
7. Select **Delete** from the drop down in the row of the VDC to remove.
8. Click **Save**.

Replication groups

Replication groups are logical constructs that define where storage pool content is protected. Replication groups can be local or global. Local replication groups protect objects within the same VDC against disk or node failures. Global replication groups protect objects against disk, node, and site failures.

Use the **Manage Replication Groups** page to view replication group details, to create new replication groups, and to modify existing replication groups. You cannot delete replication groups in this release.

Figure 6 Manage Replication Groups page



Table 3 Replication Group properties

Field	Description
Name	The replication group name.

Table 3 Replication Group properties (continued)

Field	Description
VDC	The number of VDCs in the replication group and the names of the VDCs where the storage pools are located.
Storage Pool	The names of the storage pools and their associated VDCs.
Status	States are: <ul style="list-style-type: none"> • Online • Temp Unavailable: Replication traffic to this VDC has failed. If all replication traffic to the same VDC is in the Temp Unavailable state, further investigation about the cause of the failure is recommended.
Replicate to All Sites	A replication group with this feature disabled uses default replication. With default replication, data is stored at the primary site and a full copy is stored at a secondary site chosen from the sites within the replication group. The secondary copy is protected by triple-mirroring and erasure coding. This process provides data durability with storage efficiency. A replication group with this feature enabled makes a full readable copy of all objects to all sites (VDCs) within the replication group. Having full readable copies of objects on all VDCs in the replication group provides data durability and improves local performance at all sites at the cost of storage efficiency.
Actions	Edit: Use to modify the replication group name and the set of VDCs and storage pools in the replication group.

Create replication groups

Use this procedure to create replication groups.

To create global replication groups, choose storage pools from multiple VDCs.

Procedure

1. From the **ECS Portal**, select **Manage › Replication Group** .
2. Click **New Replication Group**.
3. Type a name. For example: **ReplicationGroup1**.
4. Decide if you want to enable **Replicate to All Sites** for this replication group. This option can only be enabled at the time of creation and cannot be disabled later.
5. Click **Add VDC**.
6. Select a Virtual Data Center and Storage Pool from the dropdown.
Repeat this step to add the VDCs and Storage pools required for object protection.
7. Click **Save**.

Configuring ESRS

This process describes steps to enable ESRS configuration on ECS. ECS version 2.2 and later requires ESRS Virtual Edition.

Procedure

1. Perform install/upgrade procedures through post-upgrade/install manual configurations.
2. Run the fcli to configure customer information and serial number on the cluster.

```

provo-melon:/opt/emc/caspian/fabric/cli # echo {"customer_data":
{"serial":"ABCDE00009", "customer_name":"ABCQE_MelonA",
"customer_email":"John.Smith@xyz.com"}} | bin/fcli lifecycle
cluster.setcustomer --body
{
  "status": "OK",
  "etag": 90
}
provo-melon:/opt/emc/caspian/fabric/cli # bin/fcli lifecycle
cluster.customer
{
  "status": "OK",
  "etag": 90,
  "customer_data": {
    "serial": "ABCD00009",
    "customer_name": "ABCQE_MelonA",
    "customer_email": "John.Smith@xyz.com"
  }
}
provo-melon:/opt/emc/caspian/fabric/cli # bin/fcli lifecycle
alert.callhomeenabled
{
  "status": "OK",
  "etag": 90,
  "callhome_enabled": true
}

```

3. You can add or update an ESRS server on the ECS Portal. If you already have an ESRS server enabled, you must delete it, then add the new server. Go to **Settings > ESRS**, and then add the following information: **FQDN/IP, PORT, Username, Password**.

FOB-based passwords are not supported when configuring ESRS with ECS. Use your customer support.emc.com credentials.

Note

For 2.2, you must delete, then add any existing ESRS server that you edit with the ECS Portal. Editing the server is not functional for 2.2.

CHAPTER 5

Configure a namespace

- [Introduction](#)..... 46
- [Understanding tenants](#).....46
- [Understanding namespace settings](#)..... 47
- [Working with namespaces at the ECS portal](#)..... 50
- [Create and configure a namespace](#).....50

Introduction

Namespaces provide the mechanism by which multiple tenants can access the ECS object store and ensure that the objects and buckets written by users of a tenant are segregated from users of other tenants.

This article introduces some concepts around tenants and namespace settings:

- [Understanding tenants on page 46](#)
- [Understanding namespace settings on page 47](#)
- [Working with namespaces at the ECS portal on page 50](#)

and describes the operations required to configure a namespace using the ECS Portal:

- [Create and configure a namespace on page 50](#)

While the configuration operations described in this article use the ECS portal, the concepts described in [Understanding tenants on page 46](#) and [Understanding namespace settings on page 47](#) apply whether you are using the portal or the REST API.

Understanding tenants

ECS supports access by multiple-tenants, where each tenant is defined by a namespace and the namespace has a set of configured users who can store and access objects within the namespace.

Namespaces are global resources in ECS and a System Admin or Namespace Admin accessing ECS at any federated VDC can configure the namespace settings. In addition, object users assigned to a namespace are global and can access the object store from any federated VDC.

The key characteristic of a namespace is that users from one namespace cannot access objects belonging to another namespace. In addition, ECS enables an enterprise to configure namespaces and to monitor and meter their usage, and enables management rights to be granted to the tenant so that it can perform configuration and monitoring and metering operations.

It is also possible to use buckets as a means of creating sub-tenants. The bucket owner is the sub-tenant administrator and can assign users to the sub-tenant using access control lists. However, sub-tenants do not provide the same level of segregation as tenants; any user belonging to the tenant could be assigned privileges on a sub-tenant, so care must be taken when assigning users.

The following scenarios are supported:

Enterprise single tenant

All users access buckets and objects in the same namespace. Sub-tenants (buckets) can be created to allow a subset of namespace users to access the same set of objects. A sub-tenant could be a department within the enterprise.

Enterprise multi tenant

Different departments within an organization are assigned to different namespaces and department users are assigned to each namespace.

Cloud Service Provider single tenant

A single namespace is configured and the Service Provider provides access to the object store for users within the enterprise or outside the enterprise.

Cloud Service Provider multi tenant

The Service Provider assigns namespaces to different companies and assigns an administrator for the namespace. The namespace administrator for the tenant can then add users and can monitor and meter the use of buckets and objects.

The features provided to enable management of tenants are described in [Manage a tenant on page 74](#).

Each tenant has access to the replication groups made available by the System Admin. Depending on the access patterns of a tenant, they may require replication groups that include sites in specific geographies. For example, if a client tenant is located in China, they might prefer to access replication groups that include VDCs located in China.

Understanding namespace settings

A namespace provides a mechanism by which objects and buckets can be segregated so that an object in one namespace can have the same name as an object in another namespace. ECS will always know which object is required by the namespace qualifier. The namespace is also configured with attributes that define which users can access the namespace and what characteristics the namespace has. You can think of an ECS namespace as a tenant.

Users with the appropriate privileges can create buckets, and can create objects within buckets, in the namespace.

The way in which namespace and bucket names are used when addressing objects in ECS is described in [Address ECS object storage and use the Base URL on page 108](#).

An ECS namespace has the following attributes:

Field	Description	Can be Edited
Name	The name of the namespace. This name must be in all lowercase characters.	No
Namespace Admin - User	User Id of one or more users who you want to assign to the Namespace Admin role; a list of users should be comma separated. Namespace Admins can be local or domain users. If you want the Namespace Admin to be a domain user, you will need to ensure that an authentication provider has been added to ECS. Refer to Add users and assign roles on page 54 for details.	Yes
Namespace Admin - Domain Group	Domain group that you want to assign to the Namespace Admin role. Any member, once authenticated, will be placed in the Namespace Admin role for the namespace. The group must be assigned to the namespace by setting the Domain User Mappings for the namespace. To use this feature you will need to ensure that an authentication provider has been added to ECS. Refer to Add users and assign roles on page 54 for details.	Yes
Replication Group	The default replication group for the namespace.	Yes
Namespace Quota	Enables quotas for the namespace. The quotas will apply to the total storage used by the namespace. Soft and hard limits can be defined to notify that a defined limit has been reached and to	Yes

Field	Description	Can be Edited
	block access to the namespace when maximum storage is reached.	
Bucket Quota (Bucket Default)	Defines a default quota that will be applied to buckets created in this namespace. The default quota is a Block Quota which, when reached, will prevent write/update access to the bucket. The default bucket quota is applied at bucket create time, so changing the default bucket quota will not change the bucket quota for already created buckets.	Yes
Server-side Encryption (Bucket Default)	Defines a default value for Server-side Encryption that will apply to buckets created in this namespace. Server-side Encryption is also known as Data At Rest Encryption or D@RE. This feature encrypts data inline before storing it on ECS disks or drives. This encryption prevents sensitive data from being acquired from discarded or stolen media. If the namespace enables encryption, then all its buckets will be encrypted unless you disable encryption for the bucket at creation time. For a complete description of the feature, see the <i>ECS Security Configuration Guide</i> .	No
Access During Outage (Bucket Default)	Defines a default value for Access During Outage that will be applied to buckets created in this namespace.	Yes
Compliance (Bucket Default)	<p>ECS has object retention features enabled or defined at the object-, bucket-, and namespace-level. Compliance strengthens these features by limiting changes that can be made to retention settings on objects under retention.</p> <p>Compliance rules include:</p> <ul style="list-style-type: none"> • Compliance is enabled at the namespace-level. This means that all buckets in the namespace must have a retention period greater than zero. • Compliance can only be enabled on a namespace when the namespace is created. (Compliance cannot be added to an existing namespace.) • Compliance cannot be disabled once enabled. • All buckets in a namespace must have a retention period greater than zero. <hr/> <p>Note</p> <p>If you have an application that assigns object-level retention periods, do not use ECS to assign a retention period greater than the application retention period. This will lead to application errors.</p> <hr/> <ul style="list-style-type: none"> • A bucket cannot be deleted while it contains data regardless of its retention setting. • Using the Infinite option on a bucket mean objects in the bucket in a Compliance-enabled namespace can never be deleted. 	No

Field	Description	Can be Edited
	<ul style="list-style-type: none"> The retention period for an object cannot be deleted or shortened. Therefore, the retention period for a bucket cannot be deleted or shortened. Object and bucket retention periods can be increased. No feature can delete an object under retention. This includes the CAS privileged-delete permission. 	
Retention Policies	<p>Enables one or more retention policies to be added and configured.</p> <p>A namespace can have a number of associated retention policies, where each policy defines a retention period. By applying a retention policy to a number of objects, rather than applying a retention period directly, a change the retention policy will cause the retention period to be changed for the complete set of objects to which the policy has been applied. A request to modify an object that falls before the expiration of the retention period will be disallowed.</p> <p>It is also possible to specify retention policies and specify a quota for the namespace. Further information on using these features is provided in Retention periods and policies on page 75.</p>	Yes
Domain	<p>Enables AD/LDAP domains to be specified and the rules for including users from the domain to be configured.</p> <p>Domain users can be assigned to ECS management roles. In addition, users belonging to the domain can use the ECS self-service capability to register as object users.</p> <p>The mapping of domain users into a namespace is described in Understanding the mapping of users into a namespace on page 70</p>	Yes

The following attribute can be set using the ECS Management REST API, not from the ECS Portal.


Allowed (and Disallowed) Replication Groups

Enables a client to specify which replication groups can be used by the namespace.

Working with namespaces at the ECS portal

The namespace portal page, **Manage** > **Namespace**, enables namespaces to be created and provides a namespace table which lists the namespaces that exist and allows them to be edited.

Figure 7 Namespace management page



The namespace table comprises the following fields:

Field	Description
Name	Name of the namespace.
Replication Group	Default replication group for the namespace.
Notification Quota	Quota limit at which notification is generated.
Max Quota	Quota limit at which writes to the namespace will be blocked.
Encryption	Specifies if D@RE server-side encryption is enabled for the namespace.
Actions	Actions that can be performed on the namespace. Edit and Delete actions are available.

Create and configure a namespace

You can create a new namespace or change the configuration of an existing namespace at the **Manage** > **Namespace** page.

Before you begin

- To perform this operation, you must be assigned to the System Admin role in ECS.
- A replication group must exist. The replication group provides access to storage pools in which object data is stored.
- If you want to allow domain users to access the namespace, an authentication provider must have been added to ECS. In addition, if you intend to configure domain object users or a domain group, you should plan how you want to map users into the namespace. You can refer to [Add users and assign roles on page 54](#) for more information on mapping users.

You should ensure you are familiar with the general information about namespaces provided in [Understanding namespace settings on page 47](#).

Procedure

1. At the ECS portal, select **Manage** > **Namespace**

- To create a new namespace, select **New Namespace**. To edit the configuration of an existing namespace, choose the **Edit** action associated with the existing namespace.

The screenshot shows the 'New Namespace' configuration page. It features several input fields and toggle switches:

- Name**: A text input field.
- User Admin**: A text input field with an example: *Ex.: user1, user2@foo.com*.
- Domain Group Admin**: A text input field with an example: *Ex.: group1@foo.com, group2@test.com*.
- Replication Group**: A dropdown menu with 'rep1' selected.
- Namespace Quota**: A toggle switch currently set to 'Disabled'.
- Default Bucket Quota**: A toggle switch currently set to 'Disabled'.
- Server-side Encryption**: A toggle switch currently set to 'Disabled'.
- Access During Outage**: A toggle switch currently set to 'Disabled'.
- Compliance**: A toggle switch currently set to 'Disabled'.
- Retention Policies**: A table with columns 'Name', 'Value', and 'Actions'. An 'Add' button is located to the right of the table.
- At the bottom, there is a '+ Domain' button and 'Save' and 'Cancel' buttons.

- Set the namespace administrator by entering a domain or local user in the **User Admin** field and/or adding a domain group in the **Domain Group Admin** field.

Multiple users or groups can be added as comma separated lists.

- Specify appropriate value for each of the bucket default fields.

The following controls set the default value when a bucket is created using an object client:

- Default Bucket Quota
- Access During Outage
- Compliance

- Decide if this namespace requires Server-side Encryption. If **Yes**, every bucket in the namespace will have Server-side encryption enabled and every object in the buckets

will be encrypted. If you select **No**, you can still apply Server-side encryption to individual buckets in the namespace at the time of creation.

6. If you want to set a quota for the namespace:

- a. Set the **Namespace Quota** control to **Enabled**.
- b. Choose **Notification Only** or **Block Access**

If you choose to block access when a specified storage limit is reached, you can also specify a percentage of that limit at which a notification will be sent.

7. Add and Configure Retention Policies.

- a. In the Retention Policies area, select **Add** to add a new policy.
- b. Enter a name for the policy.
- c. Specify the period for the Retention Policy.

This can be a value in minutes or you can select the **Infinite** checkbox to ensure that buckets to which this retention policy is assigned are never deleted.

8. Specify an AD/LDAP domain whose users can log in to ECS and perform administration tasks for the namespace.

Enter the name of the domain and specify groups and attributes to provide finer grained control over the domain users that will be allowed to access ECS in the current namespace.

To perform more complex mappings using groups and attributes, you should refer to [Add users and assign roles on page 54](#)

9. Select **Save**.

CHAPTER 6

Configure authentication and manage users

- [Introduction](#)..... 54
- [Understanding users and roles in ECS](#)..... 54
- [Working with the users at the ECS Portal](#)..... 58
- [Working with the authentication providers at the ECS Portal](#)..... 64
- [Understanding the mapping of users into a namespace](#)..... 70

Introduction

This article describes the types of users supported by ECS and the roles to which they can be assigned.

It introduces the main concepts around ECS users and roles:

- [Understanding users and roles in ECS on page 54](#)
- [Working with the users at the ECS Portal on page 58](#)

and then describes how to add management users or object users:

- [Add a new object user on page 60](#)
- [Add a domain user as an object user on page 61](#)
- [Create a local management user or assign a domain user to a management role on page 61](#)
- [Create a namespace administrator on page 63](#)

In addition, it shows you how you can set up an authentication provider and perform the mapping of domain users into a namespace:

- [Add an authentication provider on page 65](#)
- [Map domain users into a namespace on page 71](#)

Understanding users and roles in ECS

ECS defines different user types and roles to determine access to ECS management facilities and to the object store.

The main concepts relating to users and roles are described in the following topics:

- [Users in ECS on page 54](#)
- [User roles on page 55](#)
- [Domain and local users on page 57](#)
- [User scope: global or namespace on page 57](#)

Users in ECS

ECS requires two types of user: management users, who can perform administration of ECS, and object users, who access the object store to read and write objects and buckets using the supported data access protocols (S3, EMC Atmos, OpenStack Swift, and CAS).

Management users can access the ECS Portal. Object users cannot access the ECS Portal but can access the object store using clients that support the ECS data access protocols.

Management users and object users are stored in different tables and their credentials are different. Management users require a local username and password, or a link to a domain user account. Object users require a username and a secret key. Hence you can create a management user and an object user with the same name, but they are effectively different users as their credentials are different.

In addition, management and object user names can be unique across the ECS system or can be unique within a namespace. This is referred to as user scope and is described in: [User scope: global or namespace on page 57](#).

Details of the supported user types are provided in the following sections:

- [Management Users on page 55](#)
- [Object users on page 55](#)
- [Root user on page 55](#)

Management Users

Management users can perform the configuration and administration of the ECS system and of tenants configured in ECS.

Management users can be local users whose credentials are stored by ECS and are authenticated by ECS against the locally held credentials, or they can be domain users defined in Active Directory (AD) or Lightweight Directory Access Protocol (LDAP) and authenticated against users held in those systems. You can find out more about domain and local users in [Domain and local users on page 57](#).

Management users are not replicated across geo-federated VDCs.

Object users

Object users are end-users of the ECS object store and access it through object clients using the ECS supported object protocols (S3, EMC Atmos, Openstack Swift, and CAS). Object users can also be assigned Unix-style permissions to access buckets exported as filesystems for HDFS.

Object users are defined by a username and a secret key that can be used to access the object store. Usernames can be local names or can be domain-style user names that include a "@" in their name.

A management user can create an object user account and can assign a secret key to the object user account when the account is created or at any time thereafter. When created by a management user, the object users secret key is distributed by email or other means.

For domain users, a secret key can be obtained by the object user using the ECS self-service capability, using a client that talks to the ECS REST API (object users do not have access to the ECS portal). You can read more about domain users in: [Domain and local users on page 57](#), and you can refer to *Data Access Guide: Obtain secret key to access object storage* for information on creating a secret key.

Object users are global resources, so an object user created at a VDC can be given privileges to read and write buckets, and objects, within the namespace to which they are assigned, from any VDC.

Root user

The root user is available at system initialization and is pre-assigned to the System Admin role.

The root user should only be used for initial access to the system. On initial access, the root user password should be changed at the **Settings > Password** page and one or more new System Admin accounts should be created.

From an audit perspective, it is important to know which user carried out changes to the system, so root should not be used, and each System Admin user should have their own account.

User roles

ECS defines roles to determine the operations that a user account can perform at the ECS Portal or when accessing ECS using the ECS Management REST API. Management users

and groups can be assigned to administration roles in ECS and can be either local users or domain users. Roles can also be assigned to Active Directory group names.

The following management roles are defined:

- [System Admin on page 56](#)
- [System Monitor on page 56](#)
- [Namespace Admin on page 56](#)

System Admin

The System Admin role can configure ECS and specify the storage used for the object store, how the store is replicated, how tenant access to the object store is configured, and which users have permissions on an assigned namespace.

The System Admin can also configure namespaces and perform namespace administration, or can assign a user who belongs to the namespace as the Namespace Admin.

The System Admin has access to the ECS Portal and system administration operations can also be performed from programmatic clients using the ECS Management REST API.

Because management users are not replicated across site, a System Admin must be created at each VDC that requires one.

System Monitor

The System Monitor role can view all ECS Portal data, but cannot make any changes.

The System Monitor role can view all ECS Portal data, but cannot provision the ECS system. The monitor cannot create, update, or delete storage pools, replication groups, namespaces, buckets, users and so on through the portal or ECS management API. Monitors cannot modify any other portal setting except their own passwords.

Because management users are not replicated across sites, a System Monitor must be created at each VDC that requires one.

Namespace Admin

The Namespace Admin is a management user who can access the ECS Portal to configure namespace settings, such as quotas and retention periods, and can map domain users into the namespace and assign local users as object users for the namespace. Namespace Admin operations can also be performed using the ECS Management REST API.

A Namespace Admin can only be the administrator of a single namespace.

Because authentication providers and namespaces are replicated across sites (they are ECS global resources), a domain user who is a Namespace Admin can log in at any site and perform namespace administration from that site.

Local management accounts are not replicated across sites, so a local user who is a Namespace Admin can only log in at the VDC at which the management user account was created. If you want the same username to exist at another VDC, the user must be created at the other VDC. As they are different accounts, changes to a same-named account at one VDC, such as a password change, will not be propagated to the account with the same name at the other VDC.

Domain and local users

ECS provides support for local and domain users.

Local users are user accounts whose credentials are stored by ECS. Both management users and object users can be defined locally to ECS. In the case of object users, the credentials are global resources and are available at all ECS VDCs.

Local users make it very simple to start using ECS, however, the use of AD/LDAP enables an existing user database to be leveraged and allows a large number of users to be given access to the object store without having to create accounts for them.

Domain users are users defined in an Active Directory AD/LDAP database and ECS must talk to the AD or LDAP server to authenticate user login request. ECS uses a construct called an authentication provider to supply the credentials it needs to talk to the AD/LDAP server and to specify the domains and groups that should be made available to ECS.

Domain users are defined in the form `user@domain.com` and ECS will attempt to authenticate user names in that form using the authentication providers that have been configured. User names without `@` will be authenticated against the local user database.

Domain users assigned to management roles can be authenticated against their AD/LDAP credentials to allow them to access ECS and perform ECS administration operations. Administration operations can be performed from the ECS Portal or using the ECS Management API.

Domain users can also be assigned as object users. To save the administrative overhead of manually creating large numbers of object user accounts in ECS, a self-service capability is provided that allows ECS to authenticate domain users and automatically add them as object users and assign a secret key to them.

To make use of this, a domain user must be mapped into a namespace and ECS provides a mechanism for mapping domain users into a namespace based on their domain and group membership and on attributes associated with their account.

User scope: global or namespace

The scope of object users depends on the user scope that has been set. The setting affects all users, in all namespaces across all federated VDCs

The user scope can be either GLOBAL or NAMESPACE. In global scope, object user names are unique across all VDCs in the ECS system. In namespace scope, object user names are unique within a namespace, so the same object user account names can exist in different namespaces.

The default setting is GLOBAL. If you intend to use ECS in a multi-tenant configuration and you want to ensure that tenants are not prevented from using names that are in use in another namespace, you should change this default configuration to NAMESPACE.

Note

The user scope setting must be made before the first object user is created.

Setting the User Scope

The user scope can be set using the PUT /config/object/properties API and passing the user scope in the payload. An example of a payload that sets the user_scope to NAMESPACE is shown below.

```
PUT /config/object/properties/

<property_update>
  <properties>
    <properties>
      <entry>
        <key>user_scope</key>
        <value>NAMESPACE</value>
      </entry>
    </properties>
  </property_update>
```

Working with the users at the ECS Portal

The ECS Portal provides a **Manage > Users** page to enable local users to be created and assigned as object users for a namespace. It also enables system administrators to create local management users and assign them to administration roles and to assign domain users to administration roles.

The **Manage > Users** page provides two sub-pages:

- [Object Users View on page 58](#)
- [Management Users View on page 59](#)


The Management Page is only accessible if you are a System Admin (or root user) for ECS.

Object Users View

The Object Users view provides an Object Users table that lists the local users that have been created, the namespace to which the users have been assigned, and the actions that can be performed on the user.

If you are a System Admin you will see the object users for all namespaces. If you are a Namespace Admin, you will only see the users belonging to your namespace.

The Object Users view is shown below.



The Object Users table provides access to the following information and operations.

Attribute	Description
Name	The name of the user.
Namespace	The namespace to which the user is assigned.
Actions	Provides a selection menu for the actions that are available. The actions that are available are: Edit and Delete .


The Object Users pane additionally provides access to the the following controls:

Control	Description
New Object User	The New Object User button enables an object user to be added.

Management Users View

The Management Users view provides a Management Users table that lists the management users that have been created and the actions that can be performed on the user. This page is only visible to users with the System Admin role.

The Management Users view is shown below.



The Management Users table provides access to the following information and operations.

Column	Description
Name	The name of the user.
Actions	Provides a selection menu for the actions that are available. The actions that are available are: Edit and Delete .

In addition, the Management Users view provides the following controls:

Control	Description
New Management User	The New Management User button enables the addition of a management user that may be assigned as the System Admin role for ECS.

Add a new object user

You can create new local users and configure them to use the supported object access protocols. Once created, you can edit a user configuration by adding or removing access to an object protocol, or by creating a new secret key for the user.

Before you begin

- If you are an ECS System Admin, you can assign users for any namespace.
- If you are a Namespace Admin, you can assign users for the namespaces for which you are the administrator.
- If you want your domain users to be enabled as object users you should refer to [Add a domain user as an object user on page 61](#).
- When assigning a password for a Swift user, the user will be added to the Swift Admin group.

Note

Do not use the ECS Portal to perform this operation if you want users to be assigned to different Swift groups.

You can refer to [Working with the users at the ECS Portal on page 58](#) for information about the **Manage › Users** page.

Procedure

1. At the ECS Portal, select **Manage › Users**.

The Object Users Page is shown by default and displays the Object Users table which lists the local users that have been created and the namespace to which they are assigned.

2. Select **New Object User**.

The New Object User page is displayed.

3. Enter a name for the user.

This is a name for a local user that will be created.

You can use domain-style names that include "@". For example, "some.name@emc.com". However, this is a convenience to enable you to keep names unique and consistent with AD names, authentication is performed using a secret key assigned to the username, not through AD or LDAP.

Note

User names must be lowercase letters, numbers and any of the following characters: ! # \$ & ' () * + , - . / : ; = ? @ _ ~

4. Select the namespace to which the local user will be assigned.

Once you have selected the namespace, you can **Save** the user and return later to edit the user and assign a secret key to access an object protocol. Alternatively, you can select **Add Passwords** and specify passwords or secret keys to access the ECS object protocols.

5. To set up secret keys for the user, select **Add Passwords**.

6. For each of the object protocols that you want to use to access the ECS object store, enter or generate a key for use in accessing the S3, Swift, or CAS, and save the key.

Select **Add Password** to save the key.

7. Specify a password for each of the object interfaces that you want the user to be able to access.

For S3 and CAS you can generate the password.

8. The secret keys and passwords are saved automatically and you can click the **Close** button to return to the Users page.

Add a domain user as an object user

You can configure domain users so that they can access ECS and generate secret keys for themselves and, by doing so, add themselves as object users.

Before you begin

Procedure

1. Ensure an authentication provider that connects to the appropriate AD/LDAP system has been configured.

Adding an authentication provider must be performed by a System Admin and is described in [Add an authentication provider on page 65](#).

2. Map domain users into the namespace as described in [Map domain users into a namespace on page 71](#).

This can be performed by the Namespace Admin.

3. Allow users to create secret keys using the instructions in [Data Access Guide: Obtain secret key to access object storage](#).

Create a local management user or assign a domain user to a management role

You can add a local management user and assign a local management user or a domain user to a management role from the ECS Portal. Management users are required to perform system-level administration (VDC administration) and namespace administration. Where a user is no longer needed to perform administration operations, you can remove the role assignment.

Before you begin

- You must be a System Admin to create a local management user or assign a management role.
- The ECS root user has the System Admin role by default and can perform the initial assignment of a user to the System Admin role.
- If you want to assign a domain user to a management role, you must first ensure that an authentication provider has been added. See [Add an authentication provider on page 65](#).
- If you want to assign a Namespace Admin, you must create a management user using the operation defined here and perform the role assignment at the portal Namespace page (see [Configure a namespace for a tenant on page 46](#)). The user will not be able to log in until they have been assigned to the Namespace Admin role (or the System Admin role).

You can refer to [Working with the users at the ECS Portal on page 58](#) for information about the **Manage** > **Users** page.

Procedure

1. At the ECS Portal, select **Manage** > **Users**.

The Object Users Page is displayed by default and you need to change to the Management Users page.

2. Select **Management Users**.

The Management Users page is displayed which shows any users that have currently been assigned and provide a **New Management User** button.

3. Select **New Management User**.

The New Management User pages is displayed which enables you to create a local user and assign the new user to the management role, or assign a domain user to the management role.

4. Select Local User or AD/LDAP User.

For a local user you will need to define a password; for a domain user, the user and password credentials that ECS will use to authenticate a user are held in AD/LDAP, so you don't need to define a password.

5. Enter the name the user.

If you have selected AD/LDAP, the user must exist and have been made available by adding an authentication provider to ECS.

If you select local user, a new local management user will be created.

Note

User names must be lowercase letters, numbers and any of the following characters: ! # \$ % & ' () * + , - . / : ; = ? @ _ ~

6. To assign the user to the System Monitor role, select **Yes** at the System Monitor selector.

7. If you want to assign the user to the System Admin role, select **Yes** at the System Administrator selector.

If you are creating a management user who will be assigned to the Namespace Admin role for a namespace, you should leave this as **No**.

If you select **Yes**, but at a later date you want to remove System Administrator privileges from the user, you can edit the user settings and change this to **No**.

8. Select **Save**.

Assign an Active Directory group name to the system admin or system monitor role

You can assign an AD domain group to the system admin or system monitor role from the ECS Portal. When an AD domain group is assigned a management role, all users in the AD group will have that role.

Before you begin

- You must be a system admin to assign a management role.
- To assign an AD domain group to a management role, you must first ensure that an authentication provider has been added. See [Add an authentication provider on page 65](#).

You can refer to [Working with the users at the ECS Portal on page 58](#) for information about the **Manage > Users** page.

Note

LDAP groups are not supported in ECS.

Procedure

1. At the ECS Portal, select **Manage** > **Users**.
The Object Users page displays by default. Change to the Management Users page.
2. Select **Management Users**.
The Management Users page is displayed which shows any users that have currently been assigned and provides a **New Management User** button.
3. Select **New Management User**.
The New Management User page displays.
4. Select AD/LDAP User or Group.
For a domain user, the user and password credentials that ECS uses to authenticate a user are held in AD/LDAP, so you don't need to define a password.
5. Change the **User** dropdown to **Group**.
6. Fill in the Group Username field with your complete AD domain group name including the domain. For example: ITadmins@somecorp.com.
7. To assign the group to the system monitor role, select **Yes** at the System Monitor selector.
8. To assign the group to the system admin role, select **Yes** at the System Administrator selector.
9. If you select **Yes** to either of these roles, you can remove the role from the group later by changing the setting to **No**.
10. Select **Save**.

Create a namespace administrator

You can assign a local or domain user as a Namespace Admin.

Before you begin

- You must be a System Admin to create a management user and assign a user to the Namespace Admin role.

You can refer to [Working with the users at the ECS Portal on page 58](#) for information about the **Manage** > **Users** page.

Procedure

1. If you want to assign a local management user to the Namespace Admin role, you need to create a management user as described in [Create a local management user or assign a domain user to a management role on page 61](#).
If you want to assign a domain user to the Namespace Admin role, you do not need to explicitly assign the user to a management role.
2. At the **Manage** > **Namespace** page.
 - a. Select the **Edit** action for the namespace.
 - b. Add the user to the Namespace Admin field. If there is more than one Namespace Admin, their usernames should be a comma separated list.
A user can only be assigned as the Namespace Admin for a single namespace.

- c. **Save** the namespace.

You can read more about configuring a namespace in: [Configure a namespace for a tenant on page 46](#).

Assign an Active Directory group name to the namespace admin role

You can assign an AD domain group to the namespace admin role from the ECS Portal. When an AD domain group is assigned a management role, all users in the AD group will have that role.

Before you begin

- You must be a system admin to assign a namespace admin role.
- To assign an AD domain group to a namespace admin, you must first ensure that an authentication provider has been added. See [Add an authentication provider on page 65](#).

You can refer to [Working with the users at the ECS Portal on page 58](#) for information about the **Manage > Users** page.

Note

LDAP groups are not supported in ECS.

Procedure

1. At the ECS Portal, select **Manage > Namespace**.
2. Select a namespace and select **Edit**, or select **New Namespace**.
3. Fill in the Domain Group Admin field with your complete AD domain group name including the domain. For example: FinanceAdmins@somecorp.com. To add more than one domain group, separate the names with commas.

Note

An AD domain group can only be the namespace admin for one namespace.


4. Complete your configuration and select **Save**.

Working with the authentication providers at the ECS Portal

The ECS Portal provides a **Manage > Authentication** page to enable authentication providers to be added.

The Authentication Provider Page is only accessible if you are a System Admin (or root user) for ECS.

The Authentication Provider Page provides an Authentication Provider table that lists the authentication provider that have been created. An example is shown below.



The table provides access to the following information and operations.

Attribute	Description
Name	The name that has been given to the authentication provider.
Type	Indicates whether the authentication provider is an Active Directory (AD) or Lightweight Directory Access Protocol (LDAP) server.
Domains	Domains that the authentication provider provides access to.
Enabled	Indicated whether the authentication provider is currently Enabled or Disabled.
Actions	Provides a selection menu for the actions that are available. The actions that are available are: Edit and Delete .

The Authentication Provider Page additionally provides access to the following controls:

Control	Description
New Authentication Provider	The New Authentication Provider button enables an authentication provider to be added.

Add an authentication provider

User authentication for domain users is performed using one or more authentication providers added to ECS. An authentication provider is a construct that enables ECS to connect an AD/LDAP server and identifies the domains and groups that the AD/LDAP should make available to ECS.

Before you begin

- To add an authentication provider you must be assigned to the System Admin role in ECS. The root user has the System Admin role.
- You need access to the authentication provider information listed in [Authentication provider settings on page 66](#). Note especially the requirements for the Manager DN user.

Procedure

1. At the ECS Portal, select **Manage** > **Authentication** > **New Authentication Providers**.
2. Enter values for the attributes. Refer to [Authentication provider settings on page 66](#)
3. **Save**.

- To verify the configuration, add a user from the authentication provider at **Manage > Users > Management Users**, then try to log in as the new user.

Authentication provider settings

You need to provide certain information when adding or editing an authentication provider.

Table 4 Authentication provider settings

Field name	Description and requirements
Name	The name of the authentication provider. You can have multiple providers for different domains.
Description	Free text description of the authentication provider.
Type	Active Directory or LDAP.
Domains	<p>Active Directory and LDAP allow administrators to organize objects of a network (such as users, computers, and devices) into a hierarchical collection of containers.</p> <p>Domains are a collection of administratively defined objects that share a common directory database, security policies, and trust relationships with other domains. In this way, each domain is an administrative boundary for objects. A single domain can span multiple physical locations or sites and can contain millions of objects.</p> <p>A typical entry in this field of the authentication provider would look like this:</p> <p><code>mycompany.com</code></p> <p>If an alternate UPN suffix is configured in the Active Directory, the Domains list should also contain the alternate UPN configured for the domain. For example, if <code>myco</code> is added as an alternate UPN suffix for <code>mycompany.com</code>, then the Domains list should contain both <code>myco</code> and <code>mycompany.com</code>.</p>
Server URLs	<p>ldap or ldaps (secure LDAP) with the domain controller IP address. Default port for ldap is 389 and ldaps is 636.</p> <p>Usage: one or more of</p> <p>ldap://<Domain controller IP>:<port> (if not default port)</p> <p>or</p> <p>ldaps://<Domain controller IP>:<port> (if not default port)</p> <p>If the authentication provider supports a multidomain forest, use the global catalog server IP and always specify the port number. Default is 3268 for ldap, 3269 for ldaps.</p> <p>Usage: ldap(s)://<Global catalog server IP>:<port></p>
Manager DN	Indicates the Active Directory Bind user account that ECS uses to connect to Active Directory or LDAP server. This account is used to search Active Directory when a ECS administrator specifies a user for role assignment, for example.

Table 4 Authentication provider settings (continued)

Field name	Description and requirements
	<p>Requirement:</p> <p>This user must have Read all inetOrgPerson information in Active Directory. The InetOrgPerson object class is used in several non-Microsoft, Lightweight Directory Access Protocol (LDAP) and X.500 directory services to represent people in an organization.</p> <p>To set this privilege in Active Directory, open Active Directory Users and Computers, right click on the domain, and select Delegate Control... . Click Next, then select the user that you are using for managerdn and click Next. The required permission is on the next screen "Read all inetOrgPerson information."</p> <p>Example:</p> <p>CN=Manager,CN=Users,DC=mydomaincontroller,DC=com</p> <p>In this example, the Active Directory Bind user is Manager, in the Users tree of the mydomaincontroller.com domain. Usually managerdn is a user who has fewer privileges than Administrator, but has sufficient privileges to query Active Directory for users attributes and group information.</p> <p>⚠ WARNING</p> <p>You must update this value in ECS if the managerdn credentials change in Active Directory.</p>
Manager Password	<p>The password of the managerdn user.</p> <p>⚠ WARNING</p> <p>You must update this value in ECS if the managerdn credentials change in Active Directory.</p>
Providers	<p>Select Disabled if you want to add the server to ECS but not immediately use it for authentication. (Regardless of whether this property is true, ECS validates that the provider's name and domain are unique.)</p>
Group Attribute	<p>Indicates the Active Directory attribute that is used to identify a group. Used for searching the directory by groups.</p> <p>Example: CN</p> <p>Active Directory only. Does not apply to other authentication providers.</p> <hr/> <p>Note</p> <p>Once this value is set for a provider, it cannot be changed, because of the tenants that are using this provider may already have role assignments and permissions configured using group names in a format using the current attribute.</p>

Table 4 Authentication provider settings (continued)

Field name	Description and requirements
Group Whitelist	<p>Optional. One or more group names as defined by the authentication provider. This setting will filter the group membership information that ECS retrieves about a user.</p> <ul style="list-style-type: none"> When a group or groups are included in the whitelist, it means that ECS will be aware of a user's membership in the specified group[s] only. Multiple values (one per line in ECS portal, comma-separated in CLI and API) and wildcards (for example MyGroup*,TopAdminUsers*) are allowed. Blank value (default) means that ECS will be aware of any and all groups that a user belongs to. Asterisk (*) is the same as blank. <p>Example:</p> <p>UserA belongs to Group1 and Group2.</p> <p>If the whitelist is blank, ECS knows that UserA is a member of Group1 and Group2.</p> <p>If the whitelist is "Group1", ECS knows that UserA is a member of Group1, but does not know that UserA is a member of Group2 (or of any other group).</p> <p>Use care when adding a whitelist value. For example, if mapping a user to a tenant is based on group membership, then ECS must be aware of the user's membership in the group.</p> <p>To restrict access to a namespace to users of certain group(s) only, one must:</p> <ul style="list-style-type: none"> add these group(s) to the namesapce user mapping , so the tenant is configured to accept only users of these group(s). add these group(s) to the whitelist, so that ECS is authorized to receive information about them <p>Note that by default, if no groups are added to the tenant user mapping, users from any groups are accepted, regardless of the whitelist configuration.</p> <p>Active Directory only. Does not apply to other authentication providers.</p>
Search Scope	<p>One Level (search for users one level under the search base) or Subtree (search the entire subtree under the search base).</p>
Search Base	<p>Indicates the Base Distinguished Name that ECS uses to search for users at login time and when assigning roles or setting ACLs.</p> <p>Example: CN=Users,DC=mydomaincontroller,DC=com</p> <p>This example searches for all users in the Users container.</p> <p>Example: CN=Users,OU=myGroup,DC=mydomaincontroller,DC=com</p>

Table 4 Authentication provider settings (continued)

Field name	Description and requirements
	<p>This example searches for all users in the Users container in the myGroup organization unit.</p> <p>Note that the structure of the searchbase value begins with the "leaf" level and goes up to the domain controller level--the reverse of the structure seen in the Active Directory Users and Computers UI.</p>
Search Filter	<p>Indicates the string used to select subsets of users. Example: userPrincipalName=%u</p> <hr/> <p>Note</p> <p>ECS does not validate this value when you add the authentication provider.</p> <hr/> <p>If an alternate UPN suffix is configured in the Active Directory, the Search Filter value must be of the format sAMAccountName=%U where %U is the username, and does not contain the domain name.</p>

Considerations when adding authentication providers

When you configure ECS to work with Active Directory, you must decide whether to manage several domains in a single authentication provider, or to add separate authentication providers for each domain.

The decision to add a single authentication provider, or multiple, depends on the number of domains in the environment, and the location on the tree from which the manager user is able to search. Authentication providers have a single search_base from which searches are conducted. They have a single manager account who must have read access at the search_base level and below.

Use a single authentication provider for multiple domains if you are managing an Active Directory forest and:

- the manager account has privileges to search high enough in the tree to access all user entries
- the search will be conducted throughout the whole forest from a single search base, not just the domains listed in the provider.

Otherwise, configure an authentication provider for each domain.

Note that even if you are dealing with a forest and you have the correct privileges, you might not want to manage all the domains with a single authentication provider. You would still use one authentication provider per domain when you need granularity and tight control on each domain, especially to set the search base starting point for the search. Since there is only one search base per configuration, it needs to include everything that is scoped in the configuration in order for the search to work.

The search base needs to be high enough in the directory structure of the forest for the search to correctly find all the users in the targeted domains.

- If the forest in the configuration contains ten domains but you target only three, do not use a single provider configuration, because the search will unnecessarily span

the whole forest, and this may adversely affect performance. In this case, use three individual configurations.

- If the forest in the configuration contains ten domains and you want to target ten domains, a global configuration is a good choice, because there is less overhead to set up.

Understanding the mapping of users into a namespace

Domain users can be added to ECS using authentication providers. To make users available as namespace users they need to be mapped into the namespace.


The authentication provider makes users belonging to specified domains and whitelisted groups available to ECS and they can be assigned to system roles.

To associate users with a namespace and make them eligible to be object users for the namespace, you must associate the domain to which the users belong with the namespace and, if necessary, apply finer grained filtering based on the groups that belong to the domain and the attributes that have been assigned to the domain users. A domain can be mapped to a single namespace or can provide users for multiple namespaces.

The ECS Portal and the ECS Management REST API provide the ability to specify mappings when a new namespace is registered and provide support for updating the mappings for all namespaces. Creating a namespace is an operation that requires System Admin privileges; modifying a tenant and performing user mappings operations can be performed by a Namespace Admin.


The user mappings assigned to different namespaces must not overlap, so if the Accounts namespace maps users from the same domain as the HR namespace, it must provide additional mappings to differentiate its users. In the example below, the Accounts namespace uses the corp.sean.com domain but maps users with specific attributes, in this case, those with their Department attribute set to Accounts in Active Directory.

Figure 8 User mappings for a tenant using AD attributes



The example below shows the use of multiple mapping criteria. All members of the corp.sean.com domain who belong to the Storage Admins group and have their Department attribute set to Accounts AND Company set to Acme, OR belong to the Storage Admins group and have their Department set to Finance, will be mapped into the namespace.

Figure 9 Using multiple mapping criteria



Map domain users into a namespace

The ECS portal provides the ability to map users into a namespace based on the AD/LDAP domain, groups, and attributes associated with users.

Before you begin

- An authentication provider must have been registered with ECS and must provide access to the domain from which you want to map users.
- The administrator of the AD must have configured the groups or users in AD before mapping the users from the ECS Portal.
- If you are using attribute mapping, each user must have the appropriate attribute value set in AD.

You should understand the concepts associated with user mapping, described in [Understanding the mapping of users into a namespace on page 70](#).

Procedure

1. At the ECS portal, select **Manage** › **Namespace**.
2. In the Namespaces table, click on the **Edit** action for the namespace to open it for editing.
3. If a domain hasn't already been specified, click **Add** to add a mapping and enter the domain name in the Domain field.
4. Specify any groups that you want to use to map users into the namespace.
The group or groups that you specify must exist in AD.
5. If you want to use attributes to map users into the namespace enter the name of the attribute and the value or values for the attribute. If you do not want to use attributes to map users into the namespace, click the delete button to remove the attribute fields from the current mapping.
For users to be mapped into the domain, the attribute value set for the user must match the attribute value specified in ECS.
6. **Save** the namespace settings.

CHAPTER 7

Manage tenants

- [Introduction](#)..... 74
- [Quotas](#)..... 74
- [Retention periods and policies](#)..... 75
- [Lock buckets and users](#)..... 76
- [Metering](#)..... 77
- [Audit buckets](#)..... 78

Introduction

ECS provides a number of features to support the management of a tenant.

The following features are supported:

Users

The ability to assign a Namespace Admin for the namespace and to create object users for the namespace is described in [Add users and assign roles on page 54](#).

Quotas

The ability to set quotas on namespaces and buckets is described in [Quotas on page 74](#).

Retention Periods

The ability to create retention policies is described in [Retention periods and policies on page 75](#).

Lock buckets and users

The ability to lock buckets and users is described in [Lock buckets and users on page 76](#).

Metering

The ability to meter the writing of data to buckets and namespaces is described in [Metering on page 77](#).

Audit buckets

The ability to audit the operations associated with buckets is described in [Audit buckets on page 78](#).

Quotas

You can set soft and hard quotas on a namespace and on buckets created within a namespace.

Soft quotas cause events to be logged to inform you that the quota has been reached; hard quotas provide a hard limit on the amount of object storage that can be used for a bucket or namespace - when the limit is reached, access to the bucket or namespace is blocked.

Quotas can be set from the ECS Portal or using the API and the CLI.

Setting quotas from the portal

You can set quotas for a namespace from the **Manage** > **Namespace** page, as described in [Configure a namespace for a tenant on page 46](#).

Quotas for a bucket are set from the **Manage** > **Bucket** page, as described in [Create and configure buckets on page 86](#).

Setting quotas using the API

The following API paths provide the ability to set quotas:

Method	Description
PUT/GET/DELETE /object/namespaces/namespace/{namespace}/quota	Sets the quota for a namespace. The payload specifies hard and soft quotas.

Method	Description
PUT/GET/DELETE /object/bucket/{bucketName}/quota	Sets the quota for a bucket. The payload specifies hard and soft quotas.

You can find more information about the ECS Management REST API in: [Data Access Guide: Use the ECS Management REST API](#) and the online reference is [here](#).

Retention periods and policies

ECS provides the ability to prevent data being modified or deleted within a specified retention period.

Retention periods and retention policies can be defined in metadata associated with objects and on buckets, and is checked each time a request to modify an object is made. Retention periods are supported on all object interfaces S3, Swift, Atmos, and CAS. However, CAS data is immutable so the retention period when applied to CAS refers to the ability to delete CAS objects only.

There are two ways of defining retention: retention periods and retention policies.

Retention Periods

Retention periods are assigned at the object and/or bucket level. Each time an attempt is made to modify or delete an object, an expiration time is calculated, where object expiration time = object creation time + retention period. Where a retention period is assigned on a bucket, the retention period for the bucket is checked and the expiration time calculated based on the retention period set on the object and the value set on the bucket, whichever is the longest.

Applying a retention period to a bucket means that the retention period for all objects in a bucket can be changed at any time, and can override the value written to the object by an object client by setting it to a longer period.

It is possible to specify that an object is retained indefinitely.

Retention Policies

Retention policies enable retention use cases to be captured and applied to objects. Retention policies are associated with a namespace and any policy associated with the namespace can be assigned to an object belonging to the namespace. A retention policy has an associated retention period.

The use of retention policies provides the flexibility to change the period associated with a policy and, in doing so, automatically change the retention period that applies to any objects that have that policy assigned.

Where a retention policy is applied to an object, when an attempt to modify or delete an object is made, the retention period associated with the policy is retrieved and used in conjunction with object and bucket retention periods to determine if the request is allowed.

As an example, a named policy could be defined for each of the following types of document and each named policy can have an appropriate retention period:

- Financial - 3 years
- Legal - 5 years
- Email - 6 months

How to create retention policies

You can configure the retention policies that are available for the namespace from the ECS Portal, refer to:

[Configure a namespace for a tenant on page 46](#)

or you can create them using the ECS Management REST API, a summary of which is provided below.

Method	Description
PUT /object/bucket/{bucketName}/retention	The retention value for a bucket defines a mandatory retention period which is applied to every object within a bucket. So, if you set a retention period of 1 year, an object from the bucket can not be modified or deleted for one year.
GET /object/bucket/{bucketName}/retention	Returns the retention period that is currently set for a specified bucket.
POST /object/namespaces/namespace/{namespace}/retention	For namespaces, the retention setting acts like a policy, where each policy is a <Name>: <Retention period> pair. You can define a number of retention policies for a namespace and you can assign a policy, by name, to an object within the namespace. This allows you to change the retention period of a set of objects that have the same policy assigned by changing the corresponding policy.
PUT /object/namespaces/namespace/{namespace}/retention/{class}	Updates the period for a retention class that is associated with a namespace.
GET /object/namespaces/namespace/{namespace}/retention	Returns the retention classes defined for a namespace.

You can find out how to access the ECS Management REST API in the following article: [Data Access Guide: Use the ECS Management REST API](#) and the online reference is [here](#).

How to apply retention policies and periods

You can apply retention periods to buckets at the ECS Portal.

When you create objects or buckets using the object service protocols, for example, when you create an S3 bucket using a client that supports the S3 protocol, you can apply the retention period or retention policy using x-ems headers.

When you create objects, you can apply the following retention period and retention policy headers:

- x-ems-retention-period
- x-ems-retention-policy

When you create a bucket, you can set the retention period using the x-ems-retention-period header.

Lock buckets and users

ECS provides the ability to prevent access to a bucket and to prevent user access.

Support for the bucket and user lock operations is provided by the ECS Management REST API. There is no support for locking buckets and users in the ECS Portal . The following calls are supported:

Method	Description
PUT /object/bucket/{bucketName}/lock	Locks a bucket so that all writes to the bucket are disallowed.
DELETE /object/bucket/{bucketName}/lock	Unlocks a bucket so that writes to the bucket are re-enabled.
PUT /object/users/{userid}/lock	Locks an object user (not AD user) such that all subsequent API operations performed by the user return an error.
DELETE /object/user/{userid}/lock	Unlocks an object user (not AD user) such that the user is re-enabled to perform further API operations.

You can find out how to access the ECS Management REST API in the following article: [Data Access Guide: Use the ECS Management REST API](#) and the online reference is [here](#).

Metering

ECS provides support for metering the use of the object storage at the namespace and bucket level.

Metering using the portal

You can use the ECS Portal to monitor the use of namespace and buckets. The **Monitor** > **Metering** page enables a namespace or a specific bucket from a namespace to be selected and its metering data displayed.

Table 5 Bucket and namespace metering

Attribute	Description
Total Size (GB)	Total size of the objects stored in the selected namespace or bucket at the end time specified in the filter.
Object Count	Number of objects associated with the selected namespace or bucket at the end time specified in the filter.
Objects Created	Number of objects created in the selected namespace or bucket in the time period.
Objects Deleted	Number of objects deleted from the selected namespace or bucket in the time period.
Bandwidth Ingress (MB)	Total of incoming object data (writes) for the selected namespace or bucket during the specified period.
Bandwidth Egress (MB)	Total of outgoing object data (reads) for the selected namespace or bucket during the specified period.

Note

Metering data is not available immediately as it can take a significant amount of time to gather the statistics for data added to the system and deleted from the system.

Refer to [Monitor storage: metering and capacity](#) for more information on accessing these details.

Metering using the API

The following API paths provide the ability to retrieve metering information:

Method	Description
GET /object/billing/buckets/{namespace}/{bucket}/info?sizeunit=<KB MB GB>	Gets the current usage for a bucket in a specified namespace.
GET /object/billing//buckets/{namespace}/{bucket}/sample? start_time=<ISO8061_format>&end_time=<ISO8061_format>&marker=<string_marker> &sizeunit=<KB MB GB>	Samples a bucket activity for the given time slice. By default, a bucket's minimum sample resolution is 5 minutes and samples will be retained for 30 days. If the start time and end time do not fall on sample boundaries, an error will be returned. If the time range spans multiple low-level samples, the data will be aggregated for the time period to produce one data point.
GET /object/billing/namespace/{namespace_name}/info? marker=<string_marker>&include_bucket_detail=<true false>&sizeunit=<KB MB GB>	Gets usage information for all of the buckets in a namespace. Note When bucket details are included, the total size on disk might be different to the total size without bucket details. This is due to bucket size being rounded and summed to give the total size.
GET /object/billing/namespace/{namespace_name}/sample? start_time=<ISO8061_format>&end_time=<ISO8061_format>&marker=<string_marker> &include_bucket_detail=<true false>&sizeunit=<KB MB GB>	Gets a snapshot for a particular time sample for a namespace. By default, buckets and namespaces will be sampled every 5 minutes and samples will be retained for 30 days. If the start time and end time do not fall on sample boundaries an error will be returned. If the time range spans multiple low-level samples, the data will be aggregated for the time period to produce one data point.

You can find more information about the ECS Management REST API in: [Data Access Guide: Use the ECS Management REST API](#) and the online reference is [here](#).

Audit buckets

The controller API provides the ability to audit the use of the S3, EMC Atmos, and OpenStack Swift object interfaces.

The following operations on object containers (S3 buckets, EMC Atmos subtenants, and OpenStack Swift containers) are logged.

- Create Bucket
- Delete Bucket
- Update Bucket

- Set Bucket ACL
- Change Bucket Owner
- Set Bucket Versioning
- Set Bucket Versioning Source
- Set Bucket Metadata
- Set Bucket Head Metadata
- Set Bucket Expiration Policy
- Delete Bucket Expiration Policy
- Set Bucket Cors Configuration
- Delete Bucket Cors Configuration

Audit logging at the portal

You can use the Portal **Monitor** > **Events** page to detect the generation of an audit log event.

The root user should only be used for initial access to the system. On initial access, the root user password should be changed at the **Settings** > **Password** page and one or more new System Admin accounts should be created. From an audit perspective, it is important to know which user carried out changes to the system, so root should not be used, and each System Admin user should have their own account.

You can refer to [Monitor events: audit portal, API, and CLI events and system alerts on page 124](#) for more information on using the events log.

Audit API

Support for bucket auditing is provided by the following ECS Management REST API calls:

Method	Description
GET /monitoring/events	Retrieves the audit events for a specified namespace and time interval.

You can find more information about the ECS Management REST API in: [Data Access Guide: Use the ECS Management REST API](#) and the online reference is [here](#).

CHAPTER 8

Remove a site

- [Fail over a site/Delete a VDC.....](#)82

Fail over a site/Delete a VDC

Use this procedure to delete a VDC. Deleting a VDC initiates site fail over when the VDC you are deleting is part of a multi-site federation.

If a disaster occurs, an entire VDC can become unrecoverable. ECS initially treats the unrecoverable VDC as a temporary site failure. If the failure is permanent, you must remove the VDC from the federation to initiate fail over processing which reconstructs and reprotects the objects stored on the failed VDC. The recovery tasks run as a background process. Review the recovery process by using the **Monitor > Geo Replication > Failover Processing**.

Procedure

1. Log in to one of the operational VDCs in the federation.
2. Go to **Manage > Replication Group**.
3. Click **Edit** for the replication group that contains the VDC to delete.
4. Click **Delete** in the row that contains the VDC and storage pool to remove.
5. Click **Save**.
6. Go to **Manage > VDC**. The status for the permanently removed VDC changes to `Permanently failed`.
7. Select **Delete** from the drop down in the row of the VDC to remove.
8. Click **Save**.

CHAPTER 9

Manage licenses

- [Licensing](#).....84
- [Obtain the EMC ECS license file](#)..... 84

Licensing

EMC ECS licensing is capacity-based.

At a minimum you need to obtain at least an ECS license and upload it to the appliance.

The **Settings > License** page provides additional details.

Obtain the EMC ECS license file

You need to obtain the license file (.lic) from the EMC license management web site for uploading to ECS.

Before you begin

In order to obtain the license file, you must have the License Authorization Code (LAC), which was emailed from EMC.

Procedure

1. Go to support.EMC.com
2. Select **Support > Service Center**.
3. Select **Get and Manage Licenses**.
4. Select **ECS** from the list of products.
5. On the LAC Request page, enter the LAC code and **Activate**.
6. Select the entitlements to activate and **Start Activation Process**.
7. Select **Add a Machine** to specify any meaningful string for grouping licenses.

The "machine name" does not have to be a machine name at all; enter any string that will help you keep track of your licenses.

8. Enter the quantities for each entitlement to be activated, or select **Activate All**. Click **Next**.

If you are obtaining licenses for a multisite (geo) configuration, you should distribute the controllers as appropriate in order to obtain individual license files for each virtual data center.

9. Optionally specify an addressee to receive an email summary of the activation transaction.
10. Click **Finish**.
11. Click **Save to File** to save the license file (.lic) to a folder on your computer.

This is the license file that is needed during initial setup of ECS, or when adding a new license later in the ECS Portal (**Settings > Licensing**).

12. From the ECS Portal, select **Settings > Licensing**.
13. Select **Choose File**, select your license file, and select **Upload**.

The license display updates.

CHAPTER 10

Create and manage buckets

- [Introduction](#) 86
- [Bucket concepts](#)86
- [Bucket attributes](#)87
- [Bucket ACLs](#) 92
- [Create a bucket using the ECS Portal](#) 93
- [Edit a bucket](#) 94
- [Set the bucket ACL permissions for a user](#) 95
- [Set the bucket ACL permissions for a pre-defined group](#) 96
- [Set custom group bucket ACLs](#) 97
- [Create bucket using the object APIs](#)98
- [Bucket and key naming conventions](#) 102

Introduction

Containers are required to store object data. In S3 these containers are called *buckets* and this term has been adopted as a general term in ECS. In Atmos, the equivalent of a bucket is a *subtenant*, in Swift, the equivalent of a bucket is a *container*, and for CAS, a bucket is a *CAS pool*.

In ECS, buckets are assigned a type which can be S3, Swift, Atmos, or CAS. In addition, S3, Atmos, or Swift buckets can be configured to support file system access (for HDFS), and a bucket configured for file system access can be read and written using its object protocol and using the HDFS protocol. This is often referred to as *cross-head support*.

Buckets can be created for each object protocol using its API, usually using a client that supports the appropriate protocol. Additional support for creating S3, HDFS, and CAS buckets is provided by the ECS Portal and the ECS Management API. The ability to create buckets from the portal makes it easy to create buckets for HDFS and CAS and makes it easy to take advantage of some of the more advanced bucket configuration options provided by ECS, such as quotas and retention periods.

Where you want to create buckets using the object protocols, you can use special x-emc headers to control bucket configuration.

This article describes how to create and edit buckets, and set ACLs for a bucket, using the ECS Portal and also describes the additional x-emc headers that you can use to control bucket configuration when using the supported object protocols.

Bucket concepts

Buckets are object containers and can be used to control access to objects and to set properties that define attributes for all contained objects, such as retention periods and quotas.

Bucket access

Buckets are associated with a replication group. Where the replication group spans multiple VDCs, the bucket contents are similarly replicated across the VDCs. Objects in a bucket that belongs to a replication group which spans two VDCs, VDC1 and VDC2, for example, can be accessed from either VDC1 or VDC2. Objects in a bucket that belongs to a replication group that is only associated with VDC1, can only be accessed from VDC1, they cannot be accessed from other VDCs in a federated ECS system.

The identity of a bucket and its metadata, such as its ACL, are global management information in ECS, which means that they are replicated across the system storage pools and can be seen from all VDCs in the federation. However, the bucket can only be listed from a VDC that is part of the replication group to which the bucket belongs.

Bucket ownership

A bucket belongs to a namespace and object users are also assigned to a namespace. Each object user can create buckets only in the namespace to which they belong, however, any ECS object user can be assigned as the owner of a bucket or object, or a grantee in a bucket ACL, even if the user does not belong to the same namespace as the bucket or object. This enables buckets and objects to be shared between users in different namespaces. For example, in an enterprise where a namespace is a department, a bucket or object can be shared between users in different departments.

When an object user wants to access a bucket in a namespace that they don't belong to, the namespace must be specified using the x-emc-namespace header.

Access to a bucket during temporary site outage

ECS provides a temporary site outage mechanism that enables objects to be retrieved even if the primary copy of the object is not available due to the site that hosts the primary being unavailable.

Because there is a risk that object data retrieved during a temporary site outage is not the most recent, the user must indicate that they are prepared to accept this by marking the bucket as available during an outage.

During the outage object data is accessible for both read and write; buckets enabled for file system access are available read-only; CAS data is available read-only as a result of the fact that it is immutable, not as a result of Access During Outage operational mode.

Bucket attributes

The ECS Portal enables buckets to be created and managed at the **Manage > Buckets** page.

The Bucket Management page provides a bucket table which displays the buckets for a selected namespace. The table displays bucket attributes and provides **Edit Bucket**, **Edit ACL**, and **Delete** actions for each bucket.

Name	Replication Group	Owner	Notification Quota (GB)	Max Quota (GB)	Encryption	Metadata	Created	Actions
sanity-bucket1-nh	sanity-rg1	wuser1@sanity.local	0	0	Disabled	Disabled	2015-11-20	Edit bucket
sanity-bucket2-nh	sanity-rg1	wuser1@sanity.local	0	0	Disabled	Disabled	2015-11-20	Edit bucket Edit ACL Delete
sanity-bucketfs-nh	sanity-rg1	wuser1@sanity.local	0	0	Disabled	Disabled	2015-11-20	Edit bucket

The attributes associated with a bucket are described in the following table. To view and change attributes that are not displayed on the Bucket Management page, you can select **Edit Bucket**.

Table 6 Bucket attributes

Attribute	Description	Can be Edited
Name	Name of the bucket. You can refer to the following topic for guidance on bucket naming: Bucket and key naming conventions on page 102 .	No
Namespace	Namespace with which the bucket is associated.	No

Table 6 Bucket attributes (continued)

Attribute	Description	Can be Edited
Replication Group	Replication group in which the bucket will be created.	No
Bucket Owner	Bucket owner.	Yes
Bucket Tagging	Tags are name-value pairs that can be defined for a bucket and enable buckets to be classified. More information on bucket tagging is provided in: Bucket tagging on page 91 .	Yes
Quota	<p>Quota for a bucket. Behavior associated with exceeding the quota can be defined by setting Hard (Block) and Soft (Notification) and quotas.</p> <p>Soft (Notification) Quota</p> <p>Quota setting at which you will be notified. This is a soft quota and can be set on its own or can be set in addition to a hard quota.</p> <p>The quota cannot be set less than 1GB.</p> <p>More information on quotas is provided in: Manage a tenant on page 74.</p> <p>Hard (Block) Quota</p> <p>Hard quota which, when reached, will cause writes/updates to the bucket to be blocked. A soft quota can be set to trigger before the hard quota is reached.</p>	Yes
Server-side Encryption	<p>Indicates whether Server-side encryption is enabled. Server-side Encryption is also known as Data At Rest Encryption or D@RE. This feature encrypts data inline before storing it on ECS disks or drives. This encryption prevents sensitive data from being acquired from discarded or stolen media. If encryption is enabled when the bucket is created, then the feature cannot be disabled later.</p> <p>If the bucket's namespace is encrypted, then every bucket will be encrypted. If the namespace is not encrypted, then you have the choice of encrypting individual buckets.</p> <p>For a complete description of the feature, see the <i>ECS Security Configuration Guide</i>.</p>	No
File System	<p>Indicates that ECS will allow the bucket to be used as a Hadoop Distributed File System (HDFS).</p> <p>To simplify access to the file system, a default group, and default permissions associated with the group, can be defined. More information can be found in Default Group on page 89</p>	No
CAS	Indicates that the bucket is enabled for CAS data.	
Metadata Search	Indicates that metadata search indexes will be created for the bucket based on specified key values.	No

Table 6 Bucket attributes (continued)


Attribute	Description	Can be Edited
	<p>If Enabled, metadata keys that will be used as the basis for indexing objects in the bucket can be defined. These keys must be specified at bucket create time.</p> <p>Once the bucket is created, search can be disabled altogether, but the configured index keys cannot be modified</p> <p>The way the attribute is defined is described in Metadata index keys on page 90.</p> <hr/> <p>Note</p> <p>Metadata that is encrypted cannot be indexed for search. Hence users cannot enable metadata search on a bucket if Server-side Encryption (D@RE) is enabled.</p>	
Access During Outage	<p>A flag set on the bucket which specifies the behavior when accessing data in the bucket when there is a temporarily unavailable zone in a geo-federated setup.</p> <p>If you set this flag to Enabled, and a temporary site outage occurs, objects that you access in this bucket might have been updated at the failed site but changes might not have been propagated to the site from which you are accessing the object. Hence, you are prepared to accept that the objects you read might not be up to date.</p> <p>If the flag is Disabled, data in the zone which has the temporary outage is not available for access from other zones and object reads for data which has its primary in the failed site will fail.</p>	Yes
Bucket Retention	<p>Sets the retention period for a bucket.</p> <p>The expiration of a retention period on an object within a bucket is calculated when a request to modify an object is made and is based on the value set on the bucket and the objects themselves.</p> <p>The retention period can be changed during the lifetime of the bucket.</p> <p>Information on retention period is provided in: Retention periods and policies on page 75.</p>	Yes

Default Group

Where a bucket is enabled for file system access, it is possible to specify a default group for the bucket. When accessed as a file system, the members of the Unix group can access the file system. Without this assignment, only the bucket owner would be able to access the file system.

In addition, files and directories created using object protocols can be assigned group permissions that will enable members of the Unix group to access them.

The File System Enabled dialog is shown below.




Metadata index keys

When Metadata Search is enabled, a set of system and/or user metadata fields/attributes can be specified as search keys for objects in a bucket. For each specified metadata search key, ECS will create an index of objects that have corresponding metadata based on the value for the metadata search key.

The metadata search facility allows S3 object clients to search for objects in a bucket based on the indexed metadata using a rich query language.

The Add Metadata Search Key dialog enables the Metadata Search Key to be selected as either System or User. For System, metadata that is automatically assigned to objects in a bucket is listed in the Key Names menu.



When a Metadata Key Type of User is selected (see below), you must specify the name of the user metadata to create an index for. In addition, you need to specify the data type so that ECS knows how to interpret the metadata values provided in search queries.

You can read more about metadata search feature in [Data Access Guide: Metadata search S3 extension](#).

Bucket tagging

Tags in the form of name-value pairs can be assigned to a bucket enabling object data stored in the bucket to be categorized. For example, bucket data can be associated with a cost-center or project.

Bucket tags and values can be read and managed using the ECS Portal or using custom clients with the ECS Management REST API. In addition, bucket tags are included in the metering data reports in the ECS Portal or ECS Management REST API.

The bucket tagging dialog is shown below.

Bucket ACLs

The privileges a user has when accessing a bucket are set using an Access Control List (ACL).

When you create a bucket and assign an owner to it, an ACL is created that assigns a default set of permissions to the bucket owner - the owner is, by default, assigned full control.

You can modify the permissions assigned to the owner or you can add new permissions for a user by selecting the Edit ACL operation for the bucket.

At the ECS Portal, the Bucket ACLs Management page provides **User ACLs**, **Group ACLs**, and **Custom Group ACLs** panels to manage the ACLs associated with individual users and pre-defined groups, and to allow groups to be defined that can be used when accessing the bucket as a file system.


The ACL attributes are provided in the following table.

Table 7 Bucket ACLs

ACL	Permission
Read	Allows user to list the objects in the bucket.
Read ACL	Allows user to read the bucket ACL.
Write	Allows user to create or update any object in the bucket.
Write ACL	Allows user to write the ACL for the bucket.
Execute	Sets the execute permission when accessed as a file system. This permission has no effect when the object is accessed using the ECS object protocols.
Full Control	Allows user to Read, Write, Read ACL, and Write ACL.
Privileged Write	Allows user to perform writes to a bucket or object when the user doesn't have normal write permission. Required for CAS buckets.
Delete	Allows user to delete buckets and objects. Required for CAS buckets.
None	User has no privileges on the bucket.

User ACLs

The User ACL panel show the ACLs that have been applied to users and enables ACLs to be assigned to a user using the **Add** operation.



Note

Because the ECS Portal supports S3, HDFS, and CAS buckets, the range of permissions that can be set are not applicable to all bucket types.

Group ACLs

You can set permissions for a set of pre-defined groups. The following groups are supported:

public

All users authenticated or not.

all users

All authenticated users.

other

Authenticated users but not the bucket owner.

log delivery

Not supported.

The permissions that can be assigned are listed in [Table 7 on page 92](#).

Custom Group ACLs

Custom group ACLs enable groups to be defined and for permissions to be assigned to the group. The main use case for assigning groups to a bucket is to support access to the bucket as a file system, for example, when making the bucket available for HDFS.

Create a bucket using the ECS Portal

The ECS portal enables the creation of buckets and provides the ability to specify the configuration of the bucket. Buckets created at the portal can be either S3, S3+HDFS, or CAS buckets.

Before you begin

- You must be a Namespace Admin or a System Admin to create a bucket at the ECS portal.
- If you are a Namespace Admin you can create buckets in your namespace.
- If you are System Admin you can create a bucket belonging to any namespace.

Procedure

1. At the ECS Portal, select **Manage** > **Buckets**.

2. Select **New Bucket**.

3. Select the namespace that the bucket and its objects will belong to.

If you are a System Admin and the ECS system has more than one namespace, select the namespace to which the bucket will belong.

If you are a Namespace Admin, you will only be able to select your own namespace.

4. Select the replication group that the bucket will be associated with.

5. Specify a bucket owner.

The bucket owner should be an ECS object user for the namespace. If you don't specify a user, you will be assigned as the owner, however, you will not be able to access the bucket unless your username is also assigned as an object user.

The user that you specify will be given Full Control.

6. Add any tags to the bucket by clicking Add at the Bucket Tagging control and add name-value pairs.

You can read more about Bucket Tagging in [Bucket tagging on page 91](#).

7. If required, specify a quota for the bucket.

The settings that you can apply are described in: [Quotas on page 74](#).

8. If you want data in the bucket to be encrypted, set Server-side Encryption to Enabled.
9. If you want the bucket to be a CAS bucket, set the CAS control to Enabled.

By default, CAS will be disabled and the bucket will be marked as an S3 bucket.

10. If you want the bucket to support operation as a file system (for HDFS), set the File System Enabled control to Enabled.

The bucket will be an S3 bucket that supports HDFS.

You can set a default Unix group for access to the bucket and for objects created in the bucket. More details are provided in: [Default Group on page 89](#)

11. If you want the bucket to support searches based on object metadata, you should set the Metadata Search control to Enabled.

If you enable Metadata Search you can add User and System metadata keys that will be used to create object indexes. More information on entering metadata search keys is provided in [Metadata index keys on page 90](#).

Note

If the bucket is to be used for CAS, you cannot enable metadata search as a similar search capability is provided in the implementation of the Centera API.

12. Set Access During Outage as Enabled if you want the bucket to be available during a temporary site outage.

13. If required, set a bucket retention period for the bucket.

You can read more about retention periods in: [Retention periods and policies on page 75](#).

14. Select **Save** to create the bucket.

Results

You can assign users to the bucket and set permissions for users (or pre-defined groups) from the buckets table Actions menu.

Edit a bucket

You can edit some bucket settings after the bucket has been created and after it has had objects written to it.

Before you begin

- You must be a Namespace Admin or a System Admin to edit a bucket.
- If you are a Namespace Admin you can edit the setting for buckets belonging to your namespace.
- If you are System Admin you can edit the settings for a bucket belonging to any namespace.

Procedure

1. At the ECS portal, select **Manage** › **Buckets**.
2. In the Buckets table, select the **Edit** action for the bucket for which you want to change the settings.
3. You can edit the following bucket attributes:
 - Quota
 - Bucket Owner
 - Bucket Tagging
 - Access During Outage
 - Bucket Retention

You cannot change the following attributes of the bucket:

- Replication Group
- Server-side Encryption
- File System Enabled
- CAS Enabled
- Metadata Search

You can find out more information about these settings in: [Bucket concepts on page 86](#).

4. Select **Save**.

Set the bucket ACL permissions for a user

The ECS portal enables the ACL for a bucket to be set for a user or for a pre-defined group.

Before you begin

- You must be a Namespace Admin or a System Admin to edit the ACL for a bucket.
- If you are a Namespace Admin you can edit the ACL settings for buckets belonging to your namespace.
- If you are System Admin you can edit the ACL settings for a bucket belonging to any namespace.

Procedure

1. At the ECS Portal, select **Manage** › **Buckets**.
2. In the Buckets table, select the **Edit ACL** action for the bucket for which you want to change the settings.
3. To set the ACL permissions for a user, select the **User ACLs** button.

To select the ACL for a group, select **Group ACLs** or **Custom Group ACLs**. You can refer to [Set the bucket ACL permissions for a pre-defined group on page 96](#) or [Set custom group bucket ACLs on page 97](#) for more information on setting group ACLs.
4. You can edit the permissions for a user that already has permissions assigned, or you can add a user that you want to assign permissions for.
 - To set (or remove) the ACL permissions for a user that already has permissions, select Edit (or Remove) from the Action column in the ACL table.

- To add a user to which you want to assign permissions, select **Add**.

The user that you have set as the bucket owner will have already have default permissions assigned.

5. If you have added an ACL, enter the username of the user that the permissions will apply to.
6. Specify the permissions that you want to apply to the user.
More information on ACL privileges is provided in [Bucket concepts on page 86](#).
7. Select **Save**.

Set the bucket ACL permissions for a pre-defined group

The ECS Portal enables the ACL for a bucket to be set for a pre-defined group.

Before you begin

- You must be a Namespace Admin or a System Admin to edit the group ACL for a bucket.
- If you are a Namespace Admin you can edit the group ACL settings for buckets belonging to your namespace.
- If you are System Admin you can edit the group ACL settings for a bucket belonging to any namespace.

Procedure

1. At the ECS portal, select **Manage** › **Buckets**.
2. In the Buckets table, select the **Edit ACL** action for the bucket for which you want to change the settings.
3. To set the ACL permissions for a pre-defined group, select the **Group ACLs** button.
You can read more about the pre-defined groups in: [Bucket concepts on page 86](#)

4. Select the privileges that you want to assign to the group.
5. Select **Save**.

Set custom group bucket ACLs

The ECS Portal enables the group ACL for a bucket to be set. Bucket ACLs can be granted for a group of users (Custom Group ACL) or for individual users, or a combination of both. For example, you can grant full bucket access to a group of users, but you can also restrict (or even deny) bucket access to individual users in that group.

Before you begin

- You must be a Namespace Admin or a System Admin to edit the group ACL for a bucket.
- If you are a Namespace Admin you can edit the group ACL settings for buckets belonging to your namespace.
- If you are System Admin you can edit the group ACL settings for a bucket belonging to any namespace.

When the bucket is accessed using HDFS, using ECS multi-protocol access, members of the Unix group will be able to access the bucket.

Procedure

1. At the ECS Portal, select **Manage** > **Buckets**.
2. In the Buckets table, select the **Edit ACL** action for the bucket for which you want to change the settings.

- To set the ACL for a custom group, select **Custom Group User ACLs**.
- At the **Custom Group User ACLs** page, select **Add**.

The screenshot shows the 'Edit Custom Group' interface. At the top, there are three tabs: 'User ACLs', 'Group ACLs', and 'Custom Group ACLs'. Below the tabs is a breadcrumb: 'Bucket Management / Bucket new-bkt / Edit Custom Group'. The main form area contains a 'Custom Group Name *' label above a text input field. Below that is a 'Permissions *' section with a help icon. The permissions are listed in two columns: Read (checked), Write (checked), Execute (checked), Full Control (unchecked) on the left; and Read ACL (unchecked), Write ACL (unchecked), Delete (unchecked), None (unchecked) on the right. A 'Privileged Write' checkbox is also present and unchecked. At the bottom of the form are 'Save' and 'Cancel' buttons.

- Enter the name for the group.
This name can be a Unix/Linux group, or an Active Directory group.
- Set the permissions for the group.
At a minimum you will want to assign Read, Write, Execute and Read ACL.
- Select **Save**.

Create bucket using the object APIs

When creating buckets using the object APIs or using tools that call the object APIs, there are a number of headers that determine the behavior.

The following `x-emc` headers are provided:

x-emc-dataservice-vpool

Determines the replication group that will be used to store the objects associated with this bucket. If you do not specify a replication group using the `x-emc-dataservice-vpool` header, ECS will choose the default replication group associated with the namespace.

x-emc-file-system-access-enabled

Configures the bucket for HDFS access. The header must not conflict with the interface that is being used. That is, a create bucket request from HDFS cannot specify `x-emc-file-system-access-enabled=false`.

x-emc-namespace

Specifies the namespace to be used for this bucket. If the namespace is not specified using the S3 convention of host/path style request, then it can be specified using the `x-emc-namespace` header. If the namespace is not specified as this header, the namespace associated with the user is used.

x-emc-retention-period

Specifies the retention period that will be applied to objects in a bucket. Each time a request is made to modify an object in a bucket, the expiration of the retention period for the object is calculated based on the retention period associated with the bucket.

x-emc-is-stale-allowed

Specifies whether the bucket can be accessed during a temporary VDC outage in a federated configuration.

x-emc-server-side-encryption-enabled

Specifies whether objects written to a bucket are encrypted.

x-emc-metadata-search

Specifies one or more user or system metadata values that will be used to create indexes of objects for the bucket. The indexes can be used to perform object searches that can be filtered based on the indexed metadata.

An example of using the S3curl tool to create a bucket is provided:

- [Create a bucket using the S3 API \(with s3curl\) on page 99](#)

Create a bucket using the S3 API (with s3curl)

You can use the S3 API to create a bucket in an replication group. Because ECS uses custom headers (x-emc), the string to sign must be constructed to include these headers. In this procedure the s3curl tool is used; there are also a number of programmatic clients you can use, for example, the S3 Java client.

Before you begin

- ECS must have at least one replication group configured.
- Perl must be installed on the Linux machine on which you will run s3curl.
- You will need to have curl installed and you will need the s3curl module, which acts as a wrapper around curl.

To use s3curl with x-emc headers, minor modifications must be made to the `s3curl` script. These modifications are described in the procedure.

Procedure

1. Obtain a secret key for the user who will create the bucket.

Refer to the article: [Data Access Guide: Obtain secret key to access object storage](#) for details.

2. Obtain the identity of the replication group in which you want the bucket to be created.

You can obtain the replication group identity by using the ECS REST API:

```
GET https://<ECS IP Address>:4443/vdc/data-service/vpools
```

The response provides the name and identity of all data services virtual pools. For example:

```
<data_service_vpools>
<data_service_vpool>
  <creation_time>1403519186936</creation_time>
  <id>urn:storageos:ReplicationGroupInfo:8fc8e19b-edf0-4e81-
bee8-79accc867f64:global</id>
  <inactive>false</inactive>
  <tags/>
  <description>IsilonVPool1</description>
  <name>IsilonVPool1</name>
  <varrayMappings>
    <name>urn:storageos:VirtualDataCenter:1de0bbc2-907c-4ede-
b133-f5331e03e6fa:vdcl</name>
    <value>urn:storageos:VirtualArray:793757ab-ad51-4038-
b80a-682e124eb25e:vdcl</value>
  </varrayMappings>
</data_service_vpool>
</data_service_vpools>
```

Here the ID is `urn:storageos:ReplicationGroupInfo:8fc8e19b-edf0-4e81-bee8-79accc867f64:global`.

3. Set up `s3curl` by creating a `.s3curl` file in which to enter the user credentials.

The `.s3curl` file must have permissions `0600 (rw-/-/-)` when `s3curl.pl` is run.

In the example below, the profile "my_profile" is used to reference the user credentials for the "user@yourco.com" account, and "root_profile" references the credentials for the root account.

```
%awsSecretAccessKeys = (
  my_profile => {
    id => 'user@yourco.com',
    key => 'sZRCTZyk93IWukHEGQ3evPJEvPUq4ASL8Nre0awN'
  },
  root_profile => {
    id => 'root',
    key => 'sZRCTZyk93IWukHEGQ3evPJEvPUq4ASL8Nre0awN'
  },
);
```

4. Add the endpoint that you want to use `s3curl` against to the `.s3curl` file.

This will be the address of your data node or the load balancer that sits in front of your data nodes.

For example:

```
push @endpoints , (
  '203.0.113.10', 'lg1w3183.lss.emc.com',
);
```

5. Modify the `s3curl.pl` script so that it includes the x-emc headers in its "string to sign".

Replace the following lines:

```
elsif ($header =~ /^(?'header'[Xx]-(([Aa][Mm][Zz])|([Ee][Mm][Cc])))-
```

```
[^:~]+): *(?'val'.+)$/) {
    my $name = lc ${header};
    my $value = ${val};
```

with:

```
elsif ($header =~ /^[Xx]-(?:([Aa][Mm][Zz])|([Ee][Mm][Cc]))-
[^:~]+): *(.+)$/) {
    my $name = lc $1;
    my $value = $2;
```

6. Create the bucket using `s3curl.pl`.

Specify the following:

- Profile of the user.
- Identity of the replication group in which to create the bucket (`<vpool_id>`). This must be set using the `x-emc-dataservice-vpool` header.
- Any custom x-emc headers.
- Name of the bucket (`<BucketName>`).

The fully specified command looks like this:

```
./s3curl.pl --debug --id=my_profile --acl public-read-write
--createBucket -- -H 'x-emc-file-system-access-enabled:true'
-H 'x-emc-dataservice-vpool:<vpool_id>' http://<DataNodeIP>:9020/
<BucketName>
```

This example uses the `x-emc-dataservice-vpool` header to specify the replication group in which the bucket will be created and the `x-emc-file-system-access-enabled` header to enable the bucket for file system access, such as for HDFS.

Note that the `-acl public-read-write` argument is optional, but can be used to set permissions to enable access to the bucket. For example, if you intend to access to bucket as HDFS from an environment that is not secured using Kerberos.

If successful (with `--debug` on) you should see output similar to the following:

```
s3curl: Found the url: host=203.0.113.10; port=9020; uri=/S3B4;
query=;
s3curl: ordinary endpoint signing case
s3curl: StringToSign='PUT\n\n\nThu, 12 Dec 2013 07:58:39 +0000\nx-
amz-acl:public-read-write
\nx-emc-file-system-access-enabled:true\nx-emc-dataservice-vpool:
urn:storageos:ReplicationGroupInfo:8fc8e19b-edf0-4e81-
bee8-79acc867f64:global:\n/S3B4'
s3curl: exec curl -H Date: Thu, 12 Dec 2013 07:58:39 +0000 -H
Authorization: AWS
root:AiTcfMDhsi6iSq2rIbHEZon0WNo= -H x-amz-acl: public-read-write -
L -H content-type:
--data-binary -X PUT -H x-emc-file-system-access-enabled:true
-H x-emc-dataservice-
```

```
vpool:urn:storageos:ObjectStore:e0506a04-340b-4e78-
a694-4c389ce14dc8: http://203.0.113.10:9020/S3B4
```

You can list the buckets using the S3 interface, using:

```
./s3curl.pl --debug --id=my_profile http://<DataNodeIP>:9020/
```

Bucket and key naming conventions

Bucket and object/key names must conform to the specification presented here.

- [S3 bucket and object naming in ECS on page 102](#)
- [OpenStack Swift container and object naming in ECS on page 103](#)
- [Atmos bucket and object naming in ECS on page 103](#)
- [CAS pool and object naming in ECS on page 103](#)

Note

If you want to use a bucket for HDFS, you should not use underscores in the bucket name as they are not supported by the URI Java class. For example, `viprfs://my_bucket.ns.site/` will not work as this is an invalid URI and is thus not understood by Hadoop.

Namespace name

The following rules apply to the naming of ECS namespaces:

- Cannot be null or an empty string
- Length range is 1..255 (Unicode char)
- Valid characters are defined by regex `/[a-zA-Z0-9- _]+/`. Hence:
 - Alphanumeric characters
 - Special characters: hyphen (-) and underscore (_).

S3 bucket and object naming in ECS

This topic details the rules that apply to the naming of buckets and objects when using the ECS S3 Object API.

Bucket name

The following rules apply to the naming of S3 buckets in ECS:

- Names must be between one and 255 characters in length. (S3 requires bucket names to be from 1 to 255 characters long).
- Names can include dot (.), hyphen (-), and underscore (_) characters and alphanumeric characters ([a-zA-Z0-9]).
- Names can start with a hyphen (-) or alphanumeric character.
- The name does not support:
 - Starting with a dot (.)
 - Containing a double dot (..)
 - Ending with a dot (.)

- Name must not be formatted as IPv4 address.

You can compare this with naming restriction specified by the S3 specification: <http://docs.aws.amazon.com/AmazonS3/latest/dev/BucketRestrictions.html>.

Object Name

The following rules apply to the naming of ECS S3 objects:

- Cannot be null or an empty string
- Length range is 1..255 (Unicode char)
- No validation on characters.

OpenStack Swift container and object naming in ECS

This topic details the rules that apply to the naming of buckets and objects when using the ECS OpenStack Swift Object API.

Container Name

The following rules apply to the naming of Swift containers:

- Cannot be null or an empty string
- Length range is 1..255 (Unicode char)
- Valid characters are defined by regex `/[a-zA-Z0-9\\.\- _]+/`
 - Alphanumeric characters
 - Special characters: dot (.), hyphen (-), and underscore (_).

Object Name

The following rules apply to the naming of Swift objects:

- Cannot be null or an empty string
- Length range is 1..255 (Unicode char)
- No validation on characters.

Atmos bucket and object naming in ECS

This topic details the rules that apply to the naming of buckets and objects when using the ECS Atmos Object API.

Subtenant (bucket)

This is created by the server, so the client does not need to know the naming scheme.

Object name

The following rules apply to the naming of Atmos objects:

- Cannot be null or an empty string
- Length range is 1..255 (Unicode char)
- No validation on characters.

Name should be percent-encoded UTF-8.

CAS pool and object naming in ECS

This topic details the rules that apply to the naming of CAS pools and objects ('clips' in CAS terminology) when using the CAS API.

CAS pool naming

The following rules apply to the naming of CAS pools in ECS:

- a maximum of 255 characters
- cannot contain: ' " / & ? * < > <tab> <newline> or <space>

Clip naming

There are no user defined keys in the CAS API. When an application using CAS API creates a clip, it opens a pool, creates a new clip, and adds tags, attributes, streams etc. After a clip is complete it is written to a device.

A corresponding clip ID is returned by CAS engine and can be referred to using <pool name>/<clip id>.

CHAPTER 11

Configure NFS file access

- [Configure file access](#)..... 106

Configure file access

ECS enables object buckets to be configured for access as NFS export file systems.

To enable Unix users to access the filesystem, ECS provides a mechanism for mapping ECS object users to Unix users. An ECS bucket always has an owner, and mapping the bucket owner to a Unix ID will give that Unix user permissions on the filesystem. In addition, ECS enables the assignment of a default custom group to the bucket so that members of a Unix group mapped to the ECS default custom group can access the bucket.

In addition, ECS supports multi-protocol access, so that files written using NFS can also be accessed using S3 and OpenStack Swift object protocols. Similarly, objects written using S3 and OpenStack Swift object protocols can be made available through NFS. In the same way as for the bucket itself, objects and directories created using object protocols can be accessed by Unix users and Unix group members by mapping the object users and groups.

Note

File support is not enabled for ECS 2.2

CHAPTER 12

Set the Base URL

- [Introduction](#)..... 108
- [Bucket addressing](#)..... 108
- [Add a Base URL](#)..... 110

Introduction

Applications that are written to use Amazon S3 can be enabled to use ECS object storage by setting the Base URL parameter. The Base URL is set by default to `amazonaws.com`. This article describes how to set the Base URL and ensure that requests are routed to ECS.

The following sections describe the addressing scheme supported by ECS, the use of the Base URL parameter, and the mechanism for setting the Base URL parameter.

- [Bucket addressing on page 108](#)
- [Add a Base URL on page 110](#)

Bucket addressing

The ECS S3 service provides a number of ways in which to identify the bucket against which the operation defined in a request should be performed.

When using the Amazon S3 service, all buckets names must be unique. However, the ECS S3 service supports the use of a namespace, which can be used in addition to the bucket name and allows buckets in different namespaces to have the same name. By assigning a namespace to each tenant, a tenant can assign bucket names without regard for the names currently used by other tenants. If no namespace is specified in a request, ECS uses the default namespace associated with the tenant to which the user making the request belongs.

The namespace that refers to the location of an object can be specified in the `x-emc-namespace` header of an HTTP request. ECS also supports extraction of the location from the host header and allows the following Amazon S3 compatible addressing schemes:

- [Virtual Host Style Addressing on page 108](#)
- [Path Based Addressing on page 108](#)

Virtual Host Style Addressing

In the virtual host addressing scheme, the bucket name appears in the hostname. For example, the bucket called "mybucket" on host `ecs1.yourco.com`, would be accessed using:

```
http://mybucket.ecs1.yourco.com
```

In addition, ECS also allows the inclusion of a namespace in the address. For example:

```
<bucketname>.<namespace>.ecs1.yourco.com
```

To use this style of addressing, you need to configure ECS so that it knows which part of the URL is the bucket name. This is done by configuring the Base URL. In addition, you need to ensure that your DNS system can resolve the address. The following sections provide more information:

- [DNS Configuration on page 109](#)
- [Base URL on page 109](#)

Path Based Addressing

In the path based addressing scheme, the bucket name is added to the end of the path. For example:

```
ecs1.yourco.com/mybucket
```

A namespace can be specified using the x-emc-namespace header.

DNS Configuration

When accessing ECS storage using the S3 service, you will need to ensure that the URL resolves to the address of the ECS data node, or the data node load balancer.

Where your application uses path-style addressing, this is simply a case of ensuring that the base name is resolvable by DNS. For example, if your application normally issues requests in the form `ecs1.yourco.com/bucket`, you will need to have a DNS entry that resolves `ecs1.yourco.com` to the IP address of your load balancer used for access to ECS nodes. If you are using the Amazon service this URI will be of the form: `s3-eu-west-1.amazonaws.com`.

Where your application is using virtual host style addressing, the URL will include the bucket name and can include a namespace. Under these circumstances, you will need to ensure that you include a DNS entry that will resolve the virtual host style address. You can do this by using a wildcard in the DNS entry.

For example, if your application normally issues requests in the form `bucket.s3.yourco.com`, you will need to have two DNS entries.

- `ecs1.yourco.com`
- `*.ecs1.yourco.com`

Or, if you are using an application that previously connected to the Amazon S3 service, using `bucket.s3.amazonaws.com`, the entries would be:

- `s3.amazonaws.com`
- `*.s3.amazonaws.com`


These entries allow the base name to be resolved when issuing service-level commands (for example, list buckets) and the virtual host style bucket address to be resolved.

If you are creating an SSL certificate for this service, it should have the wildcard entry on the name of the certificate and the non-wildcard version as a Subject Alternate Name.

Base URL

If you have an S3 application that uses virtual host style addressing and you want to use it to connect to ECS, the Base URL must be set to enable ECS to know which part of the address refers to the bucket and, optionally, namespace. The Base URL can be set using the ECS Portal, or using the ECS Management REST API, and requires the ECS System Administrator role.

The **Base URL Management** page shows the Base URLs that have been created and how ECS should use the them.



Name	Base URL	Use with Namespace	Actions
DefaultBaseUri	s3.amazonaws.com	No	Edit Delete

In order that ECS knows how to treat the bucket location prefix, the Base URL must be configured by choosing one of the following options.

- Use Base URL with namespace
- Use Base URL without namespace

When processing a request, ECS will:

1. Try to extract namespace from the x-emc-namespace header. If found, skip the steps below and process the request.
2. Get the hostname of the URL from the host header and check if the last part of the address matches any of the configured Base URLs.
3. Where there is a BaseURL match, use the prefix part of the hostname (the part left when the Base URL is removed), to obtain the bucket location.

The following examples demonstrate how ECS handles incoming HTTP requests with different structures.

Example 1

```
Host:          baseball.image.emc.finance.com
BaseURL:       finance.com
Use BaseURL with namespace enabled

Namespace:     emc
Bucket Name:   baseball.image
```

Example 2

```
Host:          baseball.image.emc.finance.com
BaseURL:       finance.com
Use BaseURL without namespace enabled

Namespace:     null (Use other methods to determine namespace)
Bucket Name:   baseball.image.emc
```

Example 3

```
Host:          baseball.image.emc.finance.com
BaseURL:       not configured

Namespace:     null (Use other methods to determine namespace.)
Bucket Name:   null (Use other methods to determine the bucket name.)

ViPR Controller treats this request as a path-style request.
```

Add a Base URL

This operation is only necessary if you use object clients that encode the location of an object, its namespace and bucket, in a URL. In that case you can specify a base URL that will be used, together with the namespace, as the path to objects in a tenant.

Before you begin

This operation requires the System Admin role in ECS.

You must ensure that the domain specified in a request that uses a URL to specify an object location resolves to the location of the ECS data node or a load balancer that sits in front of the data nodes.

Procedure

1. At the ECS Portal, select **Settings > Object Base URLs**.
2. Select **New Base URL**.

The **New Base URL** page is displayed.

3. Enter the name of the Base URL. This will provide additional information about the base URL when looking at the base URL table.
4. Enter the Base URL.

If your objects location URLs are in the form: `https://mybucket.mynamespace.acme.com` (that is, `bucket.namespace.baseurl`) or `https://mybucket.acme.com` (that is, `bucket.baseurl`), the base URL would be `acme.com`.

You can specify which format in the Namespace selector.

5. Choose the format in which your object address is encoded in the URL: with a namespace or without a namespace.
6. Select **Save**.

Set the Base URL

PART 3

Monitor

[Chapter 13, "Monitoring basics"](#)

[Chapter 14, "Monitor metering"](#)

[Chapter 15, "Monitor events"](#)

[Chapter 16, "Monitor capacity utilization"](#)

[Chapter 17, "Monitor traffic metrics"](#)

[Chapter 18, "Monitor hardware health"](#)

[Chapter 19, "Monitor node and process health"](#)

[Chapter 20, "Monitor chunk summary"](#)

[Chapter 21, "Monitor erasure coding"](#)

[Chapter 22, "Monitor recovery status"](#)

[Chapter 23, "Monitor disk bandwidth"](#)

[Chapter 24, "Monitor geo-replication"](#)

[Chapter 25, "Service logs"](#)

CHAPTER 13

Monitoring basics

- [Using monitoring pages](#)..... 116

Using monitoring pages

Introduces the basic techniques for using monitoring pages in the ECS Portal.

The ECS Portal monitoring pages share a set of common interactions. These are:

- Refresh: the refresh icon allows you to update the monitoring display with the latest data.


Figure 10 Refresh



- Filter: fill in filter fields and the date range and select **Filter** to display result rows that match all filter fields. The default date range is always yesterday and today.
- Drill down displays with breadcrumbs: Breadcrumbs let you quickly drill up when you have drilled down into detail screens. See the "Navigating with Breadcrumbs" figure below.
- History charts with left to right mouse-overs: Get detailed charts showing hourly snapshots for the last five days worth of data which you can browse through using your mouse as a left-to-right chart cursor. See the example below. See the "History chart with active cursor" figure below.

The standard monitoring filter provides the ability to narrow results by time and date. It is available on several monitoring pages. Some pages have additional filter types. Select a time range, then a date range, click apply, and the Filter panel closes and the page content updates. Select the pin icon to keep the Filter panel open after applying the filter.

Figure 11 Open Filter panel with criteria selected




When the Filter panel closes, a summary of the applied filter displays along with a Clear Filter command and a Refresh command.

Figure 12 Closed Filter panel showing summary of applied filter




Highlighted text in a table row indicates a link to a detail display. Selecting the link drills down to the next level of detail. On drill down displays, a path string shows your current location in the sequence of drill down displays. This path string is called a breadcrumb trail or breadcrumbs for short. Selecting any highlighted breadcrumb jumps up to the associated display.

Figure 13 Navigating with breadcrumbs



When you select a **History** button, all available charts for that row display below the table. Mouse over a chart from left to right to see a vertical line that helps you find a specific date-time point on the chart. A pop-up display shows the value and timestamp for that point.

Figure 14 History chart with active cursor



CHAPTER 14

Monitor metering

- [Monitor metering data](#).....120

Monitor metering data

Describes how to display metering data for namespaces or buckets within namespaces for a specified time period.

The available metering data is detailed in [Metering data on page 121](#).

Using the ECS Management REST API you can retrieve data programmatically with custom clients. Support for this feature and other features that enable a tenant to be managed is provided in [Manage a tenant on page 74](#). The ECS Management REST API Reference is provided [here](#).

Procedure

1. At the ECS Portal, select **Monitor** > **Metering**.
2. From the **Date Time Range** menu, select the period for which you want to see the metering data. Select **Current** to view the current metering data. Select **Custom** to specify a custom date-time range.

If you select **Custom**, use the **From** and **To** calendars to choose the time period for which data will be displayed.

Metering data is kept for 60 days.

3. Select the namespace for which you want to display metering data. To narrow the list of namespaces, type the first few letters of the target namespace and click the magnifying glass icon.

If you are a Namespace Admin, you will only be able to select your namespace.

4. Click the + icon next to each namespace you want to see object data for.
5. Optionally, click the + icon next to each bucket you want to see object data for.

To narrow the list of buckets, type the first few letters of the target bucket and click the magnifying glass icon.

If you do not specify a bucket, the object metering data will be the totals for all buckets in the namespace.

Figure 15 Metering page with criteria selected

The screenshot shows the 'Metering' page with the following configuration:

- Date Time Range:** Custom
- From:** 2015-12-13 09:35
- To:** 2015-12-14 09:35
- Select Namespace:** ns-base
- Select Buckets:** (Empty)
- Selected Table:**

Namespace	Bucket
ns-base	
ns-base	bkt-base

- Click **Apply** to display the metering data for the selected namespace and bucket, and time period.

Metering data

Object metering data for a specified namespace, or a specified bucket within a namespace, can be obtained for a defined time period at the ECS portal **Monitor > Metering** page.

The metering information that is provided is shown in the table below.

Table 8 Bucket and namespace metering

Attribute	Description
Total Size (GB)	Total size of the objects stored in the selected namespace or bucket at the end time specified in the filter.
Object Count	Number of objects associated with the selected namespace or bucket at the end time specified in the filter.
Objects Created	Number of objects created in the selected namespace or bucket in the time period.
Objects Deleted	Number of objects deleted from the selected namespace or bucket in the time period.
Bandwidth Ingress (MB)	Total of incoming object data (writes) for the selected namespace or bucket during the specified period.

Table 8 Bucket and namespace metering (continued)

Attribute	Description
Bandwidth Egress (MB)	Total of outgoing object data (reads) for the selected namespace or bucket during the specified period.

Note

Metering data is not available immediately as it can take a significant amount of time to gather the statistics for data added to the system and deleted from the system.

CHAPTER 15

Monitor events

- [About event monitoring](#).....124
- [Monitor audit data](#)..... 124
- [Monitor alerts](#).....125

About event monitoring

Describes the event monitoring functions of the ECS Portal.

The Events page under the **Monitor** menu displays:

- **Audit panel:** All activity by users working with portal, the ECS REST API, and the ECS CLI.
- **Alerts panel:** Alerts raised by the ECS system.

Event data through the ECS Portal is limited to 30 days. If you need to keep event data for longer periods, consider using ViPR SRM.

Monitor audit data

Use the Audit panel of the Events page to view and manage audit data.

Procedure

1. Select **Audit**.
2. Optionally, select **Filter**.
3. Specify a **Date Time Range** and adjust the **From** and **To** fields and time fields.
4. Select a **Namespace**.
5. Click **Apply**.

The screenshot shows the 'Events' interface with the 'Alerts' tab selected. A filter panel is open, showing the following settings:

- Date Time Range:** Last Month
- From:** 2015-11-10
- To:** 2015-12-10
- Namespace:** ns1

Buttons for 'Apply' and 'Clear Filter' are visible. Below the filter panel is a table of events:

User ID	Description	Namespace	Timestamp
root	Object user user1 has been created	ns1	2015-12-03 12:31:00
root	New password has been set for object user user1	ns1	2015-12-03 12:31:00
root	New password has been set for object user user1	ns1	2015-12-03 12:31:00
root	Namespace ns1 has been created	ns1	2015-12-03 12:31:00
root	Bucket bucket1 has been created	ns1	2015-12-03 12:36:00
user1	Owner of bucket1 bucket has changed	ns1	2015-12-03 12:36:00

Monitor alerts

Use the Alerts panel of the Events page to view and manage system alerts.

Procedure

1. Select **Alerts**.
2. Optionally, click **Filter**.
3. Select your filters. The alerts filter adds filtering by **Severity** and **Type**, as well as an option to **Show Acknowledged Alerts**, which retains the display of an alert even after acknowledged by the user.

Alert types must be entered exactly. Alert types are described in the table below:

Table 9 Alert types

Alert (type exactly as shown)	Description
License	Raised for license or capacity alerts.
Notify	Raised for miscellaneous alerts.
Fabric	Raised when system issues detected.
BUCKET_HARD_QUOTA_EXCEEDED	Raised when the quota on a bucket is exceeded

Table 9 Alert types (continued)

Alert (type exactly as shown)	Description
CHUNK_NOT_FOUND	Raised when chunk data is not found.
DTSTATUS_RECENT_FAILURE	Raised when the status of a data table is bad.
NAMESPACE_HARD_QUOTA_EXCEEDED	Raised when the quota on a namespace is exceeded

The screenshot shows the 'Events' monitoring interface. At the top, there are tabs for 'Audit' and 'Alerts'. Below this is a 'Filter' section with the following controls:

- Date Time Range:** A dropdown menu set to 'Last 2 Months'.
- From:** A date/time picker set to 2015-10-10 13:30.
- To:** A date/time picker set to 2015-12-10 13:30.
- Severity:** A dropdown menu set to 'Critical'.
- Type:** A dropdown menu set to 'Fabric'.
- Namespace:** A dropdown menu set to 'Select an option'.
- Show Acknowledged Alerts:** A checkbox that is checked.

Below the filter section are two buttons: 'Apply' (in blue) and 'Clear Filter'. At the bottom, a table displays the filtered events:

Severity	Description	Type	Namespace	Timestamp	Actions
Critical	Service Health Suspect Event	Fabric		2015-12-08 04:52:41	Acknowledge
Critical	Service Health Suspect Event	Fabric		2015-12-08 04:52:40	Acknowledge

4. Click **Apply**.
5. Next to each event, click the acknowledge button to acknowledge and dismiss the message (if the **Show Acknowledged Alerts** filter is not selected).

CHAPTER 16

Monitor capacity utilization

- [Monitor capacity](#) 128

Monitor capacity

You can monitor the capacity utilization of storage pools, nodes, and disks.

The capacity tables and displays are shown in [Storage capacity data on page 128](#). Each table has an associated History display that enables you to see how the table data has changed over time.

Using the ECS Management REST API you can retrieve data programmatically using custom clients. Support for this feature and other features that enable a tenant to be managed is provided in [Manage a tenant on page 74](#). The ECS Management REST API Reference is provided [here](#).

Procedure

1. At the ECS Portal, select **Monitor** > **Capacity Utilization**.
2. You can drill down into the nodes and to individual disks by selecting the appropriate link in the table.

Guidance on navigating the tables is provided in [Using monitoring pages on page 116](#).

3. To display the way in which the capacity has changed over time, select **History** for the storage pool, node, or disk that you are interested in.

Storage capacity data

Storage capacity for storage pools, nodes, and disks can be displayed at the **Monitor** > **Capacity Utilization** page.

The capacity utilization areas are described in:

- [Storage Pool Capacity on page 128](#)
- [Node Capacity Utilization on page 129](#)
- [Disk Capacity Utilization on page 130](#)

Table values represent current values when the Current Filter is selected, or average values of the metric over the period selected in the filter.

Storage Pool Capacity


Table 10 Capacity Utilization: Storage Pool

Attribute	Description
Storage Pool	Name of the storage pool.
Nodes	Number of nodes in the storage pool. Click node number to open: Node Capacity Utilization on page 129
Disks	Number of disks in the storage pool.
Usable Capacity	Total usable capacity. This is total of the capacity already used and the capacity still free for allocation.
Used Capacity	Used capacity in the storage pool.
Available Capacity	Capacity available for use.
Actions	History provides a graphic display of the data.

Table 10 Capacity Utilization: Storage Pool (continued)

Attribute	Description
	If the Current filter is selected, the History button displays default history for the last 24 hours.

The history display for the storage pool capacity utilization table is shown below.




Node Capacity Utilization

Table 11 Capacity Utilization: Node

Attribute	Description
Nodes	IP address of the node.
Disks	Number of disks associated with the node. Click node number to open: Disk Capacity Utilization on page 130
Usable Capacity	Total usable capacity provided by the disks within the node. This is total of the capacity already used and the capacity still free for allocation.
Used Capacity	Capacity used within the node.
Available Capacity	Remaining capacity available in the node.
Node Status	Check mark indicates the node status is Good.
Actions	History provides a graphic display of the data. If the Current filter is selected, the History button displays default history for the last 24 hours.

The history display for the node capacity utilization table is shown below.




Disk Capacity Utilization

Table 12 Capacity Utilization: Disk

Attribute	Description
Disks	Disk identifier.
Usable Capacity	Usable capacity provided by the disk.
Used Capacity	Capacity used on the disk.
Available Capacity	Remaining capacity available on the disk.
Disk Status	Check mark indicates the disk status is Good.
Actions	History provides a graphic display of the data. If the Current filter is selected, the History button displays default history for the last 24 hours.

The history display for the disk utilization table is shown below.



CHAPTER 17

Monitor traffic metrics

- [Monitor network traffic](#)..... 132

Monitor network traffic

Describes the ECS Portal monitoring page for network traffic.

The **Monitor > Traffic Metrics** page provides network traffic metrics at the virtual data center or the individual node level. The charts show data for the last seven days. Table values represent current values when the Current Filter is selected, or average values of the metric over the period selected in the filter.


Table 13 Network traffic metrics

Metric label	Description
R Latency (ms)	Average latency for reads in milliseconds.
W Latency (ms)	Average latency for writes in milliseconds.
R Bandwidth	Bandwidth for reads.
W Bandwidth	Bandwidth for writes.
R Transactions (per/s)	Read transactions per second.
W Transactions (per/s)	Write transactions per second.
Recent Transaction Failures per type	For each error code that occurred in the monitoring period, display that code's percent of the total errors.
History	History provides a graphic display of the data. If the Current filter is selected, the History button displays default history for the last 24 hours.

Procedure

1. Select **Monitor > Traffic Metrics**.
2. Locate the target VDC name.
3. Optionally, select the VDC name to drill down to the nodes display.
4. Select **History** button for the target VDC or node.

Figure 16 Network traffic charts for a VDC



Monitor traffic metrics

CHAPTER 18

Monitor hardware health

- [Monitor hardware](#).....136

Monitor hardware

Describes how to use the **Monitor > Hardware Health** page.

Hardware health is designated by three states:

- **Good:** The hardware component is in normal operating condition.
- **Suspect:** Either the hardware component is transitioning from good to bad because of decreasing hardware metrics, or there is a problem with a lower-level hardware component, or the hardware is not detectable by the system because of connectivity problems.
- **Bad:** The hardware needs replacement.


In the case of disks, these states have the following meanings:

- **Good:** The system is actively reading from and writing to the disk.
- **Suspect:** The system no longer writes to the disk but will read from it. Note that "swarms" of suspect disks are likely caused by connectivity problems at a node. These disks will transition back to Good when the connectivity issues clear up.
- **Bad:** The system neither reads from nor writes to the disk. Replace the disk. Once a disk has been identified as bad by the ECS system, it cannot be reused anywhere in the ECS system. Because of ECS data protection, when a disk fails, copies of the data that was once on the disk are recreated on other disks in the system. A bad disk only represents a loss of capacity to the system--not a loss of data. When the disk is replaced, the new disk does not have data restored to it. It simply becomes raw capacity for the system.

Procedure

1. Select **Monitor > Hardware Health**.
2. Locate the table row for the target storage pool.
3. Optionally, select a storage pool name to drill down to the node display.
4. Optionally, select a node endpoint to drill down to the disk display.

Figure 17 Hardware Health



Nodes	Node Status	Good Disks	Bad Disks	Suspect Disks
10.241.51.201	Good	60	0	0
10.241.51.205	Good	59	0	0
10.241.51.204	Good	60	0	0
10.241.51.203	Good	60	0	0
10.241.51.208	Good	60	0	0
10.241.51.207	Good	60	0	0
10.241.51.206	unknown	60	0	0

CHAPTER 19

Monitor node and process health

- [Monitor node and process health](#).....138

Monitor node and process health

Describes the ECS Portal monitoring page for node and process health.

The **Monitor > Node & Process Health** page provides metrics that can help assess the health of the VDC, node, or node process.

Table values represent current values when the Current Filter is selected, or average values of the metric over the period selected in the filter.

Table 14 VDC, node, and process health metrics

Metric label	Level	Description
Avg. NIC Bandwidth	VDC and Node	Average bandwidth of the network interface controller hardware used by the selected VDC or node.
Avg. CPU Usage (%)	VDC and Node	Average percent of the CPU hardware used by the selected VDC or node.
Avg. Memory Usage	VDC and Node	Average usage of the aggregate memory available to the VDC or node.
Relative NIC (%)	VDC and Node	Percent of the available bandwidth of the network interface controller hardware used by the selected VDC or node.
Relative Memory (%)	VDC and Node	Percent of the memory used relative to the memory available to the selected VDC or node.
CPU (%)	Process	Percent of the node's CPU used by the process.
Memory Usage	Process	The memory used by the process.
Relative Memory (%)	Process	Percent of the memory used relative to the memory available to the process.
Avg. # Thread	Process	Average number of threads used by the process.
Last Restart	Process	The last time the process restarted on the node.


Table 14 VDC, node, and process health metrics (continued)

Metric label	Level	Description
Actions	All	<p>History provides a graphic display of the data.</p> <p>If the Current filter is selected, the History button displays default history for the last 24 hours.</p>

Procedure

1. Locate the table row for the target VDC.
2. Optionally, select the VDC name to drill down to a table with rows for each node in the VDC.
3. Optionally, select the a node endpoint to drill down to a table with rows for each process running on the node.
4. Select the **History** button for the target VDC, node, or process.

Figure 18 Node & Process Health



CHAPTER 20

Monitor chunk summary

- [Monitor chunks](#)..... 142

Monitor chunks

Describes the ECS Portal monitoring page for chunks.

This page reports statistics for sealed chunks in the local zone. A sealed chunk is one that can no longer accept writes. It is immutable.

Table 15 Chunk tables


Table	Description
Chunk Count of Each Type	Shows number and percentage of sealed chunks for different chunk types per each storage pool configured in the local zone.
Total Length of Each Chunk Type	Shows total logical size of sealed chunks for different chunk types per each storage pool configured in the local zone.
Avg Sealed Length of Each Type	Shows average logical size of sealed chunks for different chunk types per each storage pool configured in the local zone.

Table 16 Chunk metrics

Metric label	Description
Storage pool	This column provides the list of storage pools configured in the local VDC. Each row provides chunk metrics for the specified storage pool.
User data	This column provides relevant data for the user data (repository) chunks in the storage pool.
Metadata	This column provides relevant data for the system metadata chunks in the storage pool.
Geo data	Geo chunks are chunks containing replicas of data from other zones (VDCs). This field provides relevant data for the geo-copy chunks in the storage pool.
XOR	XOR chunks are chunks that save disk space by using the XOR algorithm to compress data from other chunks and replace those chunks with an XOR chunk. This field provides relevant data for the XOR chunks in the storage pool.
Total	The total number of chunks in the storage pool.

Chunk metrics

Figure 19 Chunk Summary



Monitor chunk summary

CHAPTER 21

Monitor erasure coding

- [Monitor erasure coding](#)..... 146

Monitor erasure coding

Describes how to use the **Monitor > Erasure Coding** page.

The erasure coding display monitors the amount of total user data and erasure coded data in a local storage pool. It also shows the amount of data pending erasure coding the current rate, and estimated completion time. Charts hold seven days worth of data.


Table values represent current values when the Current Filter is selected, or average values of the metric over the period selected in the filter.

Table 17 Erasure coding metrics

Column	Description
Storage Pool	
Total Data for Erasure Coding	The total logical size of all data chunks in the storage pool, which are subject to EC.
Total Data Erasure Coded	The total logical size of all erasure-coded chunks in the storage pool.
% Erasure Coded Data	The percent of data in the storage pool that is erasure coded.
Rate of Erasure Coding	The rate at which any current data waiting for erasure coding is being processed.
Est Time to Complete	The estimated completion time extrapolated from the current erasure coding rate.
Actions	History provides a graphic display of the data. If the Current filter is selected, the History button displays default history for the last 24 hours.

Procedure

1. Select **Monitor > Erasure Coding**.
2. Locate the table row for the target storage pool.
3. Select the **History** button.



Monitor erasure coding

CHAPTER 22

Monitor recovery status

- [Monitor recovery status](#).....150

Monitor recovery status

Describes how to use the **Monitor > Recovery Status** page.

Recovery is the process of rebuilding data after any local condition that results in bad data (chunks). This table includes one row for each storage pool in the local zone.


Table values represent current values when the Current Filter is selected, or average values of the metric over the period selected in the filter.

Table 18 Recovery metrics

Column	Description
Storage Pool	Lists each storage pool in the local zone.
Total Amount Data to be Recovered	Logical size of the data yet to be recovered.
Recovery Rate	Rate data is being recovered in the specified storage pool in.
Time to Completion	Estimated time to complete the recovery extrapolated from the current recovery rate.
Actions	History provides a graphic display of the data. If the Current filter is selected, the History button displays default history for the last 24 hours.

Procedure

1. Select **Monitor > Recovery Status**.
2. Locate the table row for the target storage pool.
3. Select the **History** button.



CHAPTER 23

Monitor disk bandwidth

- [Monitor disk bandwidth](#)..... 152

Monitor disk bandwidth

Describes the ECS Portal monitoring page for disk bandwidth.

The **Monitor > Disk bandwidth** page provides disk use metrics at the virtual data center or the individual node level. There is one row for read and another for write for each VDC or node. The charts show data for the last seven days.

Table values represent current values when the Current Filter is selected, or average values of the metric over the period selected in the filter.


Table 19 Disk bandwidth metrics

Metric label	Description
Total	Total disk bandwidth used for either read or write operations.
Hardware Recovery	Rate of disk bandwidth used to recover data after hardware failures.
Erasur Encoding	Rate of disk bandwidth used in system erasure coding operations.
XOR	Rate of disk bandwidth used in the system's XOR data protection operations. Note that XOR operations occur for systems with three or more sites (VDCs).
Consistency Checker	Rate of disk bandwidth used to check for inconsistencies between protected data and its replicas.
Geo	Rate of disk bandwidth used to support geo replication operations.
User Traffic	Rate of disk bandwidth used by object users.
Actions	History provides a graphic display of the data. If the Current filter is selected, the History button displays default history for the last 24 hours.

Procedure

1. Select **Monitor > Disk Bandwidth**.
2. Locate the target VDC name and either the Read or Write table row for that VDC.
3. Optionally, select the **Node Count** to drill down to a table with rows for the nodes in the VDC.
4. Select the **History** button for the VDC or node.

Figure 20 Disk Bandwidth



Monitor disk bandwidth

CHAPTER 24

Monitor geo-replication

- [Introduction to Geo-replication monitoring](#)..... 156
- [Monitor geo-replication: Rate and Chunks](#)..... 156
- [Monitor geo-replication: Recovery Point Objective \(RPO\)](#)..... 157
- [Monitor geo-replication: Failover Processing](#)..... 158
- [Monitor geo replication: Bootstrap Processing](#)..... 159

Introduction to Geo-replication monitoring

Describes the four types of geo-replication monitoring

Geo-replication monitoring includes four different pages:

- [Rate and Chunks on page 156](#)
- [Recovery Point Objective \(RPO\) on page 157](#)
- [Failover on page 158](#)
- [Bootstrap Processing on page 159](#)

Monitor geo-replication: Rate and Chunks

Describes the monitoring metrics found in the ECS Portal **Monitor > Geo-Replication > Rate and Chunks** page.

This page provides fundamental metrics about the network traffic for geo-replication and the chunks waiting for replication by replication group or remote zone (VDC).

Table 20 Rate and Chunk columns

Column	Description
Replication Group	Lists the replication groups this zone (VDC) participates in. Click a replication group to see a table of remote zones in the replication group and their statistics. Click the Replication Groups link above the table to return to the default view.
Write Traffic	The current rate of writes to all remote zones or individual remote zone in the replication group.
Read Traffic	The current rate of reads to all remote zones or individual remote zone in the replication group.
User Data Pending Replication	The total logical size of user data waiting for replication for the replication group or remote zone.
Metadata Pending Replication	The total logical size of metadata waiting for replication for the replication group or remote zone.
Data Pending XOR	The total logical size of all data waiting to be processed by the XOR compression algorithm in the local zone for the replication group or remote zone.

Figure 21 Geo replication: Rate and Chunks



Monitor geo-replication: Recovery Point Objective (RPO)


Describes the table fields found in the ECS Portal **Monitor > Geo-Replication > RPO** page.

Recovery Point Objective (RPO) refers to the point in time in the past to which you can recover. The value here is the oldest data at risk of being lost if a local VDC fails before replication is complete.

Table 21 RPO columns

Column	Description
Remote Replication Group\Remote Zone	At the VDC level, lists all remote replication groups the local zone participates in. At the replication group level, this column lists the remote zones in the replication group. The data listed is the system identifier for the VDC or replication group as an URN.
Overall RPO (mins)	The recent time period for which data might be lost in the event of a local zone failure.

Figure 22 RPO



Monitor geo-replication: Failover Processing

Describes the metrics found in the ECS Portal **Monitor** > **Geo-Replication** > **Failover Processing** page.

The Failover Processing page provides metrics on the process to re-replicate data following permanent failure of a remote zone.

Table 22 Failover columns

Field	Description
Replication Group	Lists the replication groups that the local zone is a member of. The data listed is the system identifier for the replication group as an URN.
Failed Zone	Identifies failed zone that is part of the replication group.
User Data Pending Re-replication	Chunks which used to be replicated to the failed zone have to be re-replicated to a different zone. The field reports logical size of all user data (repository) chunks waiting re-replication to a different zone instead of the failed one"
Metadata Pending Re-replication	Chunks which used to be replicated to the failed zone have to be re-replicated to a different zone. This field reports logical size of all system data chunks waiting re-replication to a different zone instead of the failed one.
Data Pending XOR Decoding	Shows the count and total logical size of chunks waiting to be retrieved by the XOR compression scheme.
Failover State	<ul style="list-style-type: none"> • BLIND_REPLAY_DONE • REPLICATION_CHECK_DONE: The process that makes sure that all replication chunks are in an acceptable state has completed successfully. • CONSISTENCY_CHECK_DONE: The process that makes sure that all system metadata is fully consistent with other replicated data has completed successfully. • ZONE_SYNC_DONE: The synchronization of the failed zone has completed successfully. • ZONE_BOOTSTRAP_DONE: The bootstrap process on the failed zone has completed successfully. • ZONE_FAILOVER_DONE: The failover process has completed successfully.
Failover Progress	A percentage indicator for the overall status of the failover process.

Figure 23 Failover

Replication Group	Failed Zone	User Data Pending Re-replication	Metadata Pending Re-replication	Data Pending XOR Decoding	Failover State	Failover Progress
RG	VDC2	0.0 B	0.0 B	0.0 B	6 - ZONE_FAILOVER_DONE	100

Monitor geo replication: Bootstrap Processing

Describes the monitoring found in the ECS Portal **Monitor > Geo Replication > Bootstrap Processing** page.

Bootstrapping refers to the process of copying necessary metadata to a replication group that has added a zone.


Table 23 Bootstrap Processing columns

Column	Description
Replication Group	This column provides the list of replication groups the local zone participates in with new zones being added. Each row provides metrics for the specified replication group.
Added Zone	The zone being added to the specified replication group.
User Data Pending Replication	The logical size of all user data (repository) chunks waiting replication to the new zone being added.
Metadata Pending Replication	The logical size of all system metadata waiting replication to the new zone being added.
Bootstrap State	<ul style="list-style-type: none"> Started: The system has begun preparing to add the zone to the replication group. BlindReplayDone ReplicationCheckDone: The process that checks to make sure that all replication chunks are in an acceptable state has completed successfully. ConsistencyCheckDone: The process that makes sure that all system metadata is fully consistent with other replicated data has completed successfully. ZoneSyncDone: The synchronization of the failed zone has completed successfully. ZoneBootstrapDone: The bootstrap process on the failed zone has completed successfully. Done: The entire bootstrap process has completed successfully.

Table 23 Bootstrap Processing columns (continued)

Column	Description
Bootstrap Progress (%)	The completion percent of the entire bootstrap process.

Figure 24 Bootstrap processing



CHAPTER 25

Service logs

- [Service logs](#).....162
- [ECS service log locations](#).....162

Service logs

Describes the location and function of the ECS service logs.

Storage administrators can access ECS service logs if you have permission to access a node and access the logs. Using the Monitoring pages of the ECS Portal is usually a better way to understand the state of your system.

ECS service log locations

Describes the location and content of ECS service logs.

You can access ECS service logs directly by an SSH session on a node. Change to the following directory: `/opt/emc/caspian/fabric/agent/services/object/main/log` to find object service logs:

- `authsvc.log`: Records information from the authentication service.
- `blobsvc*.log`: These logs record aspects of the blob service.
- `cassvc*.log`: These logs record aspects of the CAS service.
- `coordinatorsvc.log`: Records information from the coordinator service.
- `ecsportalsvc.log`: Records information from the ECS Portal service.
- `eventsvc*.log`: These logs record aspects of the event service. This information is available in the ECS Portal Monitoring menu.
- `hdfssvc*.log`: These logs record aspects of the HDFS service.
- `objcontrolsvc.log`: Records information from the object service.
- `objheadsvc*.log`: These logs record aspects of the various object heads supported by the object service.
- `provisionsvc*.log`: These logs record aspects of the ECS provisioning service.
- `resourcesvc*.log`: These logs record information related to global resources like namespaces, buckets, object users, and so on.
- `dataheadsvc*.log`: (ECS 2.2 HF1) These logs record the aspects of the object heads supported by the object service, the file service supported by HDFS, and the CAS service.

Note

From ECS 2.2 HF1 `cassvc`, `objheadsvc` and `hdfssvc` services are combined into `dataheadsvc`.
