

Elastic Cloud Storage (ECS)

Version 3.0

Administrator's Guide

302-003-220

06

Copyright © 2013-2017 EMC Corporation All rights reserved.

Published June 2017

Dell 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." DELL 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 DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.
Published in the USA.

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

CONTENTS

Figures		9
Tables		11
Part 1	Use the ECS Portal	13
Chapter 1	Introduction	15
	Introduction to the ECS Portal.....	16
	Log in to the ECS Portal.....	16
	Change Password.....	17
	Access to portal areas.....	18
	Ordering and searching tables in the portal.....	21
	About this VDC.....	22
Chapter 2	Use the Getting Started Checklist	25
	Using the Getting Started Checklist.....	26
Chapter 3	Use the Dashboard	29
	Use the Portal Dashboard.....	30
Part 2	Configure and Manage	35
Chapter 4	Configure One or More Sites	37
	Configure storage pools, VDCs, and replication groups.....	38
	Storage pools.....	38
	Create storage pools.....	39
	Virtual data centers (VDCs).....	41
	Create a VDC for a single site.....	42
	Add a VDC to a federation.....	43
	Fail over a site/Delete a VDC.....	43
	Replication groups.....	44
	Create a replication group.....	45
	Configuring ESRS.....	46
	Verifying call home setup.....	48
Chapter 5	Configure a namespace	51
	Configure a namespace.....	52
	Understanding tenants.....	52
	Understanding namespace settings.....	53
	Working with namespaces at the ECS portal.....	56
	Create and configure a namespace.....	56
Chapter 6	Configure an authentication provider	59
	Configure an authentication provider.....	60

	Working with the authentication providers at the ECS Portal.....	60
	Add an AD or LDAP authentication provider.....	61
	AD/LDAP Authentication provider settings.....	61
	Add a Keystone authentication provider.....	66
	Keystone authentication provider settings.....	66
Chapter 7	Manage users	67
	Manage users and roles.....	68
	Understanding users and roles in ECS.....	68
	Users in ECS.....	68
	User roles.....	69
	Domain and local users.....	71
	User scope: global or namespace.....	71
	Working with the users at the ECS Portal.....	72
	Add a new object user.....	74
	Add a domain user as an object user.....	75
	Create a local management user or assign a domain user to a management role.....	76
	Assign an Active Directory group name to the system admin or system monitor role.....	77
	Create a namespace administrator.....	78
	Assign an Active Directory group name to the namespace admin role.....	78
	Understanding the mapping of users into a namespace.....	79
	Map domain users into a namespace.....	81
Chapter 8	Manage tenants	83
	Manage tenants.....	84
	Quotas.....	84
	Retention periods and policies.....	85
	Lock buckets and users.....	87
	Metering.....	87
	Audit buckets.....	89
Chapter 9	Remove a site	91
	Fail over a site/Delete a VDC.....	92
Chapter 10	Manage licenses	93
	Licensing.....	94
	Obtain the EMC ECS license file.....	94
	Upload the ECS license file.....	94
Chapter 11	Create and manage buckets	97
	Create and manage buckets.....	98
	Bucket concepts.....	98
	Bucket attributes.....	99
	Default Group.....	101
	Metadata index keys.....	102
	Bucket tagging.....	104
	Bucket ACLs.....	104
	Create a bucket using the ECS Portal.....	106
	Edit a bucket.....	107

	Set the bucket ACL permissions for a user.....	108
	Set the bucket ACL permissions for a pre-defined group.....	109
	Set custom group bucket ACLs.....	110
	Create a bucket using the S3 API (with s3curl).....	111
	Bucket HTTP headers.....	114
	Bucket and key naming conventions.....	115
	S3 bucket and object naming in ECS.....	115
	OpenStack Swift container and object naming in ECS.....	116
	Atmos bucket and object naming in ECS.....	116
	CAS pool and object naming in ECS.....	117
Chapter 12	Configure NFS file access	119
	NFS file access.....	120
	Multi-protocol access to directories and files.....	120
	Node and site failure.....	121
	ECS Portal support for NFS configuration.....	121
	Exports.....	122
	User/Group mappings.....	122
	ECS NFS configuration tasks.....	123
	Create a bucket for NFS using the ECS Portal.....	124
	Add an NFS export.....	126
	Add a user or group mapping.....	129
	Configure NFS security with Kerberos.....	130
	Mounting an NFS export : example.....	136
	Best practice when using ECS NFS.....	138
	Permissions for multi-protocol (cross-head) access.....	138
	File API Summary.....	140
Chapter 13	Configure Event Notification servers	143
	Configure Event Notification servers (SNMP or Syslog).....	144
	Working with the SNMP and Syslog servers at the ECS Portal.....	145
	Add an SNMP v2 Trap recipient.....	147
	SNMP v2 server settings.....	148
	Add an SNMP v3 Trap recipient.....	148
	SNMP v3 server settings.....	148
	Add a Syslog server.....	149
	Syslog server settings.....	150
Chapter 14	Set the Base URL	151
	Set the Base URL.....	152
	Bucket addressing.....	152
	DNS Configuration.....	153
	Base URL.....	153
	Add a Base URL.....	155
Chapter 15	Configure certificates	157
	Introduction to certificates.....	158
	Generating certificates.....	158
	Create a private key.....	159
	Generate a SAN configuration.....	159
	Create a self-signed certificate.....	160
	Create a certificate signing request.....	162
	Upload a certificate.....	164

	Authenticate with ECS Management REST API.....	164
	Upload a management certificate.....	164
	Upload a data certificate for data access endpoints.....	166
	Verifying installed certificates.....	167
	Verify the management certificate.....	168
	Verify the object certificate.....	169
Chapter 16	Locking remote access to nodes	171
	Locking remote access to nodes.....	172
	Lock and unlock nodes.....	173
Part 3	Monitor	175
Chapter 17	Monitoring basics	177
	Using monitoring pages.....	178
Chapter 18	Monitor metering	181
	Monitor metering data.....	182
	Metering data.....	183
Chapter 19	Monitor events	185
	About event monitoring.....	186
	Monitor audit data.....	186
	Monitor alerts.....	187
Chapter 20	Monitor capacity utilization	189
	Monitor capacity.....	190
	Storage capacity data.....	190
Chapter 21	Monitor traffic metrics	195
	Monitor network traffic.....	196
Chapter 22	Monitor hardware health	199
	Monitor hardware.....	200
Chapter 23	Monitor node and process health	201
	Monitor node and process health.....	202
Chapter 24	Monitor chunk summary	205
	Monitor chunks.....	206
Chapter 25	Monitor erasure coding	209
	Monitor erasure coding.....	210
Chapter 26	Monitor recovery status	213
	Monitor recovery status.....	214

Chapter 27	Monitor disk bandwidth	217
	Monitor disk bandwidth.....	218
Chapter 28	Monitor geo-replication	221
	Introduction to Geo-replication monitoring.....	222
	Monitor geo-replication: Rate and Chunks.....	222
	Monitor geo-replication: Recovery Point Objective (RPO).....	223
	Monitor geo-replication: Failover Processing.....	224
	Monitor geo replication: Bootstrap Processing.....	225
Chapter 29	Service logs	227
	Service logs.....	228
	ECS service log locations.....	228
Appendix A	Audit and Alert Messages	229
	Audit messages	230
	Alert messages	236
Appendix B	ECS Support for SNMP	241
	SNMP support in ECS.....	242
	SNMP MIBs supported for querying in ECS.....	242
	ECS-MIB SNMP Object ID hierarchy and MIB definition.....	242

CONTENTS

FIGURES

1	ECS login.....	17
2	ECS logout.....	17
3	Table Column with Sort Control Available.....	22
4	A Completed Checklist.....	27
5	Storage Pool Management page.....	38
6	VDC Management page.....	41
7	Manage Replication Groups page.....	44
8	Namespace management page.....	56
9	User mappings for a tenant using AD attributes.....	80
10	Using multiple mapping criteria.....	81
11	Event Notification page in SNMP mode.....	146
12	Event Notification page in Syslog mode.....	146
13	Refresh.....	178
14	Open Filter panel with criteria selected.....	178
15	Closed Fiter panel showing summary of applied filter.....	178
16	Navigating with breadcrumbs.....	179
17	History chart with active cursor.....	180
18	Metering page with criteria selected.....	183
19	Network traffic charts for a VDC.....	197
20	Hardware Health for nodes.....	200
21	Node & Process Health.....	204
22	Chunk Summary.....	207
23	Disk Bandwidth.....	219
24	Geo replication: Rate and Chunks	223
25	RPO.....	223
26	Failover	225
27	Bootstrap processing.....	226

FIGURES

TABLES

1	Storage pool properties.....	38
2	VDC properties.....	41
3	Replication Group properties.....	44
4	Authentication provider settings.....	61
5	Bucket and namespace metering.....	87
6	Bucket attributes.....	100
7	Bucket ACLs.....	105
8	Bucket headers.....	114
9	Bucket and namespace metering.....	183
10	Alert types.....	188
11	Capacity Utilization: Storage Pool.....	190
12	Capacity Utilization: Node.....	191
13	Capacity Utilization: Disk.....	193
14	Network traffic metrics.....	196
15	VDC, node, and process health metrics.....	202
16	Chunk tables.....	206
17	Chunk metrics.....	206
18	Erasur coding metrics.....	210
19	Recovery metrics.....	214
20	Disk bandwidth metrics.....	218
21	Rate and Chunk columns.....	222
22	RPO columns.....	223
23	Failover columns.....	224
24	Bootstrap Processing columns.....	225
25	ECS audit messages.....	230
26	ECS Object alert messages.....	236
27	ECS Fabric alert messages.....	237

TABLES

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](#)..... 16
- [Log in to the ECS Portal](#)..... 16
- [Change Password](#)..... 17
- [Access to portal areas](#)..... 18
- [Ordering and searching tables in the portal](#).....21
- [About this VDC](#)..... 22

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. That is, users assigned to the System Admin, System Monitor, and Namespace Admin roles. 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 [Manage users and roles](#) on page 68.

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

Log in to the ECS Portal

Log in to the ECS Portal from the 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

- If you are assigned to the System Admin, System Monitor, Lock Admin (emcsecurity), or Namespace Admin role, you can log in to the portal.
- A root user account, which is assigned to the System Admin role, is provided for initial access. Note that this root account is not related to node-level Linux accounts.

After the initial login with the root credentials (root/ChangeMe), you are prompted to change the password for the root account immediately.

The session ends when you close the browser, or log out. Logging out always closes the session. If you are unable to log in, contact the 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 in the following form:
`http://<node1_public_ip>`.
2. After changing the password at first login, click **Save**.

You are logged out and the standard login screen appears.

Figure 1 ECS login

Elastic Cloud Storage

User Name *

Password *


Login

EMC

© 2015 EMC Corporation. All Rights Reserved.

3. Type 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**.

Figure 2 ECS logout



Change Password

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

System Admin, Namespace Admin, Lock Admin (emcsecurity), 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 following sections detail the permissions provided for different management users.

- [System Admin and System Monitor](#) on page 18
- [Namespace Admin](#) on page 20

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 182.
	Events	View audit and alert events. For more information, see: About event monitoring on page 186.
	Capacity Utilization	Monitor storage pool, node, and disk capacity. For more information, see: Monitor capacity on page 190.
	Traffic Metrics	Monitor read and write bandwidth and latency. For more information, see: Monitor network traffic on page 196.
	Hardware Health	Monitor storage node and disk status for each storage pool. For more information, see: Monitor hardware on page 200.
	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 202.
	Chunk Summary	Monitor chunks and chunks status. For more information, see: Monitor chunks on page 206.
	Erasur Coding	Monitor erasure coding status. For more information, see: Monitor erasure coding on page 210.

Area	Menu	Operations Supported
	Recovery Status	Monitor recovery status. For more information, see: Monitor recovery status on page 214.
	Disk Bandwidth	Monitor disk bandwidth usage. For more information, see: Monitor disk bandwidth on page 218.
	Geo Replication	Monitor geo-replication activity. For more information, see: Introduction to Geo-replication monitoring on page 222.
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 .
	Virtual Data Center	
	Replication Group	
	Authentication	Add an authentication provider that can authenticate domain users. See Manage users and roles on page 68.
	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. For more information, see: Configure a namespace for a tenant
	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: Manage users and roles on page 68
	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: Bucket concepts on page 98
File	Enables buckets to be accessed as NFS filesystems. For more information, see: NFS File Access	
Settings	Object Base URL	Set the Base URL to determine which part of object address is the bucket and namespace.

Area	Menu	Operations Supported
		For more information, see: Address ECS object storage and use the Base URL
	Change Password	Change own password. For more information, see: Change Password on page 17
	ESRS	Configure sending of alerts to EMC. For more information, see: Configuring ESRS on page 46.
	Event Notification	Configure or view SNMP and Syslog servers. For more information, see: Configure Event Notification servers (SNMP or Syslog) on page 144
	Platform Locking	View the lock status of nodes. For more information, see: Locking remote access to nodes on page 172
	Licensing	View license status and upload a license. For more information, see: Obtain and upload a license file to the ECS Portal
	About this VDC	View information about the VDC's nodes: node names, rack IDs, and software versions. For more information, see About this VDC .

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 182.
Manage	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 Manage users and roles on page 68
	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 Bucket concepts on page 98

Area	Menu	Operations Supported
Settings	Change Password	Change own password. For more information, see Change Password on page 17

Lock Admin

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

Area	Menu	Operations Supported
Settings	Change Password	Change own password. For more information, see Change Password on page 17
	Platform Locking	Change the lock status of nodes. For more information, see: Locking remote access to nodes on page 172

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 3 Table Column with Sort Control Available

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](#).

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**.

The About this VDC dialog appears with the About tab open.


About this VDC 

About Nodes

EMC Elastic Cloud Storage
v3.0.0.0

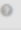
 Logged in to Node: 10.245.137.85
(rackid: strawberry)

ECS Object Version 3.0.0.0.85309.47fbd8e

 Click the Nodes tab at the top of this page for more information about all the nodes in this VDC.


Here you find information about the ECS software version and ECS Object service version for the current node.

2. Select **Nodes** to see the software version for all reachable nodes in the cluster.

About this VDC 

About Nodes

Current Node	Version	Rack ID
10.245.137.85	3.0.0.0	strawberry

Node	Node IP Address	ECS Object Version	Rack ID
 layton-strawberry.ecs.lab.emc.com	10.245.137.85	3.0.0.0.85309.47fbd8e	strawberry
lehi-strawberry.ecs.lab.emc.com	10.245.137.87	3.0.0.0.85309.47fbd8e	strawberry
logan-strawberry.ecs.lab.emc.com	10.245.137.86	3.0.0.0.85309.47fbd8e	strawberry
murray-strawberry.ecs.lab.emc.com	10.245.137.88	3.0.0.0.85309.47fbd8e	strawberry

The blue checkmark indicates the current node. Nodes with a star are nodes that have a different software version.

CHAPTER 2

Use the Getting Started Checklist

- [Using the Getting Started Checklist](#).....26

Using the Getting Started Checklist

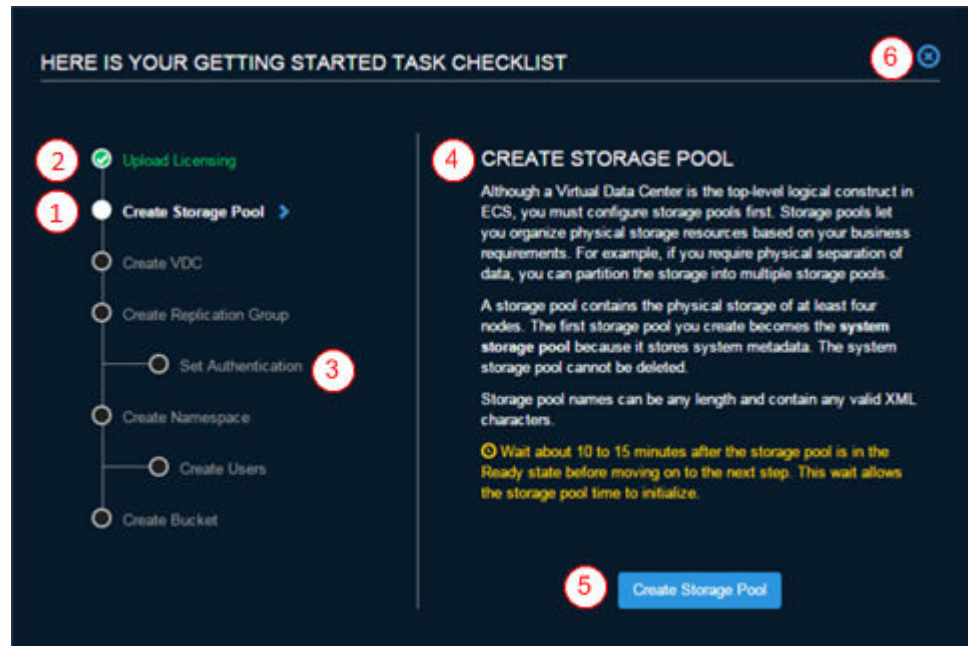
Use the Getting Started Checklist to guide you through the initial configuration of your ECS site.

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

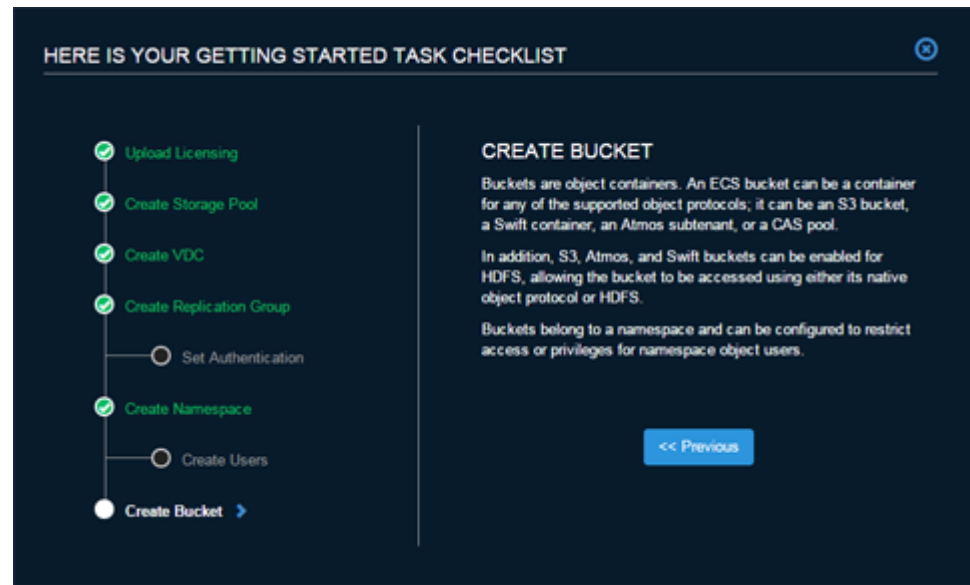
Some parts of the initial configuration may be completed as part of the installation service.



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 4 A Completed Checklist



Use the Getting Started Checklist

CHAPTER 3

Use the Dashboard

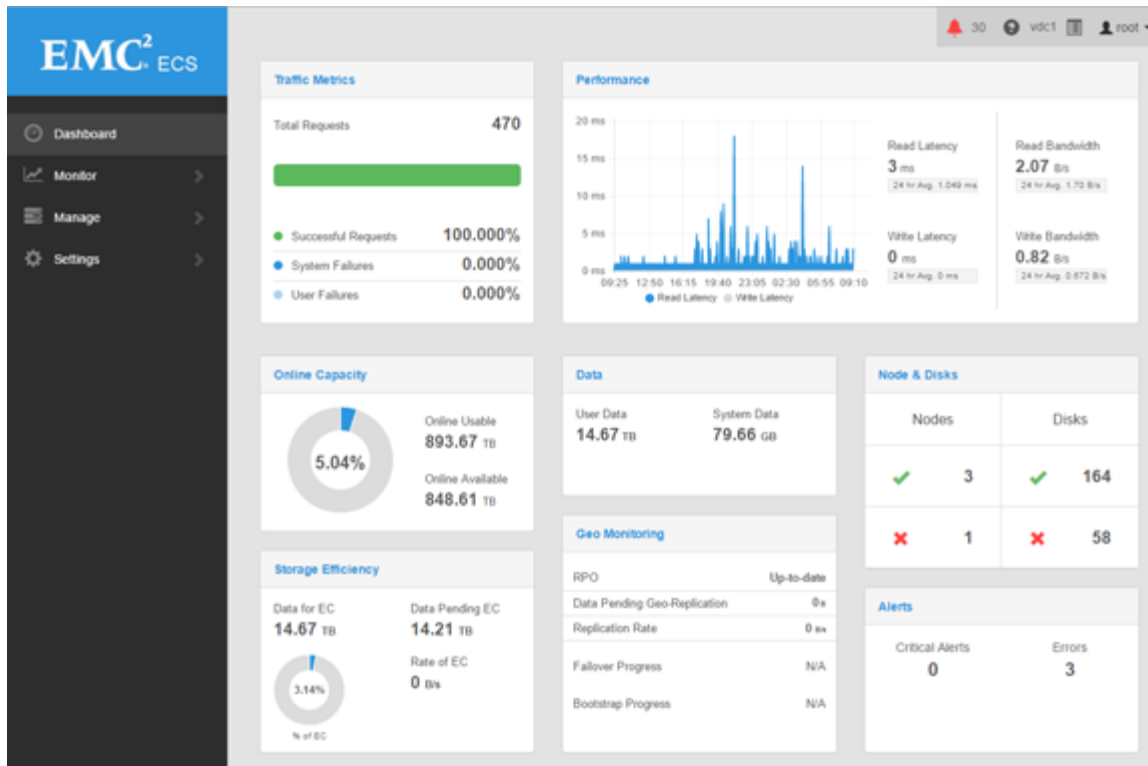
- [Use the Portal Dashboard](#)..... 30

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:

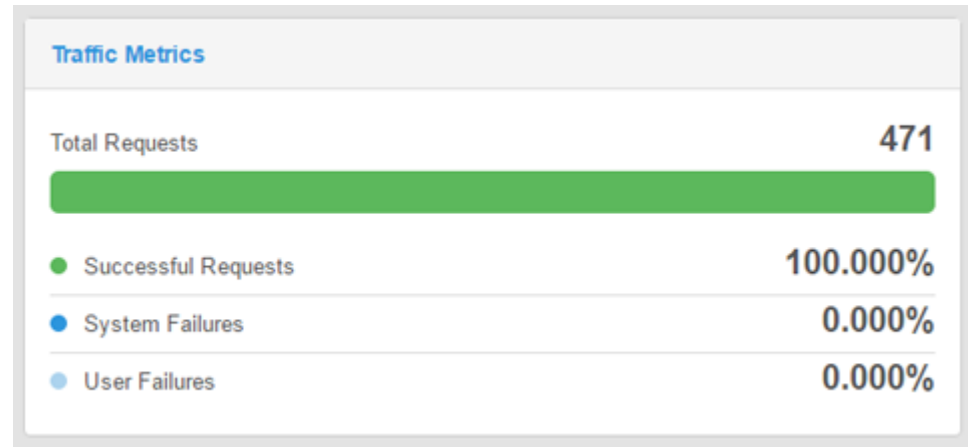
1. The Alert menu shows the most recent five alerts for the current VDC. The number indicates how many unacknowledged alerts are pending for the current VDC. The number displays "99+" if there are more than 99.
2. The Global Help icon brings up the online documentation for the current portal page.
3. 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 will 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.

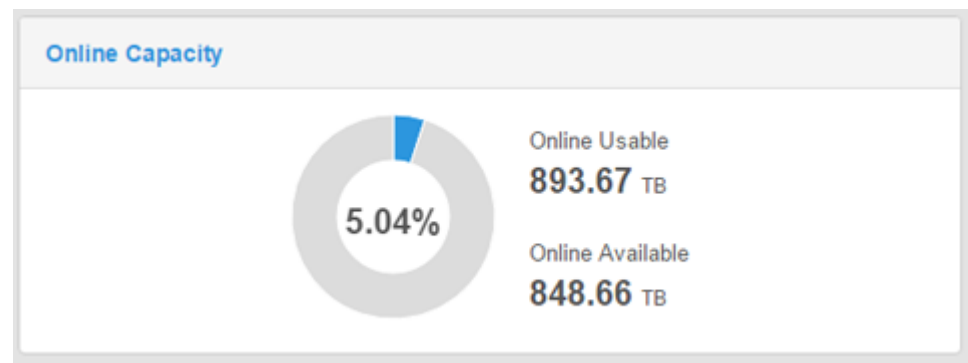
Traffic Metrics

The Traffic Metrics displays total requests and breaks that down into successful requests and failed requests by user error and failed requests by system error. Click the title to see more comprehensive traffic metrics.



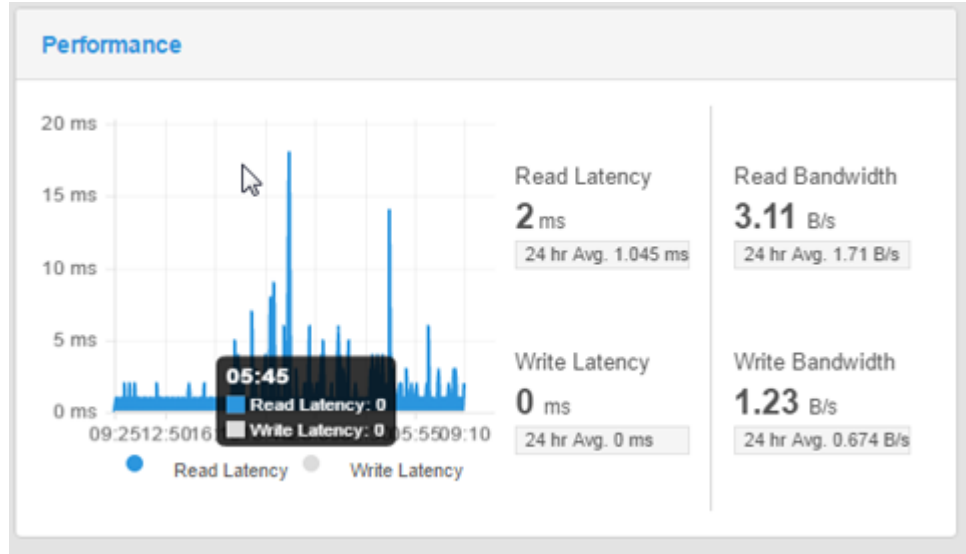
Online Capacity

The Online Capacity panel displays the online usable capacity (total capacity online) and online available capacity (total available for immediate use). The percentage shown in the center of the graph shows how much of the usable capacity is currently in use. Capacity takes into account ingested data, replicas, and system data. Click the title to see more comprehensive capacity metrics.




Performance

The Performance panel shows you how network read and write operations are performing now and the average over the last 24 hours. Click the title to see more comprehensive performance metrics.



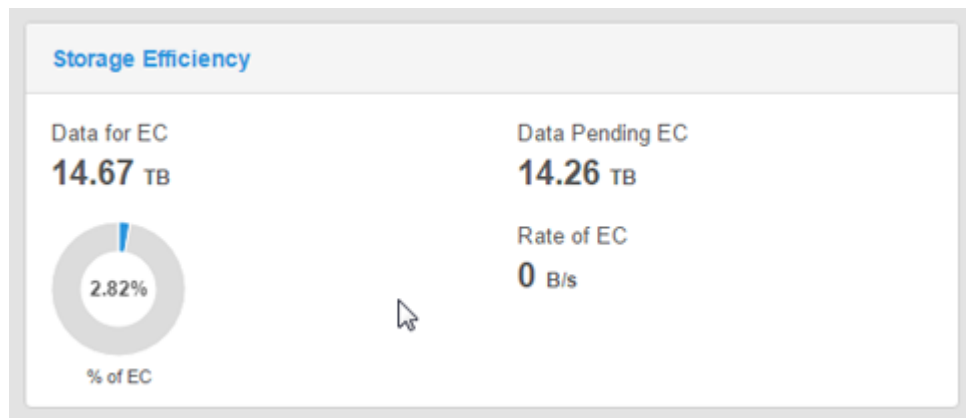
Data

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. Click the title to see more comprehensive data metrics.



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. Click the title to see more comprehensive storage efficiency metrics.



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. Click the title to see more comprehensive geo metrics.

Geo Monitoring	
RPO	Up-to-date
Data Pending Geo-Replication	0 B
Replication Rate	0 B/h
Failover Progress	N/A
Bootstrap Progress	N/A

Nodes and Disks

The Nodes and Disks panel shows the health status (Good or Bad) of disks and nodes. Click the title to see more comprehensive hardware metrics.

Node & Disks			
Nodes		Disks	
	3		164
	1		58

Alerts

The Alert panel displays a count of critical alerts and errors. Click the title 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 an authentication provider"
- Chapter 7, "Manage users"
- Chapter 8, "Manage tenants"
- Chapter 9, "Remove a site"
- Chapter 10, "Manage licenses"
- Chapter 11, "Create and manage buckets"
- Chapter 12, "Configure NFS file access"
- Chapter 13, "Configure Event Notification servers"
- Chapter 14, "Set the Base URL"
- Chapter 15, "Configure certificates"
- Chapter 16, "Locking remote access to nodes"

CHAPTER 4

Configure One or More Sites

- [Configure storage pools, VDCs, and replication groups](#)..... 38
- [Storage pools](#)..... 38
- [Virtual data centers \(VDCs\)](#)..... 41
- [Replication groups](#)..... 44
- [Configuring ESRS](#)..... 46

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 ESRS 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 5 Storage Pool Management page

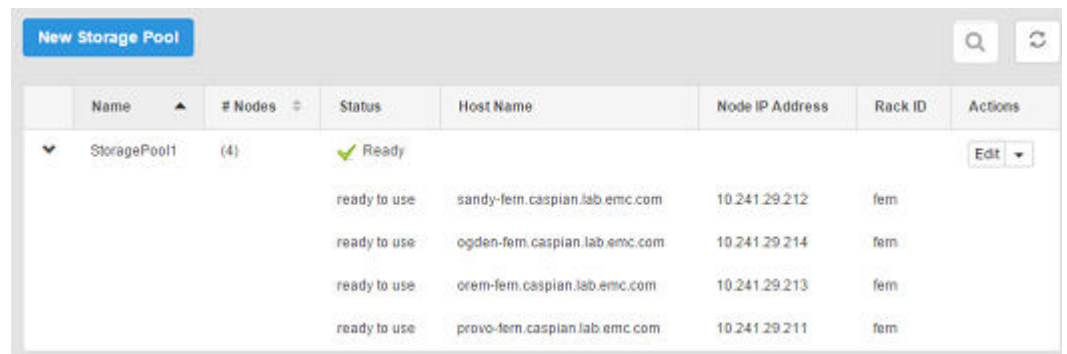


Table 1 Storage pool properties

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 <code>ready to use</code> state. Not Ready: A node in the storage pool is not in the <code>ready to use</code>. Partially Ready: There are less than four nodes and all nodes are in the <code>ready to use</code> 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.

Table 1 Storage pool properties (continued)

Field	Description
Actions	<p>Actions are:</p> <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	<p>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.</p>

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**.

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 section 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 two sites named vdc1 and vdc2.

Figure 6 VDC Management page

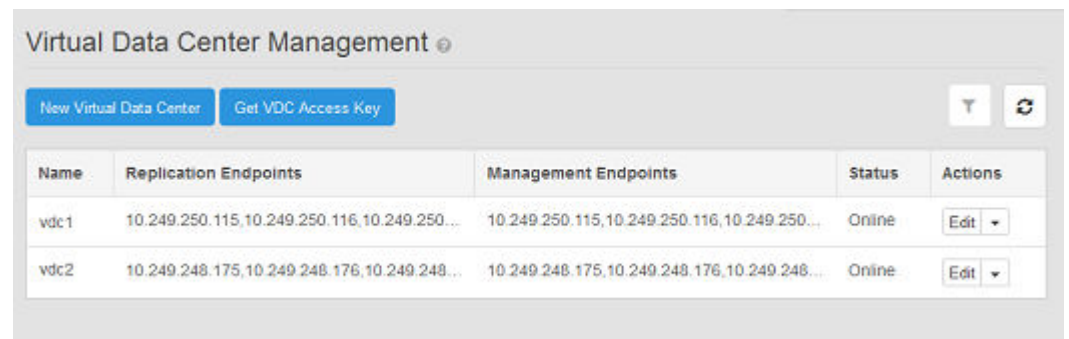


Table 2 VDC properties

Field	Description
Name	The VDC's name.
Replication Endpoints	Endpoints for communication of replication data between sites. If a separate replication network has been configured, this will be a list comprising each node's replication IP address. If no separate replication network has been configured, it will be a list comprising each node's public IP address. If neither replication or management networks have been separated, replication endpoints and management endpoints will be the same.
Management Endpoints	Endpoints for communication of management commands between sites. If a separate management network has been configured, this will be a list comprising each node's management IP address. If no separate management network has been configured, it will be a list comprising each node's public IP address. If neither management or networks have been separated, replication endpoints and management endpoints will be the same.

Table 2 VDC properties (continued)

Field	Description
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.

VDC names can be from 1 to 255 characters. Valid characters include a to z, A to Z, 0 to 9 and dash (-) and underscore (_).

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 **Replication Endpoints** text box, enter the replication IP address of each node in the VDC's storage pools. Supply them as a comma-separated list.

If network separation has been configured at installation, this will be a list comprising each node's replication IP address. If the replication network has not been separated, this will be a list comprising each node's public IP address.

6. In the **Management Endpoints** text box, enter the management IP address of the each node in the VDC's storage pools. Supply them as a comma-separated list.

If network separation has been configured at installation, this should be a list comprising each node's management address. If the management network has not been separated, this will be a list comprising each node's public IP address.

7. 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 Admin 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), or, if you have separated the management and/or replication networks, a list of the management and/or replication addresses.

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 Replication Endpoints and Management Endpoints for the nodes that comprise the site. Enter the IP addresses as comma-separated lists.

If you have not separated the replication or management networks, both of these fields will contain the same list of the public IP addresses for the nodes. If either or both the management and/or replication networks have been separated, you should list the IP addresses for the appropriate network.

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.

Note

Replication groups cannot be modified once created. If you want a replication group to replicate data across sites, you must ensure that the sites/VDCs have been federated before you attempt to create the replication group.

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 7 Manage Replication Groups page




Table 3 Replication Group properties

Field	Description
Name	The replication group name.
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

Table 3 Replication Group properties (continued)

Field	Description
	<ul style="list-style-type: none"> 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 a replication group

Use this procedure to create a replication group. Replication groups can be local to a VDC or can protect data by replicating it across sites.

Before you begin

If you want the replication group to span multiple VDCs, you must ensure that the sites have been installed and that they have been initialized with a VDC identity and a storage pool that can be part of the replication group.

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 EMC Secure Remote Support (ESRS) configuration on ECS. ECS version 2.2 and later requires ESRS Virtual Edition.

Note

This task must run once per site (VDC) in the deployment; preferably on rack 1 node 1 of the site.

Procedure

1. Perform install/upgrade procedures through post-upgrade/install manual configurations.
2. Use an editor to create customer information in a JSON file.

- a. Create a text file for the JSON-formatted customer record:

In this example, a file that is called `customer.json` is created using the `vi` editor.

```
vi /var/tmp/customer.json
```

- b. Add the customer data to the JSON file.

For example:

```
{ "customer_data":
  { "serial": "ABCDE00009",
    "customer_name": "ABCQE_MelonA",
    "customer_email": "John.Smith@xyz.com" }
}
```

- c. Change to the `cli` directory.

```
cd /opt/emc/caspian/fabric/cli
```

- d. Validate the format of the JSON file using the `mjson.tool` Python tool.

```
provo-melon:/opt/emc/caspian/fabric/cli # cat /var/tmp/
customer.json | python -mjson.tool
```

Example output:

```
{
  "status": "OK",
  "etag": 90,
  "customer_data": {
    "serial": "ABCDE00009",
    "customer_name": "ABCQE_MelonA",
    "customer_email": "John.Smith@xyz.com"
  }
}
```

3. Run the `fcli` to configure customer information and serial number on the cluster.

- a. Set the customer data using the contents of the JSON file.

```
provo-melon:/opt/emc/caspian/fabric/cli # cat /var/tmp/customer.json | bin/fcli lifecycle cluster.setcustomer --body
```

Example output:

```
{
  "status": "OK",
  "etag": 90
}
```

- b. Check the customer data has been set.

```
provo-melon:/opt/emc/caspian/fabric/cli # bin/fcli lifecycle cluster.customer
```

Example output:

```
{
  "status": "OK",
  "etag": 90,
  "customer_data": {
    "serial": "ABCD00009",
    "customer_name": "ABCQE_MelonA",
    "customer_email": "John.Smith@xyz.com"
  }
}
```

- c. Enable call home.

```
provo-melon:/opt/emc/caspian/fabric/cli # bin/fcli lifecycle alert.callhomeenabled
```

Example output:

```
{
  "status": "OK",
  "etag": 90,
  "callhome_enabled": true
}
```

4. 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.

Verifying call home setup

Verify that you set up the ESRS call home successfully.

You can test whether call home is working by generating a test alert and then checking whether the alert is received.

Procedure

1. To generate a test alert.
 - a. Authenticate with the ECS Management REST API

For example, using curl:

```
curl -L --location-trusted -k https://192.0.2.49:4443/login -u "root:ChangeMe" -v
```

The X-SDS-AUTH-TOKEN: obtained can be used to authenticate with ECS to run ECS Management REST API commands. It looks like this:

```
X-SDS-AUTH-TOKEN:
BAAcaGtktzRZU2k0SE5acVYxdFNCYU1Uby9mOVM4PQMAjAQASHVybjpgzdG9yYWdlb3M6VmlydHVhbERhdGF
DZW50ZXJEYXRhOjMwOTFjMDY1LTgzMDAtNGN1Ni1iNDY3LTU5NDFiM2MyYTBMzAIADTE0NzM3NzkxMzA4OT
MDAC51cm46VG9rZW46M2Y4NTMzMzctN2ZiZi00NWRiLTk0M2YtY2NkYTc2OWQ0ZTc4AgAC0A8=
```

- b. You can save some typing later by storing the authentication token (X-SDS-AUTH-TOKEN) in an environment variable.

```
export AUTH_TOKEN="X-SDS-AUTH-TOKEN:
BAAcaGtktzRZU2k0SE5acVYxdFNCYU1Uby9mOVM4PQMAjAQASHVybjpgzdG9yYWdlb3M6VmlydHVhbERhdGF
DZW50ZXJEYXRhOjMwOTFjMDY1LTgzMDAtNGN1Ni1iNDY3LTU5NDFiM2MyYTBMzAIADTE0NzM3NzkxMzA4OT
MDAC51cm46VG9rZW46M2Y4NTMzMzctN2ZiZi00NWRiLTk0M2YtY2NkYTc2OWQ0ZTc4AgAC0A8="
```

- c. Generate a test alert.

For example (using the \$AUTH_TOKEN environment variable):

```
curl -ks -H "$AUTH_TOKEN" -H "Content-Type: application/json" -d '{"user_str":
"test alert > for ESRS", "contact": "test_user@emc.com"}' https://
10.241.207.57:4443/vdc/callhome/alert | xmllint -format -
```

The `user_str` parameter enables you to specify a test message, and the `contact` parameter enables you to supply an email address.

2. At the ECS Portal, check that the ESRS notification has been received.
3. Check that the latest test alert is present.
 - a. SSH into the ESRS server.
 - b. Go to the location of the RSC file.

```
cd /opt/connectemc/archive/
```


c. Check for the latest RSC file, using:

```
ls -lrt RSC_<SERIAL NUMBER>*
```

d. Open the file and check if the latest test alert is present in the description.

CHAPTER 5

Configure a namespace

- [Configure a namespace](#)..... 52
- [Understanding tenants](#)..... 52
- [Understanding namespace settings](#)..... 53
- [Working with namespaces at the ECS portal](#)..... 56
- [Create and configure a namespace](#)..... 56

Configure a namespace

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 52
- [Understanding namespace settings](#) on page 53
- [Working with namespaces at the ECS portal](#) on page 56

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

- [Create and configure a namespace](#) on page 56

While the configuration operations described in this article use the ECS portal, the concepts described in [Understanding tenants](#) on page 52 and [Understanding namespace settings](#) on page 53 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](#).

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](#).

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 Manage users and roles on page 68 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 Manage users and roles on page 68 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	Yes

Field	Description	Can be Edited
	can be defined to notify that a defined limit has been reached and to 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 buckets 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)	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. Compliance rules include: <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. 	No

Field	Description	Can be Edited
	<ul style="list-style-type: none"> Using the Infinite option on a bucket mean objects in the bucket in a Compliance-enabled namespace can never be deleted. 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 polices, 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 85.</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 79</p>	Yes

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

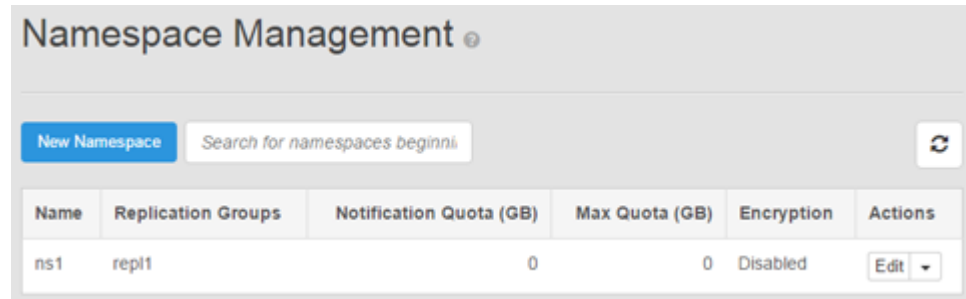
Allowed (and Disallowed) Replication Groups

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 8 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 [Manage users and roles](#) on page 68 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 53.

Procedure

1. At the ECS portal, select **Manage > Namespace**
2. 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 interface. It includes the following fields and controls:

- Name**: A text input field with a required asterisk and a help icon.
- User Admin**: A text input field with a help icon and an example: `Ex.: user1, user2@foo.com`.
- Domain Group Admin**: A text input field with a help icon and an example: `Ex.: group1@foo.com, group2@test.com`.
- Replication Group**: A dropdown menu with a required asterisk and a help icon, currently set to 'repl1'.
- Namespace Quota**: A toggle switch with 'Disabled' and 'Enabled' options.
- Default Bucket Quota**: A toggle switch with 'Disabled' and 'Enabled' options.
- Server-side Encryption**: A toggle switch with 'Disabled' and 'Enabled' options.
- Access During Outage**: A toggle switch with 'Disabled' and 'Enabled' options.
- Compliance**: A toggle switch with 'Disabled' and 'Enabled' options.
- Retention Policies**: A table with columns for Name, Value, and Actions. An 'Add' button is located to the right of the table.
- Domain**: A blue button with a plus sign and a help icon.
- Save** and **Cancel**: Buttons at the bottom of the form.

3. Enter a name for the namespace.
The name must contain lowercase characters only.
4. 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.
5. 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
6. 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.
 7. 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.
 8. 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.
 9. 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 [Manage users and roles](#) on page 68.
 10. Select **Save**.

CHAPTER 6

Configure an authentication provider

- [Configure an authentication provider](#) 60
- [Working with the authentication providers at the ECS Portal](#)..... 60
- [Add an AD or LDAP authentication provider](#) 61
- [Add a Keystone authentication provider](#)66

Configure an authentication provider

Authentication providers can be added to ECS to enable users to be authenticated by systems external to ECS.

Authentication provider is the ECS term for the a system external to ECS that can authenticate users on behalf of ECS. ECS needs to store information that will allow it to connect to the authentication provider so that it can request authentication of a user.

In ECS, there are currently two main types of authentication provider:

AD/LDAP

Used to authenticate domain users that are assigned to management roles in ECS.

Keystone

Used to authenticate OpenStack Swift object users.

Authentication providers can be created from the ECS Portal (see [Working with the authentication providers at the ECS Portal](#) on page 60) or using the ECS Management REST API or CLI. You can follow the procedures below to create AD/LDAP or Keystone authentication providers.

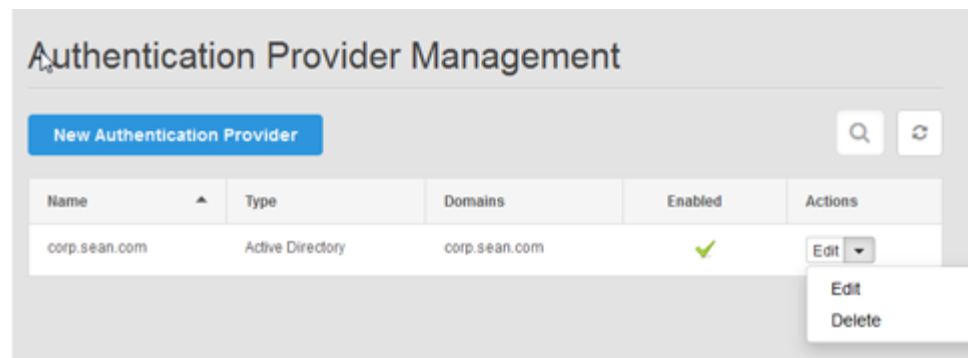
- [Add an AD or LDAP authentication provider](#) on page 61
- [Add a Keystone authentication provider](#) on page 66

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.

Attribute	Description
Type	Indicates whether the authentication provider is an Active Directory (AD), Lightweight Directory Access Protocol (LDAP), or Keystone V3 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 AD or LDAP authentication provider

User authentication for ECS domain users is performed using one or more authentication providers added 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](#). Note especially the requirements for the Manager DN user.

Procedure

- At the ECS Portal, select **Manage > Authentication > New Authentication Providers**.
- Enter values for the attributes. Refer to [Authentication provider settings](#)
- 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.

AD/LDAP 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.

Table 4 Authentication provider settings (continued)

Field name	Description and requirements
<p>Domains</p>	<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> <div style="background-color: #f0f0f0; padding: 5px; margin: 10px 0;"> <p>mycompany.com</p> </div> <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>
<p>Server URLs</p>	<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>
<p>Manager DN</p>	<p>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.</p> <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</p>

Table 4 Authentication provider settings (continued)

Field name	Description and requirements
	<p>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>
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.

Table 4 Authentication provider settings (continued)

Field name	Description and requirements
	<ul style="list-style-type: none"> 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 namespace 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> <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</p>

Table 4 Authentication provider settings (continued)

Field name	Description and requirements
	reverse of the structure seen in the Active Directory Users and Computers UI.
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.

Add a Keystone authentication provider

You can add a Keystone authentication provider that will authenticate OpenStack Swift users.

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 [Keystone authentication provider settings](#) on page 66.

Procedure

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

Keystone authentication provider settings

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

Field	Description
Name	The name of the Keystone authentication provider. This name is used to identify the provider in ECS.
Description	Free text description of the authentication provider.
Type	Keystone V3
Server URL	URI of the Keystone system that ECS will connect to obtain authentication for Swift users.
Keystone Administrator	User name for an administrator of the Keystone system. ECS will connect to the Keystone system using this username.
Admin Password	Password of the specified Keystone administrator.

CHAPTER 7

Manage users

- [Manage users and roles](#).....68
- [Understanding users and roles in ECS](#).....68
- [Working with the users at the ECS Portal](#)..... 72
- [Understanding the mapping of users into a namespace](#).....79

Manage users and roles

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 68
- [Working with the users at the ECS Portal](#) on page 72

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

- [Add a new object user](#) on page 74
- [Add a domain user as an object user](#) on page 75
- [Create a local management user or assign a domain user to a management role](#) on page 76
- [Create a namespace administrator](#) on page 78

In addition, it shows you how you perform the mapping of domain users into a namespace:

- [Map domain users into a namespace](#) on page 81

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 68
- [User roles](#) on page 69
- [Domain and local users](#) on page 71
- [User scope: global or namespace](#) on page 71

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 71.

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

- [Management Users](#) on page 69
- [Object users](#) on page 69
- [Root user](#) on page 69

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 71.

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 71, and you can refer to [ECS 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 70
- [System Monitor](#)
- [Namespace Admin](#) on page 70

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 72
- [Management Users View](#) on page 73

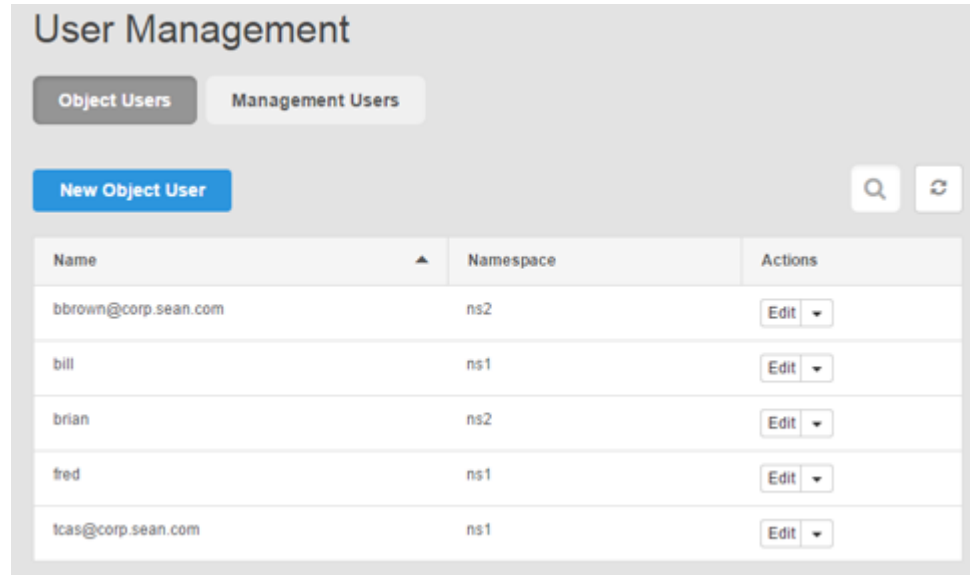
The Management Page is only accessible if you are a System Admin (or root user) for ECS.

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 .

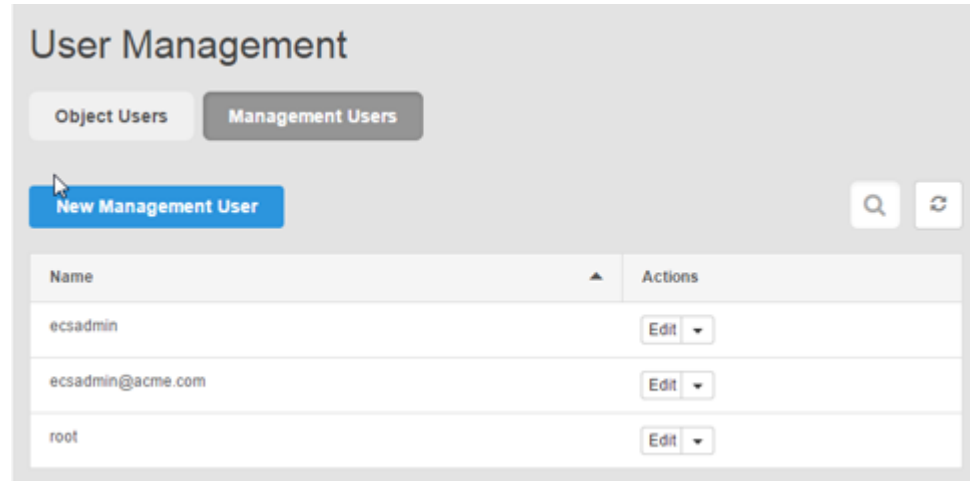
The Object Users pane additionally provides access to the the following controls:

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 75.
- 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 72 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 AD or LDAP authentication provider](#) on page 61.

2. Map domain users into the namespace as described in [Map domain users into a namespace](#) on page 81.

This can be performed by the Namespace Admin.

3. Allow users to create secret keys using the instructions in [ECS 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 AD or LDAP authentication provider](#) on page 61.
- 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](#)). 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 72 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 AD or LDAP authentication provider](#) on page 61.

You can refer to [Working with the users at the ECS Portal](#) on page 72 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 72 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 76.

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](#).

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 AD or LDAP authentication provider](#) on page 61.

You can refer to [Working with the users at the ECS Portal](#) on page 72 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**.

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 9 User mappings for a tenant using AD attributes

The screenshot shows a configuration interface for user mappings. At the top, there are two sections: 'Domain' and 'Groups'. Under 'Domain', there is a 'Domain' label with a trash icon, a text input field containing 'corp.sean.com', and a 'Groups' label with a trash icon and a text input field containing 'Comma separated groups'. Below these, there are two more input fields: 'Attribute' with the value 'Department' and 'Values' with the value 'Accounts'. On the left side, there is an 'Attribute' label with a trash icon. At the bottom left, there is a blue button with a plus sign and the text 'Domain'. In the center, there is a blue button with a plus sign and the text 'Attribute'.

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 10 Using multiple mapping criteria

The screenshot displays the ECS portal interface for mapping domain users. It features a central configuration area with the following elements:

- Domain:** A dropdown menu showing "corp.sean.com".
- Groups:** A dropdown menu showing "Storage Admins".
- Attribute/Value Pairs:** Two pairs of input fields are shown. The first pair has "Department" as the attribute and "Accounts" as the value. The second pair has "Company" as the attribute and "Acme" as the value. A red double-headed arrow labeled "AND" connects these two pairs, indicating that both criteria must be met.
- Logical Operator:** A red double-headed arrow labeled "OR" is positioned between the two attribute/Value pairs, indicating that either pair can be used for mapping.
- Buttons:** There are blue buttons with a plus sign and the text "Attribute" and "Domain" for adding new 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 79.

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 8

Manage tenants

- [Manage tenants](#)..... 84
- [Quotas](#)..... 84
- [Retention periods and policies](#)..... 85
- [Lock buckets and users](#)..... 87
- [Metering](#)..... 87
- [Audit buckets](#)..... 89

Manage tenants

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 [Manage users and roles](#) on page 68.

Quotas

The ability to set quotas on namespaces and buckets is described in [Quotas](#) on page 84.

Retention Periods

The ability to create retention policies is described in [Retention periods and policies](#) on page 85.

Lock buckets and users

The ability to lock buckets and users is described in [Lock buckets and users](#) on page 87.

Metering

The ability to meter the writing of data to buckets and namespaces is described in [Metering](#) on page 87.

Audit buckets

The ability to audit the operations associated with buckets is described in [Audit buckets](#) on page 89.

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](#).

Quotas for a bucket are set from the **Manage > Bucket** page, as described in [Bucket concepts](#) on page 98 .

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: [ECS Data Access Guide: 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.

Note

For information about CAS retention and CAS advanced retention, see [ECS Data Access Guide: CAS](#).

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](#)

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: [ECS Data Access Guide: 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: [ECS Data Access Guide: 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.

Table 5 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.

Refer to [Monitor metering data](#) on page 182 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: [ECS Data Access Guide: 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 ECS 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 [About event monitoring](#) on page 186 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: [ECS Data Access Guide: ECS Management REST API](#) and the online reference is [here](#).

CHAPTER 9

Remove a site

- [Fail over a site/Delete a VDC.....](#) 92

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 10

Manage licenses

- [Licensing](#)..... 94
- [Obtain the EMC ECS license file](#)..... 94
- [Upload the ECS license file](#)..... 94

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 can obtain a license file (.lic) from the EMC license management web site.

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. Got to the license page at: <https://support.emc.com/servicecenter/license/>
2. Select **ECS Appliance** from the list of products.
3. On the LAC Request page, enter the LAC code and **Activate**.
4. Select the entitlements to activate and **Start Activation Process**.
5. 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.

6. 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.

7. Optionally specify an addressee to receive an email summary of the activation transaction.
8. Click **Finish**.
9. 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 .

Upload the ECS license file

The ECS license file can be uploaded from the ECS Portal.

Before you begin

- Ensure you have a valid license file. You can follow the instructions provided in [Obtain the EMC ECS license file](#) on page 94 to obtain a license.
- Where you are installing more than one site in a geo configuration, the licensing scheme across sites should be the same. If the existing cluster has an encryption-enabled license, any new site added to it should have the same. Similarly, if

existing sites don't have licenses that are encryption enabled, the new sites that are added to the cluster should follow the same model.

- In a geo configuration, if the license is updated on 1 site in the cluster, it should be updated on all sites.

Procedure

1. In the ECS Portal, select **Settings > Licensing**.
2. On the **Licensing** page, at the Upload control, select **Browse** and locate the license file on your system.
3. Select **Upload**.

The license display updates.

CHAPTER 11

Create and manage buckets

- [Create and manage buckets](#).....98
- [Bucket concepts](#)..... 98
- [Bucket attributes](#)..... 99
- [Bucket ACLs](#)..... 104
- [Create a bucket using the ECS Portal](#).....106
- [Edit a bucket](#).....107
- [Set the bucket ACL permissions for a user](#) 108
- [Set the bucket ACL permissions for a pre-defined group](#)..... 109
- [Set custom group bucket ACLs](#)..... 110
- [Create a bucket using the S3 API \(with s3curl\)](#)..... 111
- [Bucket and key naming conventions](#).....115

Create and manage buckets

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 NFS and HDFS), and a bucket configured for file system access can be read and written using its object protocol and using the NFS or 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/NFS, 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/NAS 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.

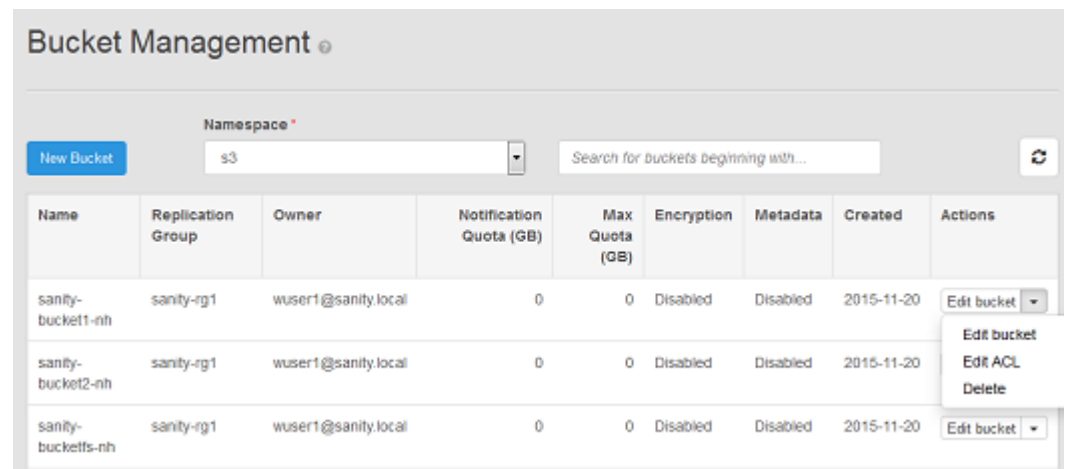
Because there is a risk that object data retrieved during a temporary site outage is not the most recent, the user must indicate that they are prepared to accept this by marking the bucket as available during an outage.

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.



The screenshot shows the 'Bucket Management' interface. At the top, there is a 'Namespace' dropdown menu set to 's3', a 'New Bucket' button, and a search bar. Below this is a table with the following columns: Name, Replication Group, Owner, Notification Quota (GB), Max Quota (GB), Encryption, Metadata, Created, and Actions. The table contains three rows of bucket data. The 'Actions' column for each row has a dropdown menu with options: 'Edit bucket', 'Edit ACL', and 'Delete'.

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 115.	No
Namespace	Namespace with which the bucket is associated.	No
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 104.	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.</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 101</p>	No

Table 6 Bucket attributes (continued)

Attribute	Description	Can be Edited
CAS	Indicates that the bucket is enabled for CAS data.	
Metadata Search	<p>Indicates that metadata search indexes will be created for the bucket based on specified key values. 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 102.</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>	No
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 85.</p>	Yes

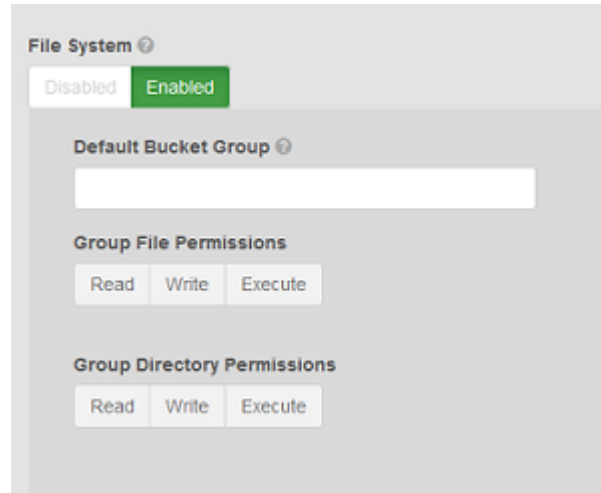
Default Group

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. When a Metadata Key Type of System is selected (see below), metadata that is automatically assigned to objects in a bucket is listed for selection in the Key Names menu.

Add Metadata Search Key [X]

Metadata Key Type
System [v] * Metadata keys cannot be modified after bucket creation

Key Name
CreateTime [v]
CreateTime
LastModified
ObjectName
Owner
Size

* Maximum keys permitted: 30

Add Cancel

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.

Add Metadata Search Key [X]

Metadata Key Type
User [v] * Metadata keys cannot be modified after bucket creation

Key Name
x-amz-meta-

* S3 Key Name prefix: x-amz-meta-

* Maximum keys permitted: 30

Data Type
string [v]
datetime
decimal
Integer
string

Add Cancel

You can read more about metadata search feature in [ECS 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.

The screenshot displays the 'Bucket Tagging' interface. At the top, there is a table with columns 'Key', 'Value', and 'Actions'. Below the table, several settings are listed: 'Quota' (Disabled), 'Server-side' (Disabled), 'File System' (Disabled), 'CAS' (Disabled/Enabled), and 'Metadata Search' (Disabled/Enabled). An 'Add Bucket Tag' modal is open, showing input fields for 'Key' and 'Value', and 'Add' and 'Cancel' buttons. A note at the bottom states: '* Cannot be enabled after bucket creation.'

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

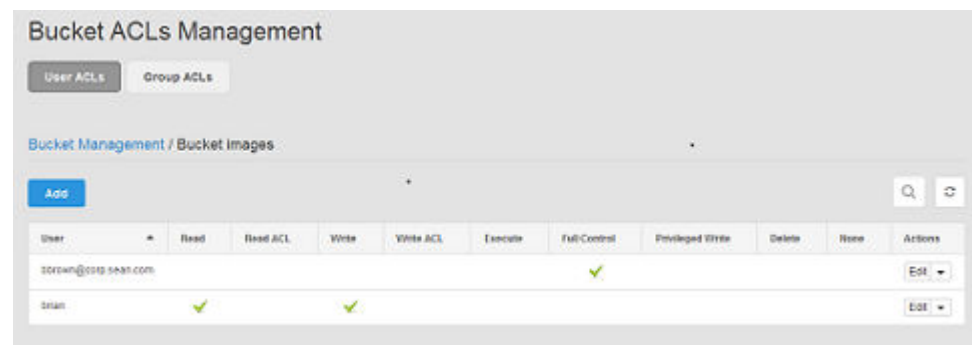
ACL	Permission
Read	Allows user to list the objects in the bucket.
Read ACL	Allows user to read the bucket ACL.
Write	Allows user to create or update any object in the bucket.
Write ACL	Allows user to write the ACL for the bucket.
Execute	Sets the execute permission when accessed as a file system. This 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.

Note

For information about ACLs with CAS buckets, see [ECS Data Access Guide: CAS](#) .

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 105.

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.
- For CAS-specific instructions on setting up a CAS bucket for a CAS object user, see [ECS Data Access Guide: CAS](#) .

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 104.

7. If required, specify a quota for the bucket.
The settings that you can apply are described in: [Quotas](#) on page 84.
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 101
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 102.

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 85.
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 98.

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 109 or [Set custom group bucket ACLs](#) on page 110 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 98.
 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 98

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.

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

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

Create a bucket using the S3 API (with s3curl)

You can use the S3 API to create a bucket in a 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

- To create a bucket, ECS must have at least one replication group configured.
- Ensure that Perl is installed on the Linux machine on which you will run `s3curl`.
- Ensure that `curl` tool and the `s3curl` tool are installed. The `s3curl` tool acts as a wrapper around `curl`.
- To use `s3curl` with `x-emc` headers, minor modifications must be made to the `s3curl` script. You can obtain the modified, ECS-specific version of `s3curl` from the [EMCECS Git Repository](#).

- Ensure that you have obtained a secret key for the user who will create the bucket. For more information, see *ECS Data Access Guide* available from the [ECS Support Site](#).

The EMC headers that can be used with buckets are described in [Bucket HTTP headers](#) on page 114.

Procedure

1. Obtain the identity of the replication group in which you want the bucket to be created, by typing the following command.

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

The response provides the name and identity of all data services virtual pools. In the following example, the ID is `urn:storageos:ReplicationGroupInfo:8fc8e19bedf0-4e81-bee8-79acc867f64:global`.

```
<data_service_vpools>
<data_service_vpool>
  <creation_time>1403519186936</creation_time>
  <id>urn:storageos:ReplicationGroupInfo:8fc8e19b-edf0-4e81-
bee8-79acc867f64:global</id>
  <inactive>>false</inactive>
  <tags/>
  <description>IsilonVPool1</description>
  <name>IsilonVPool1</name>
  <varrayMappings>
    <name>urn:storageos:VirtualDataCenter:
1de0bbc2-907c-4ede-b133-f5331e03e6fa:vdc1</name>
    <value>urn:storageos:VirtualArray:793757ab-ad51-4038-
b80a-682e124eb25e:vdc1</value>
  </varrayMappings>
</data_service_vpool>
</data_service_vpools>
```

2. 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 following example, the profile `my_profile` references 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'
  },
);
```


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

The endpoint is the address of your data node or the load balancer that sits in front of your data nodes.

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

4. Create the bucket using `s3curl.pl` and specify the following parameters:

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

The following example shows a fully specified command.

```
./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>
```

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

Note

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 NFS from an environment that is not secured using Kerberos).

If successful (with `--debug` on) output similar to the following appears:

```
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-
bee8-79accc867f64: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
```

After you finish

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

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

Bucket HTTP headers

There are a number of headers that determine the behavior of ECS when creating buckets using the objects APIs.

The following `x-emc` headers are provided.

Table 8 Bucket headers

Header	Description
<code>x-emc-dataservice-vpool</code>	Determines the replication group is used to store the objects associated with this bucket. If you do not specify a replication group using the <code>x-emc-dataservice-vpool</code> header, ECS selects the default replication group associated with the namespace.
<code>x-emc-file-system-access-enabled</code>	Configures the bucket for NFS or HDFS access. The header must not conflict with the interface that is used. That is, a create bucket request from NFS or HDFS cannot specify <code>x-emc-file-system-access-enabled=false</code> .
<code>x-emc-namespace</code>	Specifies the namespace used for this bucket. If the namespace is not specified using the S3 convention of host-style or path-style request, then it is specified using the <code>x-emc-namespace</code> header. If the namespace is not specified in this header, the namespace associated with the user is used.
<code>x-emc-retention-period</code>	Specifies the retention period that is 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.
<code>x-emc-is-stale-allowed</code>	Specifies whether the bucket is accessible during a temporary VDC outage in a federated configuration
<code>x-emc-server-side-encryption-enabled</code>	Specifies whether objects written to a bucket are encrypted.

Table 8 Bucket headers (continued)

Header	Description
x-emc-metadata-search	Specifies one or more user or system metadata values that are used to create indexes of objects for the bucket. The indexes are used to perform object searches that are filtered based on the indexed metadata.

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 115
- [OpenStack Swift container and object naming in ECS](#) on page 116
- [Atmos bucket and object naming in ECS](#) on page 116
- [CAS pool and object naming in ECS](#) on page 117

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 12

Configure NFS file access

- [NFS file access](#)..... 120
- [ECS Portal support for NFS configuration](#)..... 121
- [ECS NFS configuration tasks](#)..... 123
- [Best practice when using ECS NFS](#)..... 138
- [Permissions for multi-protocol \(cross-head\) access](#)..... 138
- [File API Summary](#)..... 140

NFS file access

ECS enables object buckets to be configured for access as NFS filesystems using NFSv3.

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, OpenStack Swift and EMC Atmos object protocols. Similarly, objects written using S3 and OpenStack Swift object protocols can be made available through NFS. For Atmos, objects created using the namespace interface can be listed using NFS, however, objects created using an object ID cannot. 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.

ECS NFS provides advisory locking and supports:

- Lock over multiple zones
- Shared and exclusive locks

ECS NFS supports Kerberos security.

Multi-protocol access to directories and files

ECS supports writing objects using the S3 protocol and accessing them as files using NFS and, conversely, writing files using NFS and accessing the files as objects using the S3 protocol. It is important to understand how directories are managed when using multi-protocol access.

The S3 protocol does not make provision for the creation of folders/directories.

To enable multi-protocol operation, ECS support for the S3 protocol formalizes the use of "/" and creates "directory" objects for all intermediate paths in an object name. So an object called "/a/b/c.txt" will result in the creation of a file object called "c.txt" and directory objects for "a" and "b". Note that these directory objects are not exposed to the customer via S3, and are only maintained to provide multi-protocol access and compatibility with filesystem based APIs. This means that when the bucket is viewed as an NFS/HDFS filesystem, ECS can display files within a directory structure.

Limitations

1. An issue can arise where both a directory object and an file object are created with the same name. This can occur in the following ways:
 - If a file "path1/path2" is created from NFS, then an object "path1/path2/path3" is created from S3. Since S3 allows creation of objects that have another object's name as prefix, this is a valid operation and is supported. However, at this point, a file and a directory called "path2" will exist.
 - If a directory "path1/path2" is created from NFS, then an object "path1/path2" is created from S3. Once again, this is a valid operation from S3 since directory

"path1/path2" is not visible via the S3 API. However, at this point, a file and a directory called "path2" will exist.

To resolve this situation, requests from S3 will always return the file and requests from NFS will always return the directory. However, this means that in the first case the file created by NFS will be hidden by the object created by S3.

2. NFS does not support filenames with a trailing "/" in them, but this is supported by S3 protocol. In such cases, NFS will not show these files.

Node and site failure

NFS provides a single namespace across all ECS nodes and can continue to operate in the event of node or site failure.

When you mount an NFS export, you can specify any of the ECS nodes as the NFS server or you can specify the address of a load balancer. Whichever node you point at, the ECS is able to resolve the filesystem path.

Node Failure

In the event of a node failure, ECS will recover data using its data fragments. If your NFS export is configured for Async writes, you will run the risk of losing data related to any transactions that have not yet been written to disk. This is the same with any NFS implementation.

If you have mounted the filesystem by pointing at an ECS node and that node fails, you will need to remount the export by specifying a different node as the NFS server. If you mounted the export using the load balancer address, failure of the node will be handled by the load balancer which will automatically direct requests to a different node.

Site Failure

In the event of a site failure, if your load balancer is able to redirect traffic to a different site, your NFS export will continue to be available. Otherwise you will need to remount the export from another, non-failed site.

When the primary site fails, and ECS is required to reconfigure to point at a secondary site, data can be lost due to NFS Async writes and also due to unfinished ECS data replication operations.

ECS Portal support for NFS configuration

The ECS Portal provides support for creating NFS exports and for mapping ECS users so that they can access the NFS export.

The ECS Portal **Manage > File** page enables an administrator to configure NFS exports and to set up user and group mappings. The page comprises Exports and a User/Group Mapping views. By default the Exports area is shown. The ability to configure NFS exports and to set up user mappings is also available using the ECS Management REST API and the CLI.

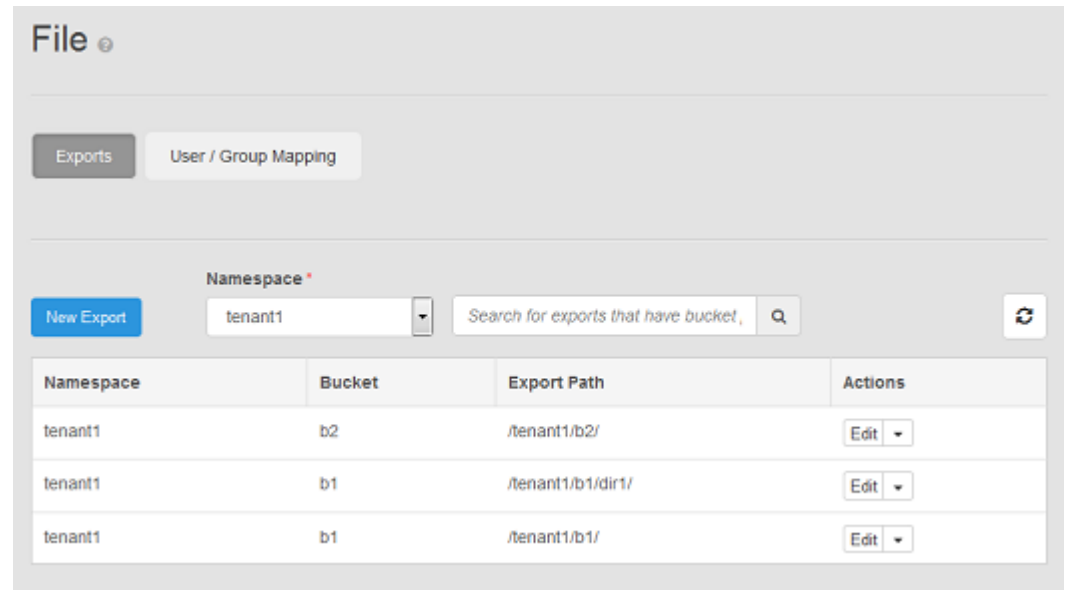
The pages are described in the following topics:

- [Exports](#) on page 122
- [User/Group mappings](#) on page 122

Exports

The Exports view shows the NFS exports that have been created and enables you to create new NFS exports and edit existing exports.

The Exports view is shown below.



The Exports view has a Namespace menu that allows you to select the namespace for which you want to see the currently defined exports. The following fields are displayed in the Exports Table:

Field	Description
Namesapce	The tenant/namespace that the underlying storage belongs to.
Bucket	The bucket that provides the underlying storage for the NFS export.
Export Path	The mount point associated with the export. This is in the form: /<namespace_name>/<bucket_name>/<export_name> . The export name will only be specified if you are exporting a directory that exists within the bucket.
Actions	The actions that can be performed on the export table entry. Comprises: <code>Edit</code> and <code>Delete</code> . You can use <code>Edit</code> to view the hosts that can access the export and the NFS options that have been set for the export.

The page also provides access to a **New Export** button to enable an export to be added.

User/Group mappings

ECS stores the owner and group for the bucket, and the owner and group for files and directories within the bucket, as ECS object username and custom group names,

respectively. The names must be mapped to Unix IDs in order that NFS users can be given access with the appropriate privileges.

The mapping enables ECS to treat an ECS object user and an NFS user as the same user but with two sets of credentials, one to access ECS using NFS, and one to access the ECS using the object protocols. Because the accounts are mapped, files written by an NFS user will be accessible as objects by the mapped object user and objects written by the object users will be accessible as files by the NFS user.

The permissions associated with the file or object will be based on a mapping between POSIX and object protocol ACL privileges. The mapping is described in detail in [Permissions for multi-protocol \(cross-head\) access](#) on page 138.

The **Manage > File** page of the ECS Portal provides access to a User/Group Mapping view that displays a User/Group Mapping table.

The view provides a Namespace selector that enables the table to show the users and groups configured for the selected namespace. The User/Group Mapping table displays the following fields:

Field	Description
User/Group Name	The object username of the user.
ID	The Unix User ID or Group ID that has been mapped to the object user.
Type	Indicates whether the ID is for a User or Group.
Actions	The actions that can be performed on the entry. They are: View and Delete

New mappings can be added using the **New User/Group Mapping** button.

ECS NFS configuration tasks

To configure ECS NFS, the following tasks must be performed.

Procedure

1. [Create a bucket for NFS using the ECS Portal](#) on page 124

2. [Add an NFS export](#) on page 126
3. [Add a user or group mapping](#) on page 129
4. [Configure NFS security with Kerberos](#) on page 130

Create a bucket for NFS using the ECS Portal

Use the ECS Portal to create a bucket configured for use with NFS.

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.

The steps provided here focus on the configuration you will need to perform to make a bucket suitable for use by NFS. The bucket you create is an S3 bucket enabled for file system use.

Procedure

1. At the ECS Portal, select **Manage > Buckets > New Bucket**.
2. Enter a name for the bucket.
3. Specify the namespace that the bucket will belong to.
4. Select a Replication Group or leave blank to use the default replication group for the namespace.
5. Enter the name of the bucket owner.
6. Do not enable CAS.

Note

A bucket that is intended for use as NFS cannot be used for CAS. The CAS control is disabled when File System is enabled.

7. Enable any other bucket features that you require.

You can enable any of the following features on a NFS bucket:

- Quota
- Server-side Encryption
- Metadata Search
- Access During Outage
- Compliance (see note)
- Bucket Retention

Refer to [Bucket concepts](#) on page 98 for information on each of these settings and how to configure them.

Note

A bucket that is compliance-enabled cannot be written to using the NFS protocol. However, data written using object protocols can be read from NFS.

8. Select Enabled for the File System.

Once enabled, controls for setting a default group for the filesystem/bucket and for assigning group permissions for files and directories created in the bucket are available.

9. At the File System panel, shown below, enter a name for the Default Bucket Group.

The screenshot shows a configuration panel titled "File System". At the top, there are two toggle buttons: "Disabled" and "Enabled", with "Enabled" being the active one. Below this is a section labeled "Default Bucket Group" with a text input field. Underneath are two sections: "Group File Permissions" and "Group Directory Permissions". Each of these sections contains three buttons: "Read", "Write", and "Execute".

This group will be the group associated with the NFS root filesystem and with any files or directories created in the NFS export. It enables users who are members of the group to access the NFS export and to access files and directories.

This group must be specified at bucket creation. If it is not, the group would have to be assigned later from the NFS client.

10. Set the default permissions for files and directories created in the bucket using the object protocol.

These settings are used to apply Unix group permissions to objects created using object protocols.

The S3 protocol does not have the concept of groups so there is no opportunity for setting group permissions in S3 and mapping them to Unix permissions. Hence, this provides a one-off opportunity for a file or directory created using the S3 protocol to be assigned to the specified default group with the permissions specified here.

- a. Set the Group File Permissions by clicking the appropriate permission buttons.

You will normally set Read and Execute permissions.

- b. Set the Group Directory Permissions by clicking the appropriate permission buttons.

You will normally set Read and Execute permissions.

11. Click **Save** to create the bucket.

Add an NFS export

The **File > Exports** page enables you to create an NFS export and set the options that control access to the export.

Before you begin

The bucket that will provide the underlying storage for the export must have been created.

Procedure

1. Select the **File > Exports > New Export** page.

The New File Export panel is displayed. An example of the panel is shown below with an export configured for access by NFS client hosts.

The screenshot shows the 'New File Export' configuration interface. At the top, there are two tabs: 'Exports' (selected) and 'User / Group Mapping'. Below the tabs is the title 'New File Export'. The configuration fields are as follows:

- Namespace ***: A dropdown menu with 'tenant1' selected.
- Bucket ***: A dropdown menu with 'b1' selected.
- Export Path ⓘ**: A text input field containing '/tenant1/b1/'.
- Export Host Options ⓘ**: A table with an 'Add' button in the top right corner.

Host	Summary	Actions
nfsclient1	rw,async,authsys	Edit ▼
nfsclient2	ro,async	Edit ▼

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

2. From the namespace field, select the namespace that owns the bucket that you want to export.
3. From the bucket field, select the bucket.
4. In the Export Path field, specify the path.

ECS automatically generates the export path based on the namespace and bucket. You only need to enter a name if you are exporting a directory that

already exists within the bucket. So if you enter `/namespace1/bucket1/dir1`, for example, you should ensure that `dir1` exists. If it does not, mounting the export will fail.

5. Add the hosts that you want to be able to access the export using the following steps.

- a. In the Export Options area, select **Add** to add.

The Add Export Host panel, shown below, is displayed.

- b. At the Add Export Host panel, specify one or more hosts that you want to be able to access the export and configure the access options.

You must choose an Authentication option. This will normally be `Sys` unless you are intending to configure Kerberos. Default values for Permissions (`ro`) and Write Transfer Policy (`async`) are already set on the Add Export Host panel and will be passed to the NFS sever. The remaining options are the same as the NFS server defaults and so will only be passed by ECS if you change them.

The parameters that you can specify when adding a host are shown in the table below.

Setting	Description
Export Host	Sets the IP address of the host or hosts that can access the export. A comma separated list is used to specify more than one host.
Permissions	Enables access to the export to be set as Read/Write or Read only. This is the same as setting <code>rw</code> or <code>ro</code> in <code>/etc/exports</code> .
Write Transfer Policy	Sets the write transfer policy as synchronous or asynchronous. The default is asynchronous. This is the same as setting <code>sync</code> or <code>async</code> for an export in <code>/etc/exports</code> .
Authentication	Sets the authentication types that will be supported by the export.
Mounting Directories Inside Export	This setting determines whether subdirectories of the export path will be allowed as mount points. This is the same as the <code>alldir</code> setting in <code>/etc/exports</code> . With the <code>alldir</code> option, if you have exported <code>/namespace1/bucket1</code> , for example, you will also be able to mount subdirectories, such as <code>/namespace1/bucket1/dir1</code> , provided the directory exists.
AnonUser	Sets the effective user ID for anonymous user access to an export and for root access where <code>root_squash</code> has been set. This is the same as setting <code>anonuid</code> in <code>/etc/exports</code> .
AnonGroup	Sets the effective group ID for anonymous group access to an export and for root access where <code>root_squash</code> has been set. This is the same as setting <code>anongid</code> in <code>/etc/exports</code> .
RootSquash	Determines whether root is allowed on the export. If root is disallowed, the UID of the root user (UID=0) is translated to the UID of the user "nobody", or to the UID you specify in AnonUser. This is the same as using <code>root_squash</code> in <code>/etc/exports</code> .

c. Select **Add** to finish defining the host options.

- If you want to add more hosts that can access the export, but with different options, repeat the previous step.

- Click **Save** to save the NFS export definition.

Add a user or group mapping

To provide NFS access to the filesystem (the bucket), you must map an object user who has permissions on the bucket to a Unix User Id (UID) in order to provide access for the mapped Unix user, alternatively, you can map an ECS custom group that has permissions on the bucket to a Unix Group Id (GID) to provide access for members of the Unix group.

Before you begin

- For the mapping to work, the UID must exist on the NFS client and the username must be an ECS object username.
- For group members to have access to the filesystem a default custom group must have been assigned to the bucket.
- For group members to have access to objects and directories created using object protocols, default object and directory permissions must have been assigned to the bucket.

Procedure

- At the **Manage > File** page, select the **User / Group Mapping** view.
- Select **New Mapping** to display the New User Mapping form, shown below.

The screenshot shows a web interface for managing file exports. At the top, there's a 'File' header with a help icon. Below it are two tabs: 'Exports' and 'User / Group Mapping'. The 'User / Group Mapping' tab is active. Underneath, there's a section titled 'New User / Group Mapping'. This section contains several form fields: 'User / Group Name' with a red asterisk and a help icon, an empty text input field; 'Namespace' with a red asterisk, a dropdown menu showing 'tenant1', and a small downward arrow; 'ID' with a red asterisk and a help icon, an empty text input field; and 'Type' with a help icon, two radio buttons labeled 'User' and 'Group', where 'User' is selected. At the bottom of the form are two buttons: 'Save' (blue) and 'Cancel' (white).

- In the User/Group field, enter the ECS user name or group name that you want to map.

4. Specify the namespace that the ECS object user or custom group, to which you are going to map the Unix user or group, belongs.
5. In the UID field, enter to user ID or group ID that you want the ECS username to map to.
6. Select the Type of mapping: User or Group.
7. Select **Save**.

Configure NFS security with Kerberos

You can secure access to your NFS export using Kerberos. The following scenarios are supported:

- ECS client to single ECS node. The keytab on each ECS that you want to use as the NFS server must be specific to that node.
- ECS client to load balancer. Keytab on all ECS nodes will be the same, and will use the hostname of the load balancer.

Refer to [Configure ECS NFS with Kerberos security](#) on page 130.

Depending on your internal IT setup, you can use a KDC or you can Active Directory (AD) as your KDC.

To use AD, follow the steps in the following tasks: Refer to [Register an ECS node with Active Directory](#) on page 134 and Refer to [Register a Linux NFS client with Active Directory](#) on page 135.

Configure ECS NFS with Kerberos security

To configure Kerberos authentication to secure ECS NFS, you need to perform configuration on both the ECS node(s) and the NFS client, and create keytabs for the NFS server principal and for the NFS client principal.

Procedure

1. Make sure that the hostname of the ECS node can be resolved.

You can use the `hostname` command to ensure that the FQDN of the ECS node is added to `/etc/HOSTNAME`.

```
dataservice-10-247-142-112:~ # hostname ecsnode1.yourco.com
dataservice-10-247-142-112:~ # hostname -i
10.247.142.112
dataservice-10-247-142-112:~ # hostname -f
ecsnode1.yourco.com
dataservice-10-247-142-112:~ #
```

2. Create the Kerberos configuration file (`krb5.conf`) on the ECS node as `/opt/emc/caspian/fabric/agent/services/object/data/hdfs/krb5.conf`. Unless HDFS has already been configured, you will need to create the `hdfs` directory with `655 (drwxr-xr-x)` permissions (`chmod 655 hdfs`) and make user with uid 444 and group with gid 444 as the owner (`chown 444:444 hdfs`).

Change the file permissions to `644` and make the user with id 444(`storageos`) the owner of the file.

In the example below, the following values are used and will need to be replaced with your own settings.

Kerberos REALM

Set to NFS-REALM in this example.

KDC

Set to kdcname.yourco.com in this example.

KDC Admin Server

In this example, the KDC acts as the admin server.

```
[libdefaults]
    default_realm = NFS-REALM.LOCAL
[realms]
    NFS-REALM.LOCAL = {
        kdc = kdcname.yourco.com
        admin_server = kdcname.yourco.com
    }
[logging]
    kdc = FILE:/var/log/krb5/krb5kdc.log
    admin_server = FILE:/var/log/krb5/kadmind.log
    default = SYSLOG:NOTICE:DAEMON
```

Note

If HDFS for Kerberos is already configured, instead of replacing `/opt/emc/caspian/fabric/agent/services/object/data/hdfs/krb5.conf`, merge the REALM information, if it is different, into the existing `krb5.conf` file. Usually there will be no change to this file as REALM will have been configured by HDFS. In addition, the default permissions and owner will have been already configured by HDFS and will not require any change.

3. Add a host principal for the ECS node and create a keytab for the principal.

In this example, the FQDN of the ECS node is `ecsnode1.yourco.com`

```
$ kadmin
kadmin> addprinc -randkey nfs/ecsnode1.yourco.com
kadmin> ktadd -k /datanode.keytab nfs/ecsnode1.yourco.com
kadmin> exit
```

4. Copy the keytab (`datanode.keytab`) to `/opt/emc/caspian/fabric/agent/services/object/data/hdfs/krb5.keytab`. Unless HDFS has already been configured, you will need to create the `hdfs` directory with 655 (`drwxr-xr-x`) permissions (`chmod 655 hdfs`) and make user with uid 444 and group with gid 444 as the owner (`chown 444:444 hdfs`).

Change its file permissions to 644 and make the user with id 444(`storageos`) the owner of the file.

If HDFS is already configured, instead of replacing `/opt/emc/caspian/fabric/agent/services/object/data/hdfs/krb5.keytab`, merge the `datanode.keytab` file into the existing keytab file using `ktutil`. Default permissions and owner will have been already configured by HDFS and will not require any change.

- Download the "unlimited" JCE policy archive from oracle.com and extract it to the `/opt/emc/caspian/fabric/agent/services/object/data/jce/unlimited` directory.

Kerberos may be configured to use a strong encryption type, such as AES-256. In that situation, the JRE within the ECS nodes must be reconfigured to use the 'unlimited' policy.

Note

This step should be performed only if you are using a strong encryption type.

If HDFS is already configured, this step would have been completed by HDFS Kerberos configuration.

- Run the following command from inside the object container.

```
service storageos-dataservice restarthdfs
```

- To set up the client, begin by making sure that the hostname of the client can be resolved.

You can use the `hostname` command to ensure that the FQDN of the ECS node is added to `/etc/HOSTNAME`.

```
dataservice-10-247-142-112:~ # hostname ecsnode1.yourco.com
dataservice-10-247-142-112:~ # hostname -i
10.247.142.112
dataservice-10-247-142-112:~ # hostname -f
ecsnode1.yourco.com
dataservice-10-247-142-112:~ #
```

- If your client is running SUSE Linux make sure that line `NFS_SECURITY_GSS="yes"` is uncommented in `/etc/sysconfig/nfs`.
- If you are on Ubuntu make sure to have line `NEED_GSSD=yes` in `/etc/default/nfs-common`.
- Install `rpcbind` and `nfs-common`.

Use `apt-get` or `zypper`. On SUSE Linux, for `nfs-common`, use:

```
zypper install yast2-nfs-common
```

By default these are turned off in Ubuntu client.

- Set up your Kerberos configuration file.

In the example below, the following values are used and will need to be replaced with your own settings.

Kerberos REALM

Set to `NFS-REALM` in this example.

KDC

Set to `kdcname.yourco.com` in this example.

KDC Admin Server

In this example, the KDC acts as the admin server.

```
[libdefaults]
    default_realm = NFS-REALM.LOCAL
[realms]
    NFS-REALM.LOCAL = {
        kdc = kdcname.yourco.com
        admin_server = kdcname.yourco.com
    }
[logging]
    kdc = FILE:/var/log/krb5/krb5kdc.log
    admin_server = FILE:/var/log/krb5/kadmind.log
    default = SYSLOG:NOTICE:DAEMON
```

12. Add a host principal for the NFS client and create a keytab for the principal. In this example, the FQDN of the NFS client is `nfsclient.yourco.com`

```
$kadmin
kadmin> addprinc -randkey host/nfsclient.yourco.com
kadmin> ktadd -k /nkclient.keytab host/nfsclient.yourco.com
kadmin> exit
```

13. Copy the keytab file (`nfsclient.keytab`) from the KDC machine to `/etc/krb5.keytab` on the NFS client machine.

```
scp /nkclient.keytab root@nfsclient.yourco.com:/etc/krb5.keytab
ssh root@nfsclient.yourco.com 'chmod 644 /etc/krb5.keytab'
```

14. Create a principal for a user to access the NFS export.

```
$kadmin
kadmin> addprinc yourusername@NFS-REALM.LOCAL
kadmin> exit
```

15. Log in as root and add the following entry to your `/etc/fstab` file.

```
HOSTNAME:MOUNTPOINT    LOCALMOUNTPOINT      nfs
rw,user,nolock,noauto,vers=3,sec=krb5 0    0
```

For example:

```
ecsnodel.yourco.com:/s3/b1    /home/kothan3/lb1
nfs    rw,user,nolock,noauto,vers=3,sec=krb5    0 0
```

16. Log in as non root user and `kinit` as the non-root user that you created.

```
kinit yourusername@NFS-REALM.LOCAL
```

17. You can now mount the NFS export.

Note

Mounting as the root user will not require you to use kinit. However, when using root, authentication is done using the client machine's host principal rather than your Kerberos principal. Depending upon your operating system, you can configure the authentication module to fetch the Kerberos ticket when you login, so that there is no need to fetch the ticket manually using kinit and you can mount the NFS share directly.

Register an ECS node with Active Directory

To use Active Directory (AD) as the KDC for your NFS Kerberos configuration, you need to create accounts for the client and server in AD and map the account to a principal. For the NFS server, the principal represents the NFS service accounts, for the NFS client, the principal represents the client host machine.

Before you begin

You must have administrator credentials for the AD domain controller.

Procedure

1. Log in to AD.
2. In Server Manager, go to **Tools > Active Directory Users and Computers**.
3. Create a user account for the NFS principal using the format "nfs-<host>". For example: "nfs-ecsnodel". Set the password to never expire.
4. Create an account for yourself (optional and one time).
5. Execute the following command to create a keytab file for the NFS service account.

```
ktpass -princ nfs/<fqdn>REALM.LOCAL +rndPass -mapUser nfs-
<host>@REALM.LOCAL -mapOp set -crypto All -ptype
KRB5_NT_PRINCIPAL -out filename.keytab
```

For example, to associate the nfs-ecsnodel account with the principle nfs/ecsnodel.yourco.com@NFS-REALM.LOCAL, you can generate a keytab using:

```
ktpass -princ nfs/ecsnodel.yourco.com@NFS-REALM.LOCAL
+rndPass -mapUser nfs-ecsnodel@NFS-REALM.LOCAL -mapOp set -
crypto All -ptype KRB5_NT_PRINCIPAL -out nfs-ecsnodel.keytab
```

6. Import the keytab to the ECS node.

```
ktutil
ktutil> rkt <keytab to import>
ktutil> wkt /etc/krb5.keytab
```

7. Test registration by running.

```
kinit -k nfs/<fqdn>@NFS-REALM.LOCAL
```

8. See the cached credentials by running the `klist` command.
9. Delete the cached credentials by running the `kdestroy` command.
10. View the entries in the keytab file by running the `klist` command.

For example:

```
klist -kte /etc/krb5.keytab
```

11. Follow steps 2 on page 130, 4 on page 131, and 5 on page 132 from [Configure ECS NFS with Kerberos security](#) on page 130 to place the Kerberos configuration files (`krb5.conf`, `krb5.keytab` and `jce/unlimited`) on the ECS node.

Register a Linux NFS client with Active Directory

To use Active Directory (AD) as the KDC for your NFS Kerberos configuration, you need to create accounts for the client and server in AD and map the account to a principal. For the NFS server, the principal represents the NFS service accounts, for the NFS client, the principal represents the client host machine.

Before you begin

You must have administrator credentials for the AD domain controller.

Procedure

1. Log in to AD.
2. In Server Manager, go to **Tools > Active Directory Users and Computers**.
3. Create a computer account for the client machine. For example: "nfsclient". Set the password to never expire.
4. Create an account for a user (optional and one time)
5. Execute the following command to create a keytab file for the NFS service account.

```
ktpass -princ host/<fqdn>@REALM.LOCAL +rndPass -mapUser  
<host>@REALM.LOCAL -mapOp set -crypto All -ptype  
KRB5_NT_PRINCIPAL -out filename.keytab
```

For example, to associate the `nfs-ecsnode1` account with the principle `host/nfsclient.yourco.com@NFS-REALM.LOCAL`, you can generate a keytab using:

```
ktpass -princ host/nfsclient.yourco.com@NFS-REALM.LOCAL  
+rndPass -mapUser nfsclient$@NFS-REALM.LOCAL -mapOp set -  
crypto All -ptype KRB5_NT_PRINCIPAL -out nfsclient.keytab
```

6. Import the keytab to the client node.

```
ktutil  
ktutil> rkt <keytab to import>  
ktutil> wkt /etc/krb5.keytab
```

7. Test registration by running.

```
kinit -k host/<fqdn>@NFS-REALM.LOCAL
```

8. See the cached credentials by running the `klist` command.
9. Delete the cached credentials by running the `kdestroy` command.
10. View the entries in the keytab file by running the `klist` command.

For example:

```
klist -kte /etc/krb5.keytab
```

11. Follow steps [2](#) on page 130, [4](#) on page 131, and [5](#) on page 132 from [Configure ECS NFS with Kerberos security](#) on page 130 to place the Kerberos configuration files (`krb5.conf`, `krb5.keytab` and `jce/unlimited`) on the ECS node.

Mounting an NFS export : example

When mounting an export, it is important that:

- The bucket owner name has been mapped to a Unix UID.
- A default group has been assigned to the bucket. For the default group to show as the associated Linux group when the export is mounted, a mapping between its name and a Linux GID must have been created.

The following steps provide and an example of how to mount an ECS NFS export filesystem.

1. Create a directory on which to mount the export. The directory should belong to the same owner as the bucket.
In this example, we will use the user "fred" to create a directory `/home/fred/nfsdir` on which to mount an export.

```
su - fred
mkdir /home/fred/nfsdir
```

2. As the root user, mount the export in the directory mount point that you created.
For example:

```
mount -t nfs -o "vers=3,noexec" 10.247.179.162:/s3/tc-nfs6 /home/fred/nfsdir
```

When mounting an NFS export, you can specify the name or IP address of any of the nodes in the VDC or the address of the load balancer.

It is important that you specify `-o "vers=3"`

3. Check that you can access the filesystem as user "fred".
 - a. Change to user "fred".

```
$ su - fred
```


- b. Check you are in the directory in which you created the mount point directory.

```
$ pwd
/home/fred
```

- c. List the directory.

```
fred@lrmh229:~$ ls -al
total
drwxr-xr-x  7 fred  fredsgroup  4096 May 31 05:38 .
drwxr-xr-x 18 root   root         4096 May 30 04:03 ..
-rw-----  1 fred  fred         16 May 31
05:31 .bash_history
drwxrwxrwx  3 fred  anothergroup  96 Nov 24 2015 nfsdir
```

You can see that, in this example, the bucket owner is "fred" and a default group, "anothergroup", was associated with the bucket.

If no group mapping had been created, or no default group has been associated with the bucket, you will not see a group name but a large numeric value, as shown below.

```
fred@lrmh229:~$ ls -al
total
drwxr-xr-x  7 fred  fredssgroup  4096 May 31 05:38 .
drwxr-xr-x 18 root   root         4096 May 30 04:03 ..
-rw-----  1 fred  fred         16 May 31 05:31 .bash_history
drwxrwxrwx  3 fred  2147483647   96 Nov 24 2015 nfsdir
```

If you have forgotten the group mapping, you can rectify this by creating the appropriate mapping at the ECS Portal.

You can find the group ID by looking in `/etc/group`.

```
fred@lrmh229:~$ cat /etc/group | grep anothergroup
anothergroup:x:1005:
```

And adding a mapping between the name and GID (in this case: `anothergroup => GID 1005`).

If you try and access the mounted filesystem as the root user, or another user that does not have permissions on the filesystem, you will see `?`, as below.

```
root@lrmh229:~# cd /home/fred
root@lrmh229:/home/fred# ls -al
total
drwxr-xr-x  8 fred  fredsgroup  4096 May 31 07:00 .
drwxr-xr-x 18 root   root         4096 May 30 04:03 ..
-rw-----  1 fred  fred         1388 May 31 07:31 .bash_history
d????????? ? ?      ?           ?           ? nfsdir
```

Best practice when using ECS NFS

The following recommendations apply when mounting ECS NFS exports.

Use `async`

Whenever possible you should use the "async" mount option. Using this option dramatically reduces latency and improves throughput and reduces the number of connections from the client.

Set `wsize` and `rsize` to reduce round trips from the client

Where you are expecting to read and/or write large files, you should ensure that the read or write size of files is set appropriately using the `rsize` and `wsize` mount options. It is generally recommended that you set the `wsize` and `rsize` to the highest possible value to reduce the number of round trips from the client. This is typically 512KB (524288 B).

For example, to write a 10MB file, if the `wsize` is set to 524288 (512KB) the client would make 20 separate calls, whereas, if the write size had been set as 32KB this would result in 16 times as many calls.

When using the mount command, you can supply the read and write size using the options (-o) switch. For example:

```
# mount 10.247.97.129:/home /home -o
"vers=3,nolock,rsize=524288,wsize=524288"
```

Permissions for multi-protocol (cross-head) access

Objects can be accessed using NFS and using the object service. Object Access Control List (ACL) permissions and File System permissions are stored for each object.

When an object is saved using the object protocol, the permissions associated with the object owner are mapped to NFS permissions and the corresponding permissions are stored. Similarly, when an object is created or modified using NFS, the NFS permissions of the owner are mapped to object permissions and stored.

The S3 object protocol does not have the concept of groups, so changes to group ownership or permissions from NFS do not need to be mapped to corresponding object permissions. However, when a bucket is created, or objects (the equivalent of files and directories) are created within a bucket, ECS can assign Unix group permissions so that they can be accessed by NFS users.

The following ACL attributes are stored for NFS:

- Owner
- Group
- Other

and for object access, the following ACLs are stored:

- Users
- Custom Groups
- Groups (Pre-defined)
- Owner (a specific user from Users)

- Primary Group (a specific group from Custom Groups)

Note

You can find out more about bucket ACLs [Bucket ACLs](#) on page 104.

The table below show the way in which NFS ACL attributes map to object ACL attributes.

NFS ACL Attribute	Object ACL Attribute
Owner	User who is also Owner
Group	Custom Group that is also Primary Group
Others	Pre-Defined Group

Examples of this mapping are discussed later in this topic.

The following Access Control Entries (ACE) can be assign to each ACL attribute.

NFS ACEs:

- Read (R)
- Write (W)
- Execute (X)

Object ACEs:

- Read (R)
- Write (W)
- Execute (X)
- ReadAcl (RA)
- WriteAcl (WA)
- Full Control (FC)

Creating and modifying an object using NFS and accessing using the object service

When you create an object using the NFS protocol, the owner RWX permissions are mirrored to the ACL of the object user who is designated as the owner of the bucket. If the NFS owner has RWX permissions, this is translated to Full Control in the object ACL.

The permissions assigned to the group that the NFS file or directory belongs are reflected onto a custom group of the same name, if it exists. Permissions associated with Others are reflected into pre-defined groups permissions.

The example below illustrates this scenario.

NFS ACL	Setting	Object ACL	Setting
Owner	John : RWX	Users	John : Full Control
Group	ecsgroup : R-X --->	Custom Groups	ecsgroup : R-X
Other	RWX	Groups	All_Users : R, RA
		Owner	John
		Primary Group	ecsgroup

When a user accesses ECS using NFS and changes the ownership of an object, the new owner inherits the owner ACL permissions and, in addition, is given Read_ACL and Write_ACL. The previous owner permissions are kept in the object,users ACL.

When a chmod operation is performed, the permissions are reflected in the same way as when creating an object. Write_ACL is preserved in group and others if it already exists in the object user's ACL.

Creating and modifying objects using the object service and accessing using NFS

When you create an object using the object service, the owner of the object is automatically granted Full Control of the object. Where the object owner is granted Full Control, the file owner is granted RWX permissions. If the owner permissions are set to other than Full Control, the object RWX permissions are reflected onto the file RWX permissions, so an object owner with RX permissions will result in an NFS file owner with RX permissions.

The object primary group, which is set using the Default Group on the bucket, becomes the custom group that the object belongs to and the permissions are set based on the default permissions that have been set. These permissions are reflected onto the NFS.group permissions. In the same way as for owner permissions, if the object custom group has Full Control, these become RWX for the NFS group.

If pre-defined groups are specified on the bucket, these are applied to the object and are reflected onto the Others for the NFS ACLs.

The example below illustrates this scenario.

Object ACL Setting	Setting		NFS ACL
Users	John : Full Control		Owner John : RWX
Custom Groups	ecsgroup : R-X	---->	Group ecsgroup : R-X
Groups	All_Users : R, RA		Other RWX
Owner	John		
Primary Group	ecsgroup		

If a new owner is assigned, the permissions associated with that owner are applied to the object.

File API Summary

NFS access can be configured and managed using the ECS Management REST API.

The table below provides a summary of the available API.

Method	Description
POST /object/nfs/exports	Creates an export. The payload specifies the export path, the hosts that can access the export, and a string that defines the security settings for the export.
PUT/GET/DELETE /object/nfs/exports/{id}	Performs the selected operation on the specified export
GET /object/nfs/exports	Retrieves all user exports that have been defined for the current namespace.
POST /object/nfs/users	Creates a mapping between an ECS object user name or group name and a Unix user or group ID.

Method	Description
PUT/GET/DELETE /object/nfs/users/{mappingid}	Performs the selected operation on the specified user or group mapping.
GET /object/nfs/users	Retrieves all user mappings that have been defined for the current namespace.

The API documentation provides full details of the API and the documentation for the NFS export methods can be accessed [ECS API Reference](#).

Configure NFS file access

CHAPTER 13

Configure Event Notification servers

- [Configure Event Notification servers \(SNMP or Syslog\)](#) 144
- [Working with the SNMP and Syslog servers at the ECS Portal](#) 145
- [Add an SNMP v2 Trap recipient](#) 147
- [Add an SNMP v3 Trap recipient](#) 148
- [Add a Syslog server](#) 149

Configure Event Notification servers (SNMP or Syslog)

Add SNMP and Syslog server configurations to enable ECS to provide SNMP and Syslog data to systems external to ECS.

In ECS, there are two types of Event Notification servers:

SNMP

Simple Network Management Protocol (SNMP) servers, also known as SNMP Agents, provide data about network managed device status and statistics to SNMP Network Management Station clients.

To allow communication between SNMP Agents and SNMP Network Management Stations, you must configure both sides to use the same credentials. To use SNMP v2, both sides must use the same Community name. To use SNMP v3, both sides must use the same EngineID, username, authentication protocol and authentication passphrase, and privacy protocol and privacy passphrase.

To authenticate traffic between SNMP servers and SNMP Network Management Stations using SNMP v3 messaging, in order to verify message integrity between hosts, ECS supports the SNMP v3 standard use of these possible cryptographic hash functions:

- Message Digest 5 (MD5)
- Secure Hash Algorithm 1 (SHA-1)

To encrypt all traffic between SNMP servers and SNMP Network Management Stations, ECS supports encryption of SNMP v3 traffic using these cryptographic protocols:

- Digital Encryption Standard (using 56-bit keys)
- Advanced Encryption Standard (using 128-bit, 192-bit or 256-bit keys)

Note

Support for advanced security modes (AES192/256) provided by the ECS SNMP Trap feature may be incompatible with certain SNMP targets (e.g. iReasoning).

Syslog

Syslog servers provide a method for centralized storage and retrieval of system log messages. ECS supports forwarding of alerts and audit messages to remote syslog servers, and supports operations using these application protocols:

- BSD Syslog
- Structured Syslog

Alerts and Audit messages that are sent to Syslog servers are also displayed on the ECS portal, with the exception of OS level Syslog messages (such as node SSH login messages), which are only sent over to Syslog servers and not displayed in the ECS Portal.

Each node runs a Syslog server as a fabric object container. The message traffic occurs over either TCP or UDP, with UDP as the default.

ECS sends Audit log messages to Syslog servers, including the severity level, using this format:


```

${serviceType} ${eventType} ${namespace} ${userId} $
{message}

```

ECS sends Alert logs to Syslog servers using the same severity as appears in the ECS Portal, using this format:

```

${alertType} ${symptomCode} ${namespace} ${message}

```

ECS sends Fabric alerts using this format:

```

Fabric {symptomCode} "{description}"

```

ECS forwards OS logs to Syslog servers without any modification.

You can configure information about Event Notification servers from the ECS Portal (see [Working with the SNMP and Syslog servers at the ECS Portal](#) on page 145) or using the ECS Management REST API or CLI. Follow the procedures below to configure SNMP v2, SNMP v3, and Syslog servers in ECS.

- [Add an SNMP v2 Trap recipient](#) on page 147
- [Add an SNMP v3 Trap recipient](#) on page 148
- [Add a Syslog server](#) on page 149

Working with the SNMP and Syslog servers at the ECS Portal

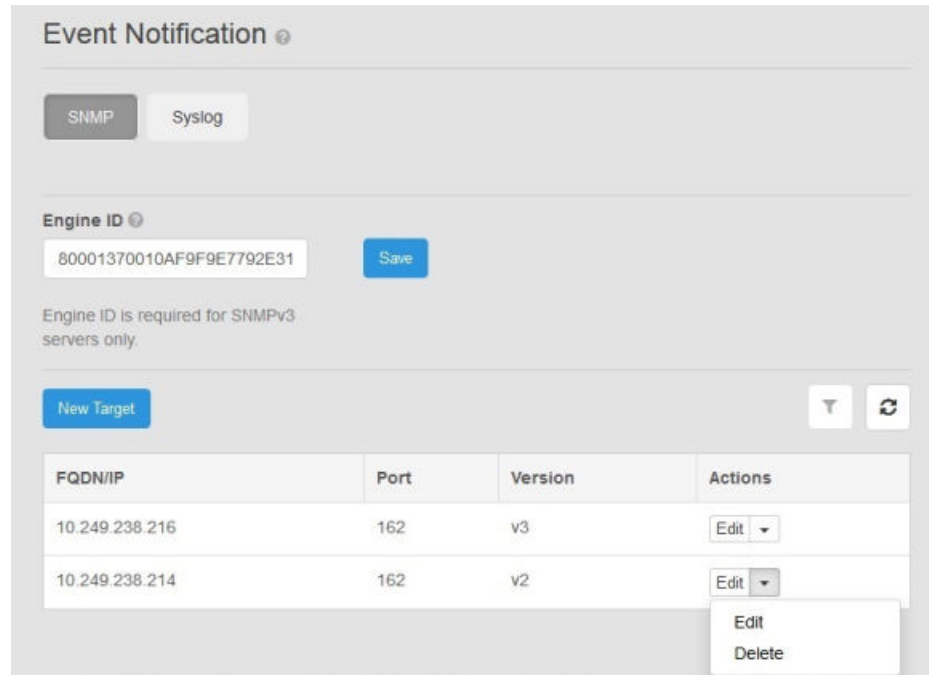
The ECS Portal provides a **Settings > Event Notification** page to configure SNMP servers.

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

The Event Notification Page provides tools to display and configure both SNMP server targets and Syslog servers, depending on the setting of the mode buttons. When you choose the **SNMP** mode button, the page displays an SNMP server table that lists the SNMP server configurations that have been created. When you choose the **Syslog** mode button, the page displays a Syslog server table that lists the Syslog server configurations that have been created.

An example of each appears shown below.

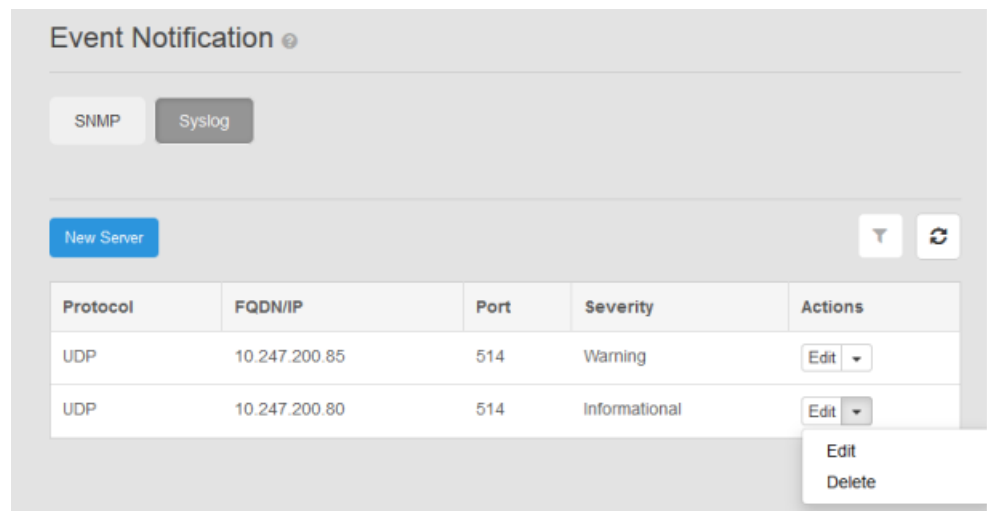
Figure 11 Event Notification page in SNMP mode



The SNMP server table provides access to the following information and operations.

Attribute	Description
FQDN/IP	A Fully Qualified Domain Name or IP address for the SNMP Trap recipient that runs the <code>snmptrapd</code> server.
Port	The port number of the <code>snmptrapd</code> on the SNMP Trap recipient.
Version	The version of SNMP that this <code>snmptrapd</code> server supports.
Actions	Provides a selection menu for the actions that are available. The actions that are available are: Edit and Delete .

Figure 12 Event Notification page in Syslog mode



The Syslog server table provides access to the following information and operations.

Attribute	Description
Protocol	The protocol you use to communicate with the Syslog server, either UDP or TCP.
FQDN/IP	A Fully Qualified Domain Name or IP address for the node that runs the Syslog server.
Port	The port number of the syslog service on the Syslog server.
Severity	The severity of threshold for messages to send to the log.
Actions	Provides a selection menu for the actions that are available. The actions that are available are: Edit and Delete .

The Event Notification Page additionally provides access to the following controls:

Control	Description
SNMP	The SNMP mode button toggles the page to present SNMP server information.
Syslog	The Syslog mode button toggles the page to present Syslog server information.
New Target	The New Target button allows you to configure a new SNMP server. This control only appears for SNMP servers.
New Server	The New Server button allows you to configure a new Syslog server. This control only appears for Syslog servers.

Add an SNMP v2 Trap recipient

You can configure Network Management Stations as SNMP v2 Trap recipients for SNMP Traps generated by the ECS Fabric using SNMP v2 standard messaging.

Before you begin

- To add an SNMP v2 Trap recipient, you must be assigned to the System Admin role in ECS. The root user has the System Admin role.
- You need access to the SNMP v2 credentials listed in [SNMP v2 server settings](#) on page 148.

Procedure

1. At the ECS Portal, select **Settings > Event Notification**.
2. Enter values for the attributes. Refer to [SNMP v2 server settings](#) on page 148
3. **Save**.

SNMP v2 server settings

You need to provide certain information when adding or editing an SNMP v2 server configuration for an SNMP v2c Trap recipient.

Field	Description
FQDN/IP	A Fully Qualified Domain Name or IP address for the SNMP v2c Trap recipient node that runs the <code>snmptrapd</code> server.
Port	The port number of the SNMP v2c <code>snmptrapd</code> running on the Network Management Station. Default= 162.
Version	The version of SNMP for this server. Use v2.
Community Name	The SNMP community name. Both the SNMP server and any Network Management Stations that access it must use the same community name in order to ensure authentic SNMP message traffic, as defined by the standards in RFC 1157 and RFC 3584. Default= public.

Add an SNMP v3 Trap recipient

You can configure Network Management Stations as SNMP v3 Trap recipients for SNMP Traps generated by the ECS Fabric using SNMP v3 standard messaging.

Before you begin

- To add an SNMP v3 Trap recipient, you must be assigned to the System Admin role in ECS. The root user has the System Admin role.
- You need access to the SNMP v3 credentials listed in [SNMP v3 server settings](#) on page 148.

Procedure

- At the ECS Portal, select **Settings > Event Notification**.
- Enter values for the attributes. Refer to [SNMP v3 server settings](#) on page 148
- Save**.

SNMP v3 server settings

You need to provide certain information when adding or editing an SNMP v3 server configuration for an SNMP v3 Trap recipient.

Field	Description
FQDN/IP	A Fully Qualified Domain Name or IP address for the SNMP v3 Trap recipient node that runs the <code>snmptrapd</code> server.
Port	The port number of the SNMP 3c <code>snmptrapd</code> running on the Network Management Station. Default= 162.

Field	Description
Version	The version of SNMP for this server. Use v3.
Username	A username to employ in authentication and message traffic as per the User-based Security Model (USM) defined by RFC 3414. Both the SNMP server and any Network Management Stations that access it must specify the same username in order to ensure communication. This is a octet string of up to 32 characters in length.
Authentication Protocol	The cryptographic hash function to use to verify message integrity between hosts. Choose between Message Digest 5 (MD5) or Secure Hash Algorithm 1 (SHA-1). Default= MD5.
Authentication Passphrase	The string to use as a secret key for authentication between SNMP v3 USM standard hosts, when calculating a message digest. The Passphrase can be 16 octets long for MD5 and 20 octets long for SHA-1.
Privacy Protocol	The cryptographic protocol to use in encrypting all traffic between SNMP servers and SNMP Network Management Stations. Choose between Digital Encryption Standard (DES) using 56-bit keys or Advanced Encryption Standard, using 128-bit, 192-bit or 256-bit keys. Default= DES.
Privacy Passphrase	The string to use in the encryption algorithm as a secret key for encryption between SNMP v3 USM standard hosts. The length of this key MUST be 16 octets for DES and longer for the AES protocols.

Note

When you create the first SNMP v3 configuration, the ECS system creates an SNMP Engine ID to use for SNMP v3 traffic. The **Event Notification** page displays that SNMP Engine ID in the **Engine ID** field. You could instead obtain an Engine ID from a Network Monitoring tool and specify that Engine ID in the **Engine ID** field. The important issue is that the SNMP server and any SNMP Network Management Stations that need to communicate with it using SNMP v3 traffic must use the same SNMP Engine ID in that traffic.

Add a Syslog server

You can configure a Syslog server to remotely store ECS logging messages.

Before you begin

- To add an Syslog server you must be assigned to the System Admin role in ECS. The root user has the System Admin role.

- You need access to the Syslog server credentials listed in [Syslog server settings](#) on page 150.

Procedure

1. At the ECS Portal, select **Settings > Event Notification**.
2. Enter values for the attributes. Refer to [Syslog server settings](#) on page 150.
3. **Save**.

Syslog server settings

You need to provide certain information when adding or editing a Syslog server configuration.

Field	Description
Protocol	The IP protocol to use for communication, either UDP or TCP. Default= UDP.
FQDN/IP	A Fully Qualified Domain Name or IP address for the node that runs the Syslog server.
Port	The port number for the Syslog server on which you want to store log messages. Default= 514.
Severity	<p>The severity of threshold for messages to send to the log. Choose from these severity levels:</p> <ul style="list-style-type: none"> • Debug • Informational • Notice • Warning • Error • Critical • Alert • Emergency

CHAPTER 14

Set the Base URL

- [Set the Base URL](#)..... 152
- [Bucket addressing](#)..... 152
- [Add a Base URL](#)..... 155

Set the Base URL

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 152
- [Add a Base URL](#) on page 155

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-ems-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 152
- [Path Based Addressing](#) on page 152

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 153
- [Base URL](#) on page 153

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-ems-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`

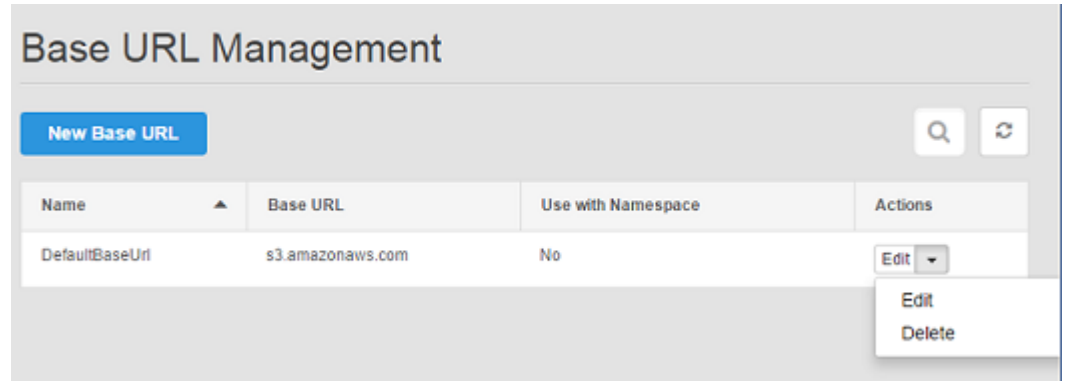
These entries allow the base name to be resolved when issuing service-level commands (for example, list buckets) and the virtual host style bucket address to be resolved.

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.



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.)
```

ECS 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.

Set the Base URL

6. Select **Save**.

CHAPTER 15

Configure certificates

- [Introduction to certificates](#)..... 158
- [Generating certificates](#)..... 158
- [Upload a certificate](#)..... 164
- [Verifying installed certificates](#)..... 167

Introduction to certificates

ECS ships with an SSL certificate installed in the keystore for each node. This certificate is not trusted by applications that talk to ECS, or by the browser when users access ECS through the ECS Portal.

To prevent users from seeing an untrusted certificate error, or to allow applications to communicate with ECS, you should install a certificate signed by a trusted Certificate Authority (CA). You can generate a self-signed certificate to use until you have a CA signed certificate. The self-signed certificate is installed into the certificate store of any machines that will access ECS.

ECS uses the following types of SSL certificates:

Management certificates

Used for management requests using the ECS Management REST API. These HTTPS requests use port 4443.

Object certificates

Used for requests using the supported object protocols. These HTTPS requests use ports 9021 (S3), 9023 (Atmos), 9025 (Swift).

You can upload a self-signed certificate, a certificate signed by a CA authority, or, for an object certificate, you can request ECS to generate a certificate or you. The key/certificate pairs can be uploaded to ECS by using the ECS Management REST API on port 4443.

The following topics explain how to create, upload, and verify certificates:

- [Generating certificates](#) on page 158
- [Upload a certificate](#) on page 164
- [Verifying installed certificates](#) on page 167

Generating certificates

You can generate a self-signed certificate, or you can purchase a certificate from a certificate authority (CA). The CA-signed certificate is strongly recommended for production purposes because it can be validated by any client machine without any extra steps.

Certificates must be in PEM-encoded x509 format.

When you generate a certificate, you typically specify the hostname where the certificate is used. Because ECS has multiple nodes, and each node has its own hostname, installing a certificate created for a specific hostname could cause a common name mismatch error on the nodes that do not have that hostname. You can create certificates with alternative IPs or hostnames called Subject Alternative Names (SANs).

For maximum compatibility with object protocols, the Common Name (CN) on your certificate must point to the wildcard DNS entry used by S3, because S3 is the only protocol that utilizes virtually-hosted buckets (and injects the bucket name into the hostname). You can specify only one wildcard entry on an SSL certificate and it must be under the CN. The other DNS entries for your load balancer for the Atmos and Swift protocols must be registered as a Subject Alternative Names (SANs) on the certificate.

The topics in this section show how to generate a certificate or certificate request using `openssl`, however, your IT organization may have different requirements or procedures for generating certificates.

Create a private key

Create a private key that is required to sign self-signed certificates and is used to create signing requests.

SSL uses public-key cryptography which requires a private and a public key. The first step in configuring it is to create a private key. The public key is created automatically, using the private key, when you create a certificate signing request or a certificate. The following steps describe how to use the `openssl` tool to create a private key.

Procedure

1. Log in to an ECS node or to a node that you can connect to the ECS cluster.
2. Use the `openssl` tool to generate a private key.

For example, to create a key called `server.key`, use:

```
openssl genrsa -des3 -out server.key 2048
```

3. When prompted, enter a passphrase for the private key and reenter it to verify. You will need to provide this passphrase when creating a self-signed certificate or a certificate signing request using the key.

You must create a copy of the key with the passphrase removed before uploading the key to ECS. Instructions on doing this are provided in the uploading section.

4. Set the permissions on the key file.

```
chmod 0400 server.key
```

Generate a SAN configuration

If you want your certificates to support Subject Alternative Names (SANs), you must define the alternative names in a configuration file.

OpenSSL does not allow you to pass Subject Alternative Names (SANs) through the command line, so you must add them to a configuration file first. To do this, you must locate your default OpenSSL configuration file. On Ubuntu, it is located at `/usr/lib/ssl/openssl.cnf`.

Procedure

1. Create the configuration file.

```
cp /usr/lib/ssl/openssl.cnf request.conf
```

2. Edit the configuration file with a text editor and make the following changes.
 - a. Add the `[alternate_names]`.

For example:

```
[ alternate_names ]
DNS.1 = os.example.com
```

```
DNS.2 = atmos.example.com
DNS.3 = swift.example.com
```

Note

There is a space between the bracket and the name of the section.

If you are uploading the certificates to ECS nodes rather than to a load balancer, the format is:

```
[ alternate_names ]
IP.1 = <IP node 1>
IP.2 = <IP node 2>
IP.3 = <IP node 3>
...
```

b. In the section [v3_ca], add the following lines:

```
subjectAltName      = @alternate_names
basicConstraints    = CA:FALSE
keyUsage            = nonRepudiation, digitalSignature, keyEncipherment
extendedKeyUsage    = serverAuth
```

The following line is likely to already exist in this [v3_ca] section. If you create a certificate signing request, you must comment it out as shown:

```
#authorityKeyIdentifier=keyid:always,issuer
```

c. In the [req] section, add the following lines:

```
x509_extensions = v3_ca      #for self signed cert
req_extensions   = v3_ca     #for cert signing req
```

d. In the section [CA_default], uncomment or add the line:

```
copy_extension=copy
```

Create a self-signed certificate

Create a self-signed certificate.

Before you begin

- Create a private key using the procedure in [Create a private key](#) on page 159.
- To create certificates that use SAN, you should create a SAN configuration file using the procedure in [Generate a SAN configuration](#) on page 159.

Procedure

1. Use the private key to create a self-signed certificate.

Two ways of creating the signing request are shown. One for use if you have already prepared a SAN configuration file to specify the alternative server name, another if you have not.

If you are using SAN:

```
openssl req -x509 -new -key server.key -config request.conf -
out server.crt
```

If you are not, use:

```
openssl req -x509 -new -key server.key -out server.crt
```

Example output.

```
Signature ok
subject=/C=US/ST=GA/
```

2. Enter the pass phrase for your private key.
3. At the prompts, enter the fields for the DN for the certificate.

Most fields are optional. You must enter a Common Name (CN).

Note

The CN should be a FQDN. Even if you install the certificate on the ECS nodes, you must use an FQDN and all of the IP addresses must be in the alternate names section.

You will see the following prompts:

```
You are about to be asked to enter information that will be
incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished
Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [AU]:US
State or Province Name (full name) [Some-State]:
Locality Name (eg, city) []:
Organization Name (eg, company) [Internet Widgits Pty
Ltd]:Acme
Organizational Unit Name (eg, section) []:
Common Name (e.g. server FQDN or YOUR name) []:*.acme.com
Email Address []:
```

4. Enter the Distinguished Name (DN) details when prompted. More information on the DN fields are provided in [Distinguished Name \(DN\) fields](#) on page 162.

5. View the certificate.

```
openssl x509 -in server.crt -noout -text
```

Distinguished Name (DN) fields

The following table describes the fields that comprise the Distinguished Name (DN).

Name	Description	Example
Common Name (CN)	The fully qualified domain name (FQDN) of your server. This is the name that you specified when you installed the ECS appliance.	*.yourco.com ecs1.yourco.com
Organization	The legal name of your organization. This must not be abbreviated and should include suffixes such as Inc, Corp, or LLC.	Yourco Inc.
Organizational Unit	The division of your organization handling the certificate.	IT Department
Locality/City	The state/region where your organization is located. This must not be abbreviated.	Mountain View
State/Province	The city where your organization is located.	California
Country	The two-letter ISO code for the country where your organization is located.	US
Email address	An email address to contact your organization.	contact@yourco.com

Create a certificate signing request

Create a certificate signing request that you can submit to a CA to obtain a signed certificate.

Before you begin

- Create a private key using the procedure in [Create a private key](#) on page 159.
- To create certificates that use SAN, you must create a SAN configuration file using the procedure in [Generate a SAN configuration](#) on page 159.

Procedure

1. Use the private key to create a certificate signing request.

Two ways of creating the signing request are shown. One for if you have already prepared a SAN configuration file to specify the alternative server name, another if you have not.

If you are using SAN:

```
openssl req -new -key server.key -config request.conf -out server.csr
```

If you are not, use:

```
openssl req -new -key server.key -out server.csr
```

When creating a signing request, you are asked to supply the Distinguished Name (DN) which comprises a number of fields. Only the Common Name is required and you can accept the defaults for the other parameters.

2. Enter the pass phrase for your private key.
3. At the prompts, enter the fields for the DN for the certificate.

Most fields are optional. However, you must enter a Common Name (CN).

Note

The CN should be a FQDN. Even if you install the certificate on the ECS nodes, you must use an FQDN and all of the IP addresses must be in the alternate names section.

You will see the following prompts:

```
You are about to be asked to enter information that will be
incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished
Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [AU]:US
State or Province Name (full name) [Some-State]:
Locality Name (eg, city) []:
Organization Name (eg, company) [Internet Widgits Pty
Ltd]:Acme
Organizational Unit Name (eg, section) []:
Common Name (e.g. server FQDN or YOUR name) []:*.acme.com
Email Address []:
```

More information on the DN fields are provided in [Distinguished Name \(DN\) fields](#) on page 162.

4. You are prompted to enter an optional challenge password and a company name.

```
Please enter the following 'extra' attributes
to be sent with your certificate request
A challenge password []:
An optional company name []:
```

5. View the certificate.

```
openssl req -in server.csr -text -noout
```

Results

You can submit the certificate signing request to your CA who will return a signed certificate file.

Upload a certificate

You can upload management or data certificates to ECS. Whichever type of certificate you upload, you must authenticate with the API.

- [Authenticate with ECS Management REST API](#) on page 164
- [Upload a management certificate](#) on page 164
- [Upload a data certificate for data access endpoints](#) on page 166

Authenticate with ECS Management REST API

To run ECS Management REST API commands, you must first authenticate with the API service and obtain an authentication token.

Procedure

1. Authenticate with the ECS Management REST API and obtain an authentication token that can be used when using the API to upload or verify certificates.
 - a. Run the following command:

```
export TOKEN=`curl -s -k -v -u <user>:<password> https://$(hostname -i):4443/login 2>&1 | grep X-SDS-AUTH-TOKEN | awk '{print $2, $3}'`
```

The username and password are those used to access the ECS Portal. The `public_ip` is the public IP address of the node.

- b. Verify the token exported correctly.

```
echo $TOKEN
```

Example output:

```
X-SDS-AUTH-TOKEN:
BAAcTGZjUjJ2ZmliYURsUFZzKzhBSVVPQVFDRUUpQMAjAQASHVybjpgzdg9y
YwDlb3M6VmlydHVhbERhdGFZDW50ZXJEYXRhOjcxYjA1ZTgwLTNkNzktND
dmMC04OThhLWI2OTU4NDk1YmVmYgIADTE0NjQ3NTM2MjgzMTIDAC51cm46VG
9rZW46YWMwN2Y0NGYtMjE5OS00ZjA4LTgyM2EtZTAwNTc3ZWl0NDAYAgAC
0A8=
```

Upload a management certificate

You can upload a management certificate which is used to authenticate access to management endpoints, such as the ECS Portal and the ECS Management REST API.

Before you begin

- Ensure that you have authenticated with the ECS Management REST API and stored the token in a variable (`$TOKEN`) as described in [Authenticate with ECS Management REST API](#) on page 164.
- Ensure that the machine that you use has a suitable REST client (such as `curl`) and can access the ECS nodes using the ECS Management REST API.
- Ensure your private key and certificate are available on the machine from which you intend to perform the upload.

Procedure

1. Ensure that your private key does not have a passphrase.

If it does, you can create a copy with the passphrase stripped, by typing the following command:

```
openssl rsa -in server.key -out server_nopass.key
```

2. Upload the keystore for the data path using your private key and signed certificate.

Using curl:

```
curl -svk -H "$TOKEN" -H "Content-type: application/xml" -H
"X-EMC-REST-CLIENT: TRUE" -X PUT -d "<rotate_keycertchain>
<key_and_certificate><private_key>`cat privateKeyFile`
<private_key>`</private_key><certificate_chain>`cat
certificateFile`</certificate_chain></key_and_certificate></
rotate_keycertchain>" https://<ecs_node_address>:4443/vdc/
keystore
```

Using the ECS command line interface (ecscli.py):

```
python ecscli.py vdc_keystore update -hostname <ecs host ip> -
port 4443
-cf <cookiefile> -privateKey privateKeyFile -certificateChain
certificateFile
```

The `privateKeyFile`, for example `<path>/server_nopass.key`, and `certificateFile`, for example `<path>/server.crt`, must be replaced with the path to the key and certificate files.

3. Log in to one of the ECS nodes as the admin user.
4. Verify that the MACHINES file has all nodes in it.

The MACHINES file is used by ECS wrapper scripts that execute commands on all nodes, such as `vipreexec`.

For 2.2.1 and later, the MACHINES file is in `/home/admin`. For ECS versions before 2.2.1, the MACHINES file is in `/root`.

- a. Display the contents of the MACHINES file.

For 2.2.1, use:

```
cat /home/admin/MACHINES
```

For earlier versions, use:

```
cat /root/MACHINES
```

- b. If the MACHINES file does not contain all nodes, recreate it.

For 2.2.1, use:

```
/usr/sbin/getrackinfo -c MACHINES
```

Note

In 2.2.1, access to nodes as root was removed and access is through the admin user account.

The full path to the `getrackinfo` command is specified so that you do not have to use `sudo`. Using `sudo` would make the `MACHINES` file owned by root.

For earlier versions, use:

```
getrackinfo -c /root/MACHINES
```

Verify that the `MACHINES` file now contains all nodes.

5. Restart the `objcontrolsvc` and `nginx` once the management certificates are applied.
 - a. Restart the object service.

```
viprexec -f ~/MACHINES -i 'pidof objcontrolsvc;
kill `pidof objcontrolsvc`; sleep 60; pidof objcontrolsvc'
```

- b. Restart the `nginx` service.

Run the following commands:

```
viprexec 'docker exec -it object-main service nginx restart'
```

Repeat for each node.

After you finish

You can verify the certificate has uploaded correctly using the following procedure: [Verify the management certificate](#) on page 168.

Upload a data certificate for data access endpoints

You can upload a data certificate which is used to authenticate access for the S3, EMC Atmos, or OpenStack Swift protocols.

Before you begin

- Ensure that you have authenticated with the ECS Management REST API and stored the token in a variable (`$TOKEN`). See [Authenticate with ECS Management REST API](#) on page 164.
- Ensure that the machine that you use has a suitable REST client (such as `curl`) and can access the ECS nodes using the ECS Management REST API.
- Ensure your private key and certificate are available on the machine from which you intend to perform the upload.

Procedure

1. Ensure that your private key does not have a pass phrase.

If it does, you can create a copy with the pass phrase stripped, using:

```
openssl rsa -in server.key -out server_nopass.key
```

2. Upload the keystore for the data path using your private key and signed certificate.

```
curl -svk -H "$TOKEN" -H "Content-type: application/xml" -H
"X-EMC-REST-CLIENT: TRUE" -X PUT -d "<rotate_keycertchain>
<key_and_certificate><private_key>`cat privateKeyFile`</
private_key><certificate_chain>`cat
certificateFile`</certificate_chain></key_and_certificate></
rotate_keycertchain>"
https://<ecs_node_address>:4443/object-cert/keystore
```

Using the ECS command line interface (ecscli.py):

```
python ecscli.py keystore update -hostname <ecs host ip> -
port 4443 -cf <cookiefile>
-pkvf privateKeyFile -cvf certificateFile -ss false
```

The `privateKeyFile`, for example `<path>/server_nopass.key`, and `certificateFile`, for example `<path>/server.crt`, must be replaced with the path to the key and certificate files.

3. The certificate is distributed when the `dataheadsvc` is restarted. You can do this with the commands below.

Note

You do not need to restart the services when changing data certificate, `dataheadsvc` is restarted automatically on each node 2 hours from certificate update.

```
ssh admin@<ecs_ip_where_cert_uploaded>
```

```
sudo kill `pidof dataheadsvc`
```

After you finish

You can verify that the certificate has correctly uploaded using the following procedure: [Verify the object certificate](#) on page 169.

Verifying installed certificates

The object certificate and management certificate each has an ECS Management API GET request to retrieve the installed certificate.

- [Verify the management certificate](#) on page 168
- [Verify the object certificate](#) on page 169

Verify the management certificate

The ECS Management REST API provides methods for retrieving the installed certificates.

Before you begin

- Ensure that you have authenticated with the ECS Management REST API and stored the token in a variable (`$TOKEN`). See [Authenticate with ECS Management REST API](#) on page 164.
- If you have restarted services, the certificate is available immediately. Otherwise, you must wait 2 hours to be sure that the certificate is propagated to all nodes.

Procedure

1. Use the `GET /vdc/keystore` method to return the certificate.

Using the `curl` tool, the method can be run by typing the following:

```
curl -svk -H "X-SDS-AUTH-TOKEN: $TOKEN" https://x.x.x.x:4443/vdc/keystore
```

Using the ECS command line interface (`ecscli.py`):

```
python ecscli.py vdc_keystore get -hostname <ecs host ip> -port 4443 -cf <cookiefile>
```

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<certificate_chain><chain>-----BEGIN CERTIFICATE-----
MIIDgjCCAmoCCQCEDeNwcGsttTANBgkqhkiG9w0BAQUFADCBgjELMAkGA1UEBh
MC&#xD;
VVMxCzAJBgNVBAGMAkdBMQwwCgYDVQQHDANBVEwxDDAKBgNVBAoMA0VNQzEMMA
oG&#xD;
A1UECwwDRU5HMQ4wDAYDVQQDDAVjaHJpczEsMCoGCSqGS1b3DQEJARYdY2hyaX
N0&#xD;
b3BoZXIuZ2hva2FzaWFuQGVtYy5jb20wHhcNMTYwNjAxMTg0MTIyWhcNMTCwNj
Ay&#xD;
MTg0MTIyWjCBgjELMAkGA1UEBhMCMVVMxCzAJBgNVBAGMAkdBMQwwCgYDVQQHDA
NB&#xD;
VEwxDDAKBgNVBAoMA0VNQzEMMAoGA1UECwwDRU5HMQ4wDAYDVQQDDAVjaHJpcz
Es&#xD;
MCoGCSqGS1b3DQEJARYdY2hyaXN0b3BoZXIuZ2hva2FzaWFuQGVtYy5jb20wgG
Ei&#xD;
MA0GCSqGS1b3DQEBAAUAA4IBDwAwggEKAoIBAQBd9WtdcW5HJpIDOUtB7o7ic0
RK&#xD;
dwA4dY/
nJXrk6Ikae5zDW08XH4noQNhAu8FnEwS5kjtBK1hgI2GEFBtLkIH49AUp&#xD;
c4KrMmotDmbCeHvOhNCqBLZ5JM6DACfO/elHpb2hgBENTd6zyp7mz/
7MUf52s9Lb&#xD;
x5pRRcPl1LDw3s15iodZ5GL8pRT62puJVK1do9mPfMoL22woR3YB2+
+AkSdAgEFH&#xD;
1XLI sFGkBsEJObbDBoEMEjEIivnTRPiyocyWki6gfLh50u9Y9B2GRzLazIlgNi
Es&#xD;
L/vyyrHcwOs4up9QqhAlvMn3A101VF+OH0omQECSchBdsc/R/
Bc35FAEVdmTAgMB&#xD;
AAEwDQYJKoZIhvcNAQEFBQADggEBAAyYcvJtEhOq
+n87wukjPMGc719n7rgvaTmo&#xD;
tzpQhtt6kFoSBO7p//
76DNzXRXhBDADwpUGG9S4tgHChAFu9DpHFzvnjNGGw83ht&#xD;
qcJ6JYgB2M31OQAssgW4fU6VD2bfQbGRWKy9G1rPYGVsmKQ59Xeuvf/
```



```
cWvplkwW2&#xD;
bKnZmAbWEfElcEOqt+5m20qGPcf45B7DPp2J
+wVdDD7N8198Jj5HJBjt3T3aUEwj&#xD;
kvnPx1PtFM9YORKXFX2InF3UOdMs0zJUkhBZT9cJ0gASi1w0vEnx850secu1CP
LF&#xD;
WB9G7R5qHWOX1kbAVPuFN01Tav+yrr8RgTawAcsv9LhkTTOUcqI=&#xD;
-----END CERTIFICATE-----</chain></certificate_chain>
```

2. You can verify the certificate using `openssl` on all nodes.

```
openssl s_client -showcerts -connect <NODE_IP>:<port>
```

Note

The management port is 4443.

For example:

```
openssl s_client -showcerts -connect 10.1.2.3:4443
```

Verify the object certificate

The ECS Management REST API provides methods for retrieving the installed certificates.

Before you begin

- Ensure that you have authenticated with the ECS Management REST API and stored the token in a variable (`$TOKEN`). See [Authenticate with ECS Management REST API](#) on page 164.
- If you have restarted services, the certificate will be available immediately. Otherwise, you need to wait 2 hours to be sure that the certificate has propagated to all nodes.

Procedure

1. Use the GET `/object-cert/keystore` method to return the certificate.

Using the curl tool, the method can be run by typing the following:

```
curl -svk -H "X-SDS-AUTH-TOKEN: $TOKEN" https://x.x.x.x:4443/object-cert/keystore
```

Using the ECS command line interface (`ecscli.py`):

```
python ecscli.py keystore show -hostname <ecs host ip> -port 4443 -cf <cookiefile>
```

2. You can verify the certificate using `openssl` on all nodes.

```
openssl s_client -showcerts -connect <NODE_IP>:<port>
```

Note

Ports are: s3: 9021, Atmos: 9023, Swift: 9025

Example:

```
openssl s_client -showcerts -connect 10.1.2.3:9021
```

CHAPTER 16

Locking remote access to nodes

- [Locking remote access to nodes](#)..... 172
- [Lock and unlock nodes](#)..... 173

Locking remote access to nodes

Use the ECS Portal to lock remote access to nodes.

Access types

ECS can be configured in the following ways:

1. Using the ECS Portal or the ECS Management API.
2. By directly connecting to a node through the management switch with a service laptop and using SSH or the CLI to directly access the node's operating system.
3. By remotely connecting to a node over the network using SSH or the CLI to directly access the node's operating system.

Node locking provides another layer of security against remote node access from all accounts. Without node locking, any privileged node-level account, such as the `admin`, `service`, or `emc` accounts, can remotely access nodes at any time to collect data, configure hardware, and run Linux commands. If all the nodes in a cluster are locked, then remote access can be planned and scheduled for a defined window minimizing the opportunity for unauthorized activity.

Using the ECS Portal or the ECS Management API, you can lock selected nodes in a cluster or all the nodes in the cluster. Doing so only affects the ability to remotely access (SSH to) the locked nodes. Locking does not change the way the ECS Portal and ECS Management APIs access nodes and it does not affect the ability to directly connect to a node.

Lock Admin

To lock and unlock nodes requires the Lock Admin user. The Lock Admin is a pre-provisioned local user called `emcsecurity`. Lock Admins can only change their passwords and lock and unlock nodes. The Lock Admin role cannot be assigned to another user.

System Admins and System Monitors can view the lock status of the nodes.

For instructions on locking and unlocking nodes see: [Lock and unlock nodes](#).

Maintenance

If node maintenance using remote access is periodically required, you can unlock a single node to allow remote access to the entire cluster using SSH with the `admin` or `emc` account. Once the authorized user successfully logs into the unlocked node using SSH, the user can SSH from that node to any other node in the cluster by way of the private network.

You will also need to unlock a node to remotely use commands that provide OS-level read-only diagnostics.

Auditing

Node lock and unlock events are captured in audit logs and also sent to Syslog. Errors from lock or unlock attempts are also logged.

ECS Management API

The following APIs allow you to manage node locks.

Resource	Description
GET /vdc/nodes	Gets the data nodes that are currently configured in the cluster

Resource	Description
GET /vdc/lockdown	Gets the locked/unlocked status of a VDC
PUT /vdc/lockdown	Sets the locked/unlocked status of a VDC
PUT /vdc/nodes/{nodeName}/lockdown	Sets the Lock/unlock status of a node
GET /vdc/nodes/{nodeName}/lockdown	Gets the Lock/unlock status of a node

Lock and unlock nodes

Use the portal to lock and unlock remote SSH access to ECS nodes.

Before you begin

This task can only be done by the Lock Admin (login: emcsecurity).

Locking a node only prevents remote access to the operating system of the node by SSH or the CLI. Locking or unlocking a node has no affect on ECS Portal or REST Management API functions or on directly connecting to a node locally and then using SSH or the CLI.


Procedure

1. Login in as emcsecurity.

If this is the first log in from this account, you will be required to change the password and re-login.

2. From the left-hand navigation, select **Settings > Platform Locking**.

The screen lists the nodes in the cluster and displays their lock status.



The node states are:

- Unlocked: Displays an open green lock icon and the **Lock** action button.
- Locked: Displays a closed red lock icon and the **Unlock** action button.
- Offline: Displays the circle-with-slash icon and no action button because the node is unreachable and the lock state cannot be determined.

3. Choose:

Option	Description
Lock	To lock an unlocked node. Any user who is currently remotely logged in by SSH or CLI will have about five minutes to exit before their session will be terminated. An impending shutdown message appears on the user's terminal screen.
Unlock	To unlock a locked node. A privileged user will now be able to remotely login to the node by SSH or the CLI after a few minutes.
Lock the VDC	This convenience feature locks all unlocked nodes in the VDC as long as they are online. It does not set a state where any new or offline node will be automatically locked once detected.

PART 3

Monitor

[Chapter 17, "Monitoring basics"](#)

[Chapter 18, "Monitor metering"](#)

[Chapter 19, "Monitor events"](#)

[Chapter 20, "Monitor capacity utilization"](#)

[Chapter 21, "Monitor traffic metrics"](#)

[Chapter 22, "Monitor hardware health"](#)

[Chapter 23, "Monitor node and process health"](#)

[Chapter 24, "Monitor chunk summary"](#)

[Chapter 25, "Monitor erasure coding"](#)

[Chapter 26, "Monitor recovery status"](#)

[Chapter 27, "Monitor disk bandwidth"](#)

[Chapter 28, "Monitor geo-replication"](#)

[Chapter 29, "Service logs"](#)

CHAPTER 17

Monitoring basics

- [Using monitoring pages](#).....178

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 13 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 14 Open Filter panel with criteria selected



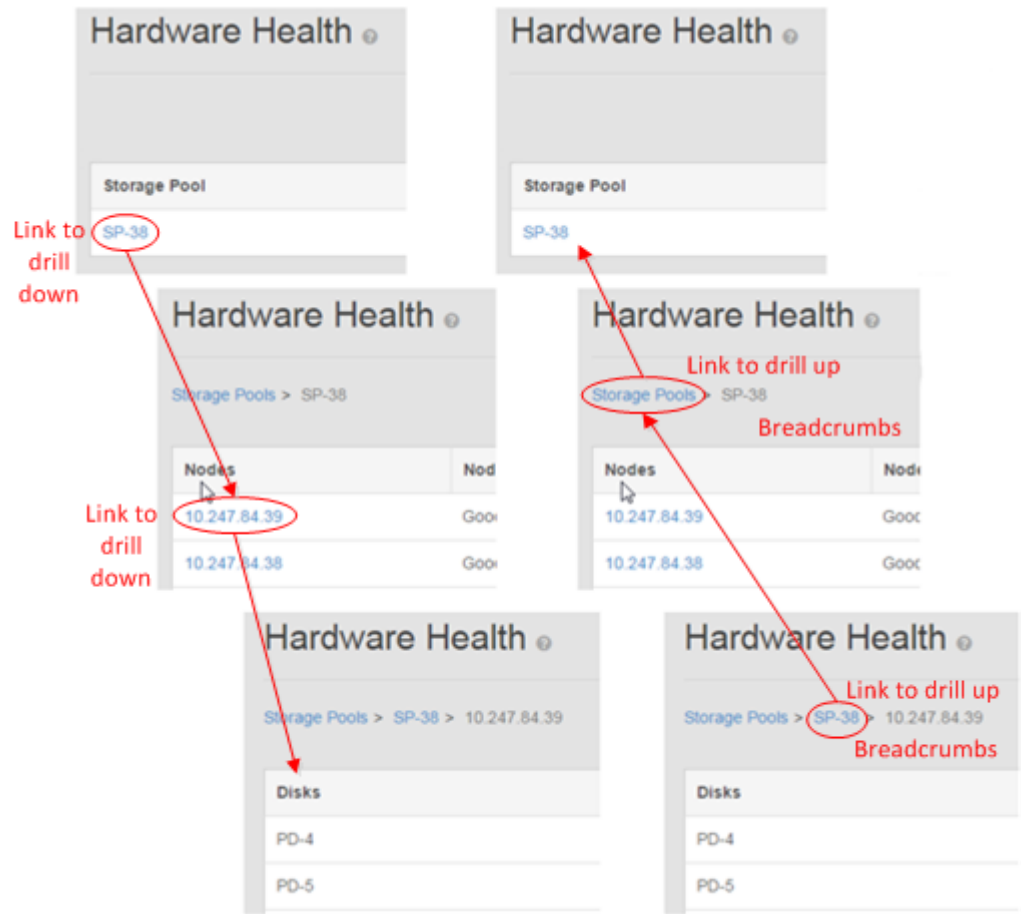
When the Filter panel closes, a summary of the applied filter displays along with a Clear Filter command and a Refresh command.

Figure 15 Closed Filter panel showing summary of applied filter



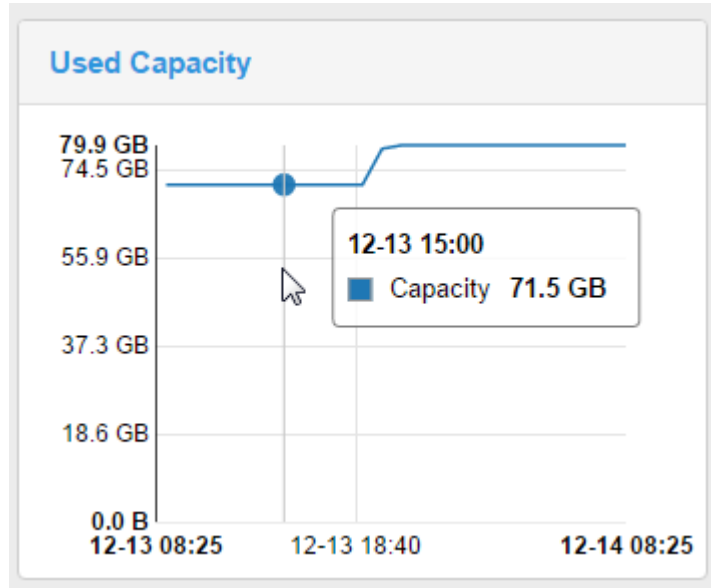
Highlighted text in a table row indicates a link to a detail display. Selecting the link drills down to the next level of detail. On drill down displays, a path string shows your current location in the sequence of drill down displays. This path string is called a breadcrumb trail or breadcrumbs for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

Figure 16 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 17 History chart with active cursor



CHAPTER 18

Monitor metering

- [Monitor metering data](#)..... 182

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 183.

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](#). 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 18 Metering page with criteria selected

The screenshot shows the 'Metering' interface. It includes a 'Date Time Range' section with a 'Custom' dropdown, 'From' and 'To' time pickers set to 2015-12-13 09:35 and 2015-12-14 09:35 respectively. Below is a 'Select Namespace' section with a search bar and a list containing 'ns-base'. To the right is a 'Selected' table with columns 'Namespace' and 'Bucket', showing 'ns-base' and 'bkl-base'. At the bottom left is an 'Apply' button.

- 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 9 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 9 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 19

Monitor events

- [About event monitoring](#)..... 186
- [Monitor audit data](#)..... 186
- [Monitor alerts](#)..... 187

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.

See [Appendix A: Audit messages](#).

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' page with the 'Alerts' tab selected. A 'Filter' panel is visible, allowing users to filter events by 'Date Time Range' (set to 'Last Month'), 'From' (2015-11-10), 'To' (2015-12-10), and 'Namespace' (ns1). 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.

See [Appendix A: Alert messages](#).

Alert message Severity labels have the following meanings:

- **CRITICAL:** Messages about conditions that require immediate attention.
- **ERROR:** Messages about error conditions that report either a physical failure or a software failure.
- **WARNING:** Messages about less than optimal conditions.
- **INFO:** Routine status messages.

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 as described in the table below:

Table 10 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
NAMESPACE_HARD_QUOTA_EXCEEDED	Raised when the quota on a namespace is exceeded.
DTSTATUS_RECENT_FAILURE	Raised when the status of a data table is bad.
CHUNK_NOT_FOUND	Raised when chunk data is not found.
FILE_NOT_FOUND	Raised when a file is not found.
FILE_DATA_CORRUPTED	Raised when a file contains corrupted data.
VPOOL_FREE_SPACE_CRITICAL	Raised when the storage pool is 90% or more full.

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

- Date Time Range:** A dropdown menu set to 'Last 2 Months'.
- From:** A date and time selector set to '2015-10-10 13:30'.
- To:** A date and time selector 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 panel are two buttons: 'Apply' (in blue) and 'Clear Filter'. Below the buttons is a table with the following columns: Severity, Description, Type, Namespace, Timestamp, and Actions.

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

CHAPTER 20

Monitor capacity utilization

- [Monitor capacity](#)..... 190

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 190. 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](#). 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 178.

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 190
- [Node Capacity Utilization](#) on page 191
- [Disk Capacity Utilization](#) on page 193

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

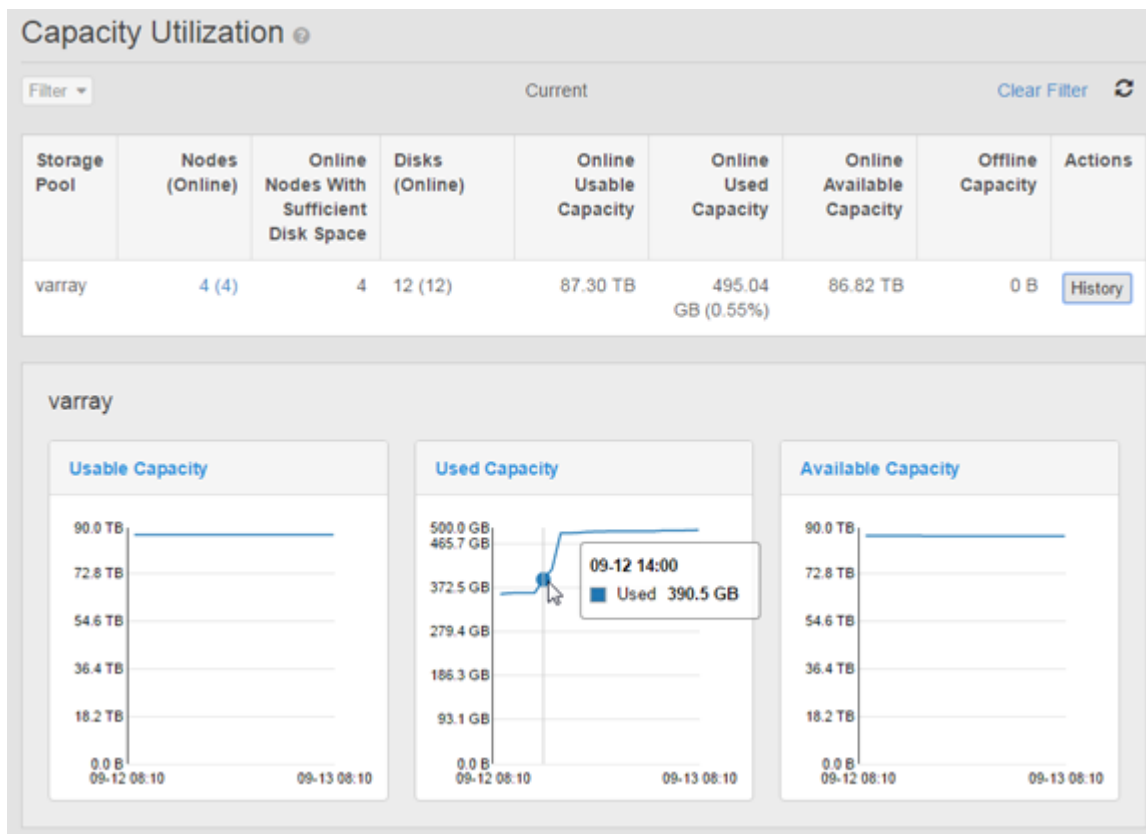
Table 11 Capacity Utilization: Storage Pool

Attribute	Description
Storage Pool	Name of the storage pool.
Nodes (Online)	Number of nodes in the storage pool followed by the number of those nodes currently online. Click node number to open: Node Capacity Utilization on page 191.
Online Nodes with Sufficient Disk Space	Number of online nodes that have sufficient disk space to accept new data. If too many disks are too full to accept new data, the performance of the system may be impacted.
Disks (Online)	Number of disks in the storage pool followed by the number of those disks that are currently online.

Table 11 Capacity Utilization: Storage Pool (continued)

Attribute	Description
Online Usable Capacity	Total usable capacity of the storage pool that is currently online. This is the total of the capacity already used and the capacity still free for allocation.
Online Used Capacity	Used online capacity in the storage pool.
Online Available Capacity	Online capacity available for use.
Offline Capacity	The total capacity that is currently offline.
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 storage pool capacity utilization table is shown below.



Node Capacity Utilization

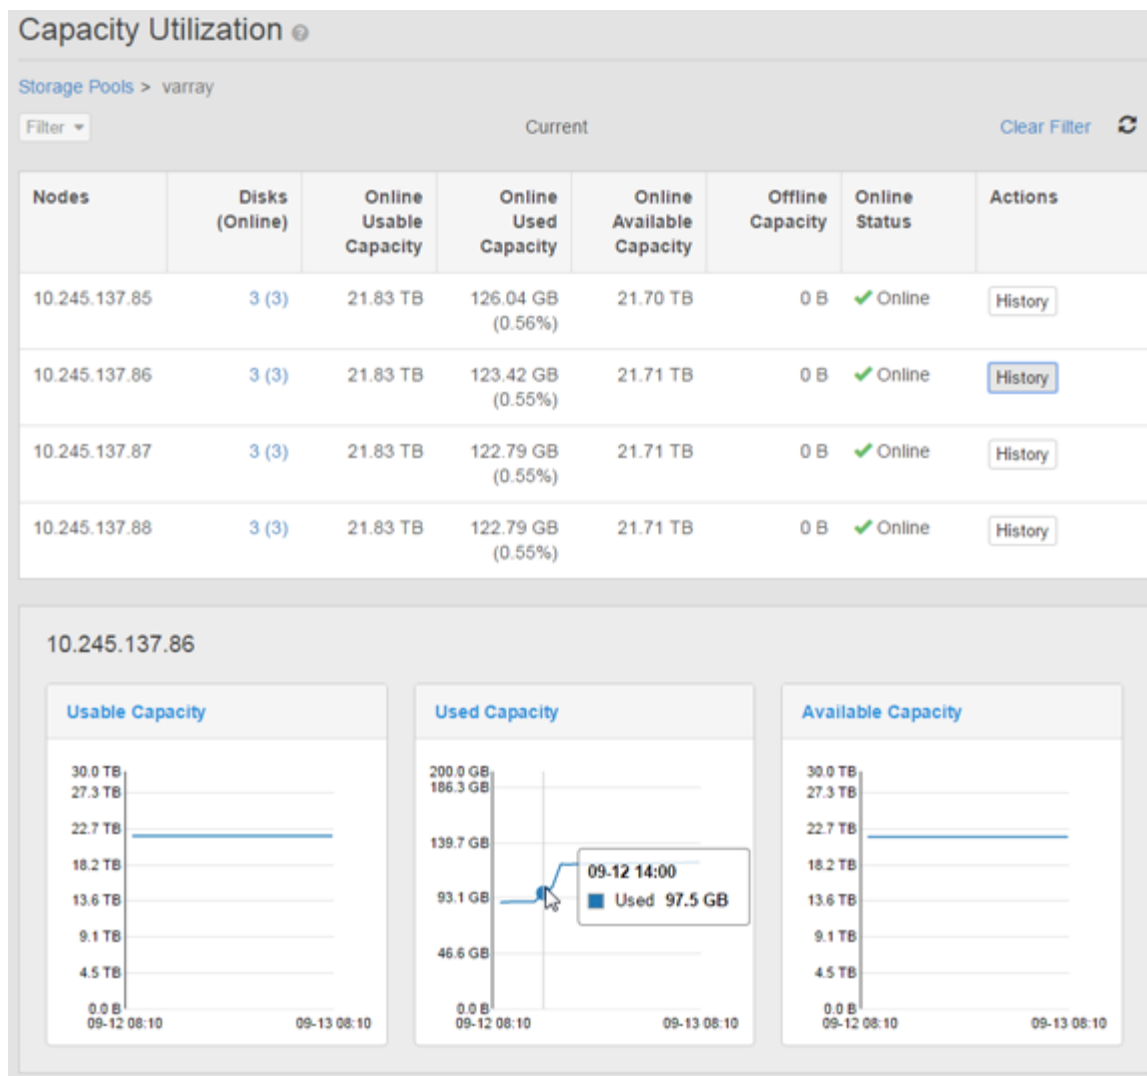
Table 12 Capacity Utilization: Node

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 193

Table 12 Capacity Utilization: Node (continued)

Attribute	Description
Online Usable Capacity	Total usable online capacity provided by the online disks within the node. This is the total of the capacity already used and the capacity still free for allocation.
Online Used Capacity	Online capacity used within the node.
Online Available Capacity	Remaining online capacity available in the node.
Offline Capacity	Total capacity of the node that is currently offline.
Online Status	Indicates whether the node is online or offline. A 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.

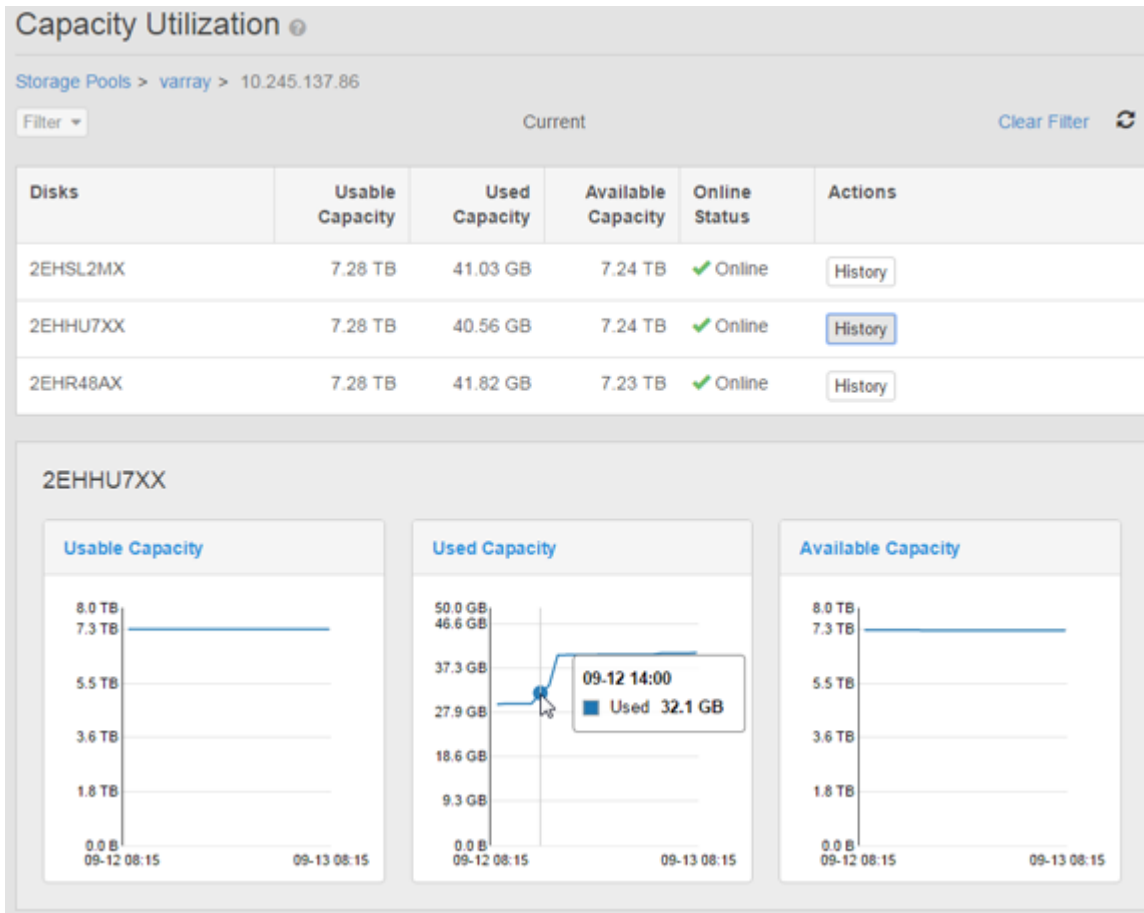


Disk Capacity Utilization

Table 13 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.
Online Status	Indicates whether the disk is online or offline. The 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.



Monitor capacity utilization

CHAPTER 21

Monitor traffic metrics

- [Monitor network traffic](#)..... 196

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.

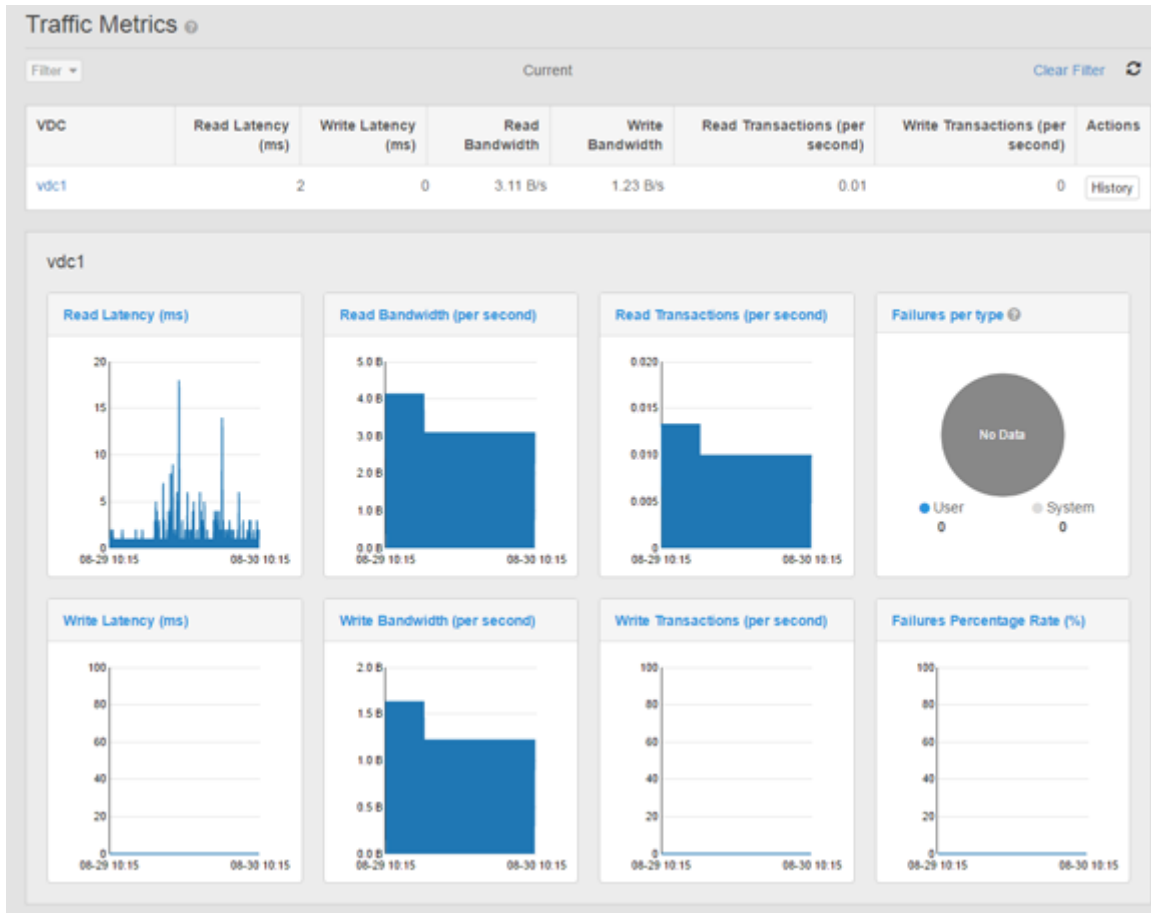
Table 14 Network traffic metrics

Metric label	Description
VDC	Click a VDC name to see traffic metrics by node.
Read Latency (ms)	Average latency for reads in milliseconds.
Write Latency (ms)	Average latency for writes in milliseconds.
Read Bandwidth	Bandwidth for reads.
Write Bandwidth	Bandwidth for writes.
Read Transactions (per second)	Read transactions per second.
Write Transactions (per second)	Write transactions per second.
History	<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>

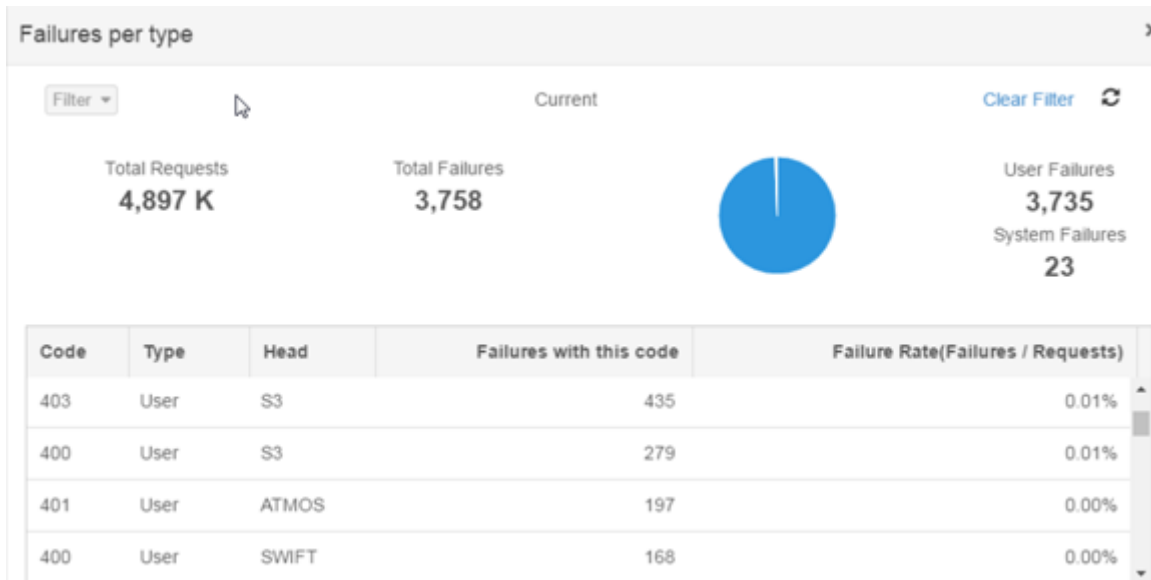
Procedure

1. 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 19 Network traffic charts for a VDC



5. Select the Failures by type panel.



ECS Failures by type lists the error code with highest frequency on top. The user has the ability to sort on any of the columns.

User errors are known error types originating from the heads (typically an HTTP error code of 4xx).

System errors are failed requests associated with hardware or service errors (typically an HTTP error code of 5xx).

CHAPTER 22

Monitor hardware health

- [Monitor hardware](#)..... 200

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 as well as two more states:

- **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.
- **Missing:** The disk is a known disk that is currently unreachable. The disk may be transitioning between states, disconnected, or pulled.
- **Removed:** The disk is one that the system has completed recovery on and removed from the storage engine's list of valid disks.

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 20 Hardware Health for nodes

The screenshot shows the 'Hardware Health' page for a storage pool named 'varray'. The page displays a table with four columns: 'Nodes', 'Node Status', 'Online Disks', and 'Offline Disks'. All four nodes listed have a 'Good' status, 3 online disks, and 0 offline disks.

Nodes	Node Status	Online Disks	Offline Disks
10.245.137.85	✓ Good	3	0
10.245.137.86	✓ Good	3	0
10.245.137.87	✓ Good	3	0
10.245.137.88	✓ Good	3	0

CHAPTER 23

Monitor node and process health

- [Monitor node and process health](#)..... 202

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 15 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.

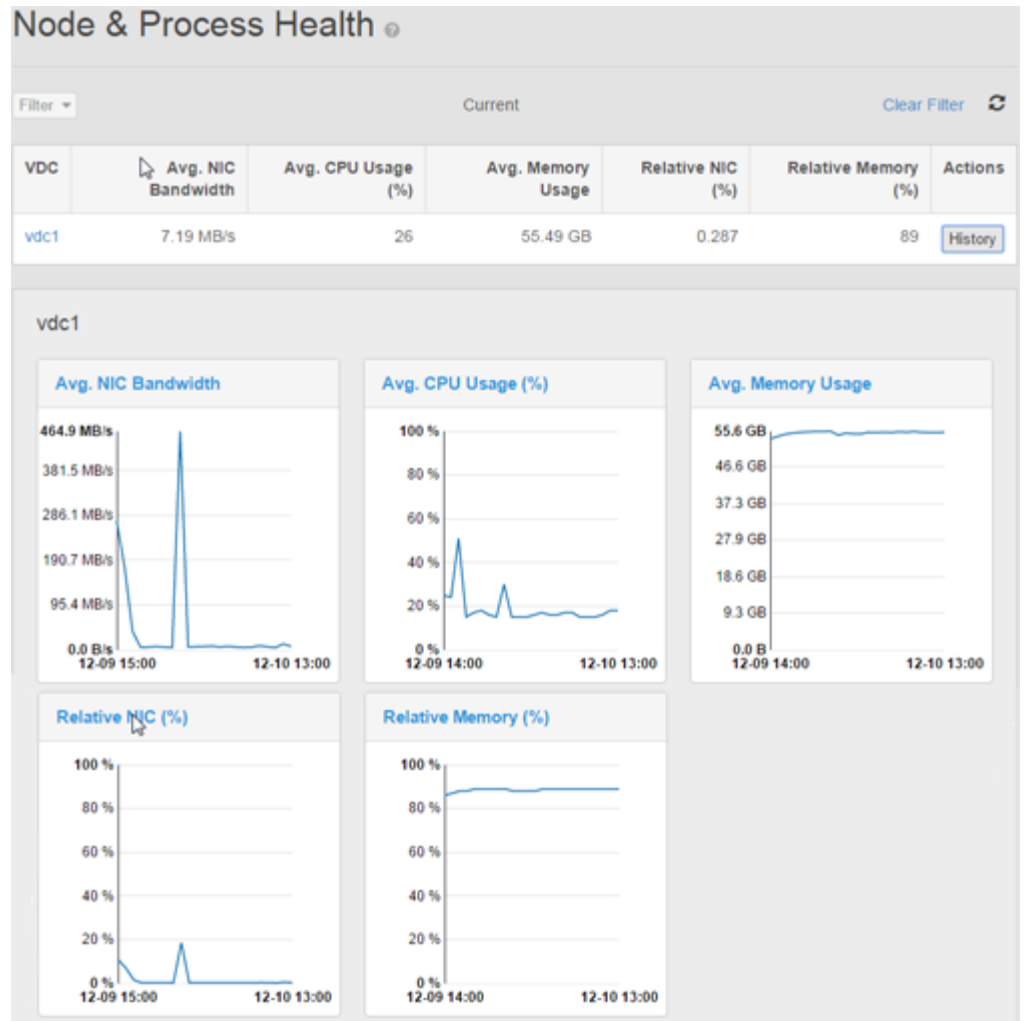
Table 15 VDC, node, and process health metrics (continued)

Metric label	Level	Description
Last Restart	Process	The last time the process restarted on the node.
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 21 Node & Process Health



CHAPTER 24

Monitor chunk summary

- [Monitor chunks](#).....206

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 16 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.

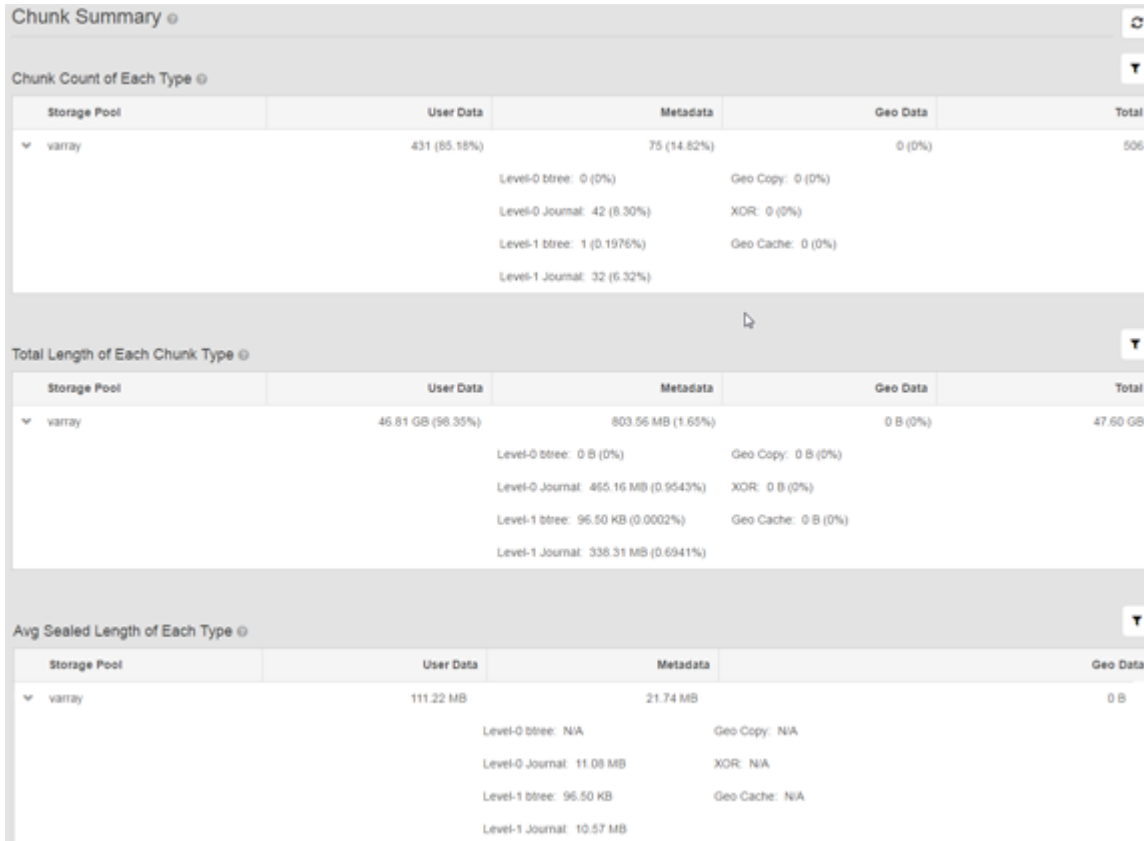
Table 17 Chunk metrics

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). There are three types: <ul style="list-style-type: none"> • Geo copies: chunks replicated from other sites. • XOR: chunks resulting from the XOR-ing of original chunks • Geo Cache: a temporary chunk needed to facilitate read operations when the site does not have an original or copy of a required chunk.
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.

Table 17 Chunk metrics (continued)

Metric label	Description
	This field provides relevant data for the XOR chunks in the storage pool.
Total	The total number of chunks in the storage pool.

Figure 22 Chunk Summary



Monitor chunk summary

CHAPTER 25

Monitor erasure coding

- [Monitor erasure coding](#).....210

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.

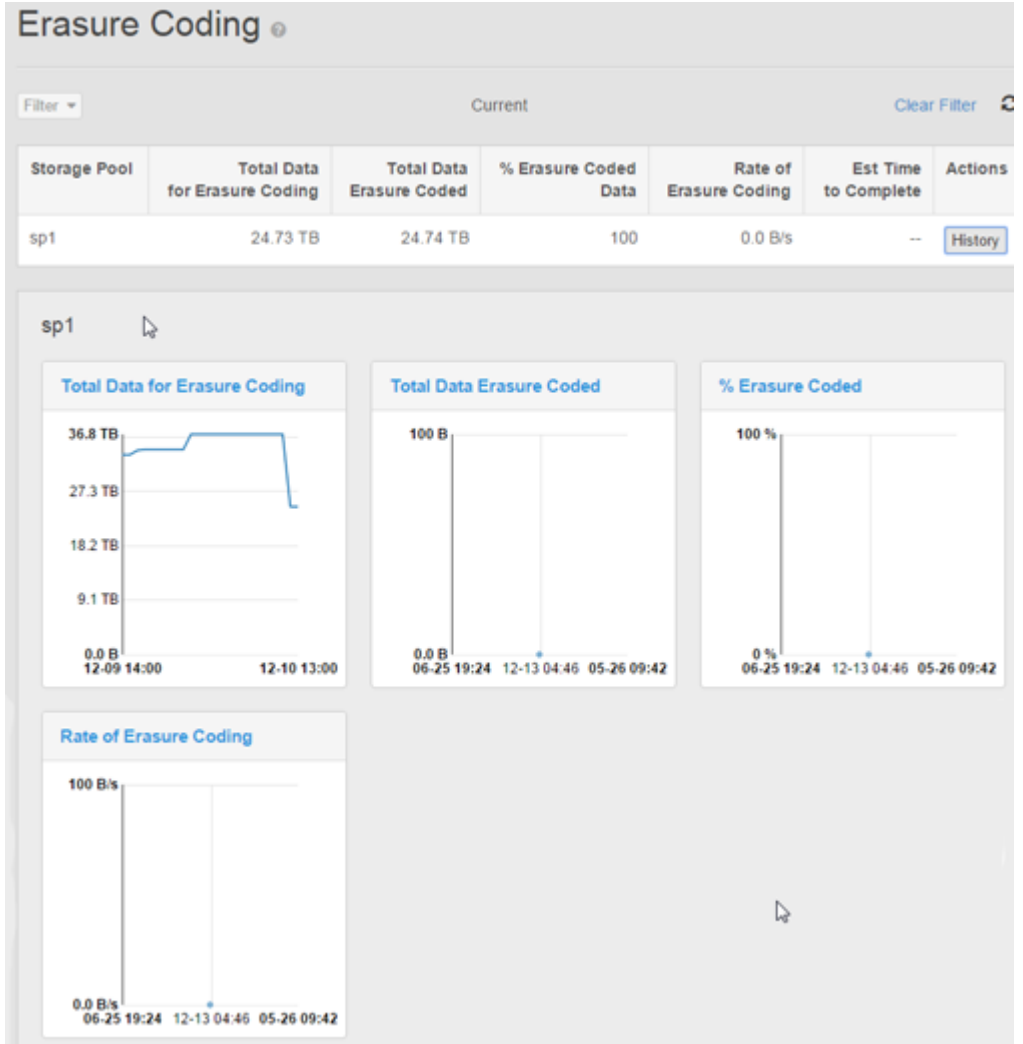
Table values represent current values when the Current Filter is selected, or average values of the metric over the period selected in the filter.

Table 18 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	<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. Select **Monitor > Erasure Coding**.
2. Locate the table row for the target storage pool.
3. Select the **History** button.



Monitor erasure coding

CHAPTER 26

Monitor recovery status

- [Monitor recovery status](#)..... 214

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.

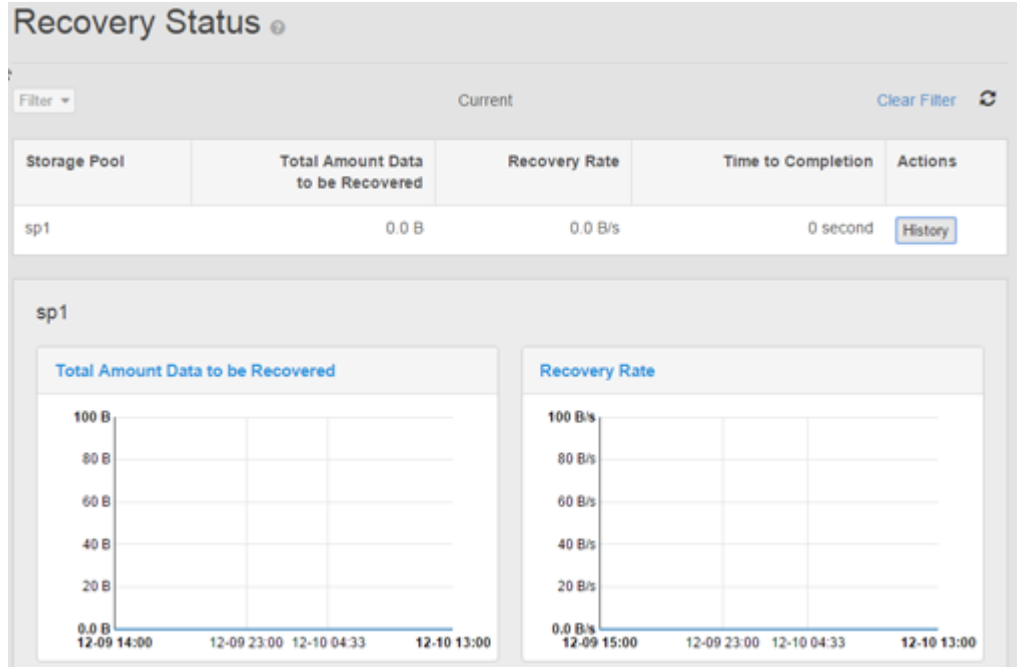
Table values represent current values when the Current Filter is selected, or average values of the metric over the period selected in the filter.

Table 19 Recovery metrics

Column	Description
Storage Pool	Lists each storage pool in the local zone.
Total Amount Data to be Recovered	<p>With the Current filter selected, this is the logical size of the data yet to be recovered. When a historical period is selected as the filter, the meaning of Total Amount Data to be Recovered is the average amount of data pending to be recovered during that selected period of time.</p> <p>For example, if the first hourly snapshot of the data showed 400 GB of data that had to be recovered in a historical time period and every other snapshot showed 0 GB waiting to be recovered, the value of this field would be 400 GB divided by the total number of hourly snapshots in the period.</p>
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	<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. Select **Monitor > Recovery Status**.
2. Locate the table row for the target storage pool.
3. Select the **History** button.



Monitor recovery status

CHAPTER 27

Monitor disk bandwidth

- [Monitor disk bandwidth](#)..... 218

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.

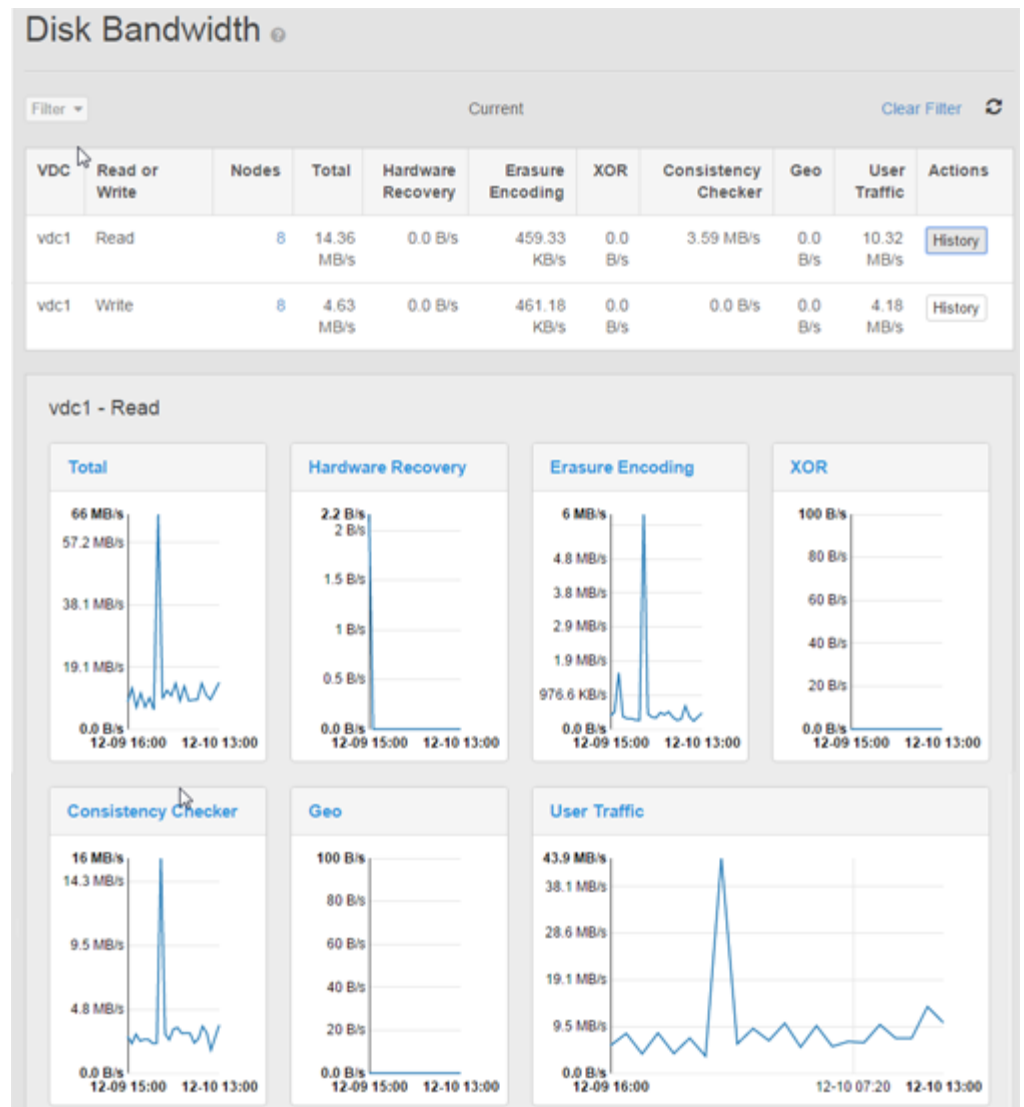
Table 20 Disk bandwidth metrics

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	<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. 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 23 Disk Bandwidth



Monitor disk bandwidth

CHAPTER 28

Monitor geo-replication

- [Introduction to Geo-replication monitoring](#)..... 222
- [Monitor geo-replication: Rate and Chunks](#)..... 222
- [Monitor geo-replication: Recovery Point Objective \(RPO\)](#)..... 223
- [Monitor geo-replication: Failover Processing](#)..... 224
- [Monitor geo replication: Bootstrap Processing](#)..... 225

Introduction to Geo-replication monitoring

Describes the four types of geo-replication monitoring

Geo-replication monitoring includes four different pages:

- [Rate and Chunks](#)
- [Recovery Point Objective \(RPO\)](#)
- [Failover](#)
- [Bootstrap Processing](#)

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 21 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 24 Geo replication: Rate and Chunks



Monitor geo-replication: Recovery Point Objective (RPO)

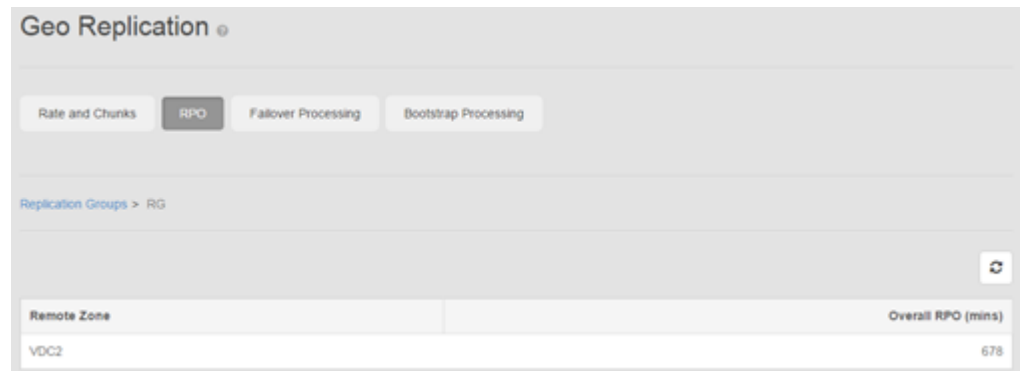
Describes the table fields found in the ECS Portal **Monitor > Geo-Replication > RPO** page.

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 22 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 25 RPO



Monitor geo-replication: Failover Processing

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

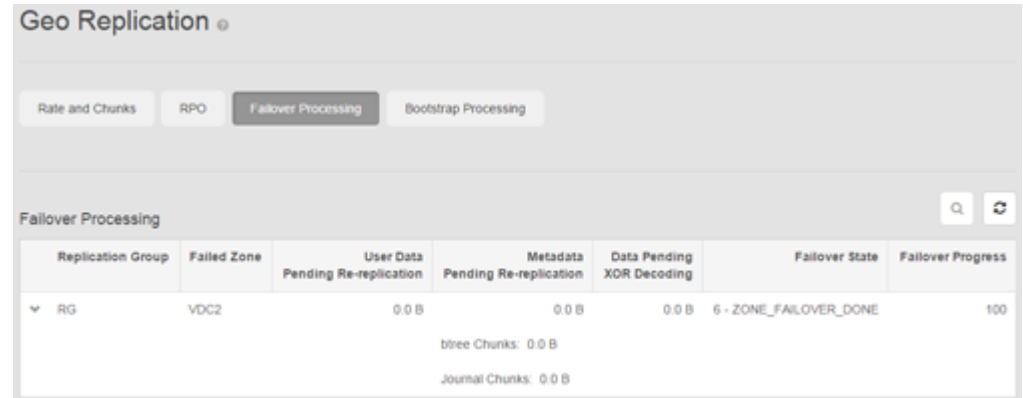
The Failover Processing page provides metrics on the process to re-replicate data following permanent failure of a remote zone.

Table 23 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.

Table 23 Failover columns (continued)

Field	Description
Failover Progress	A percentage indicator for the overall status of the failover process.

Figure 26 Failover

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.

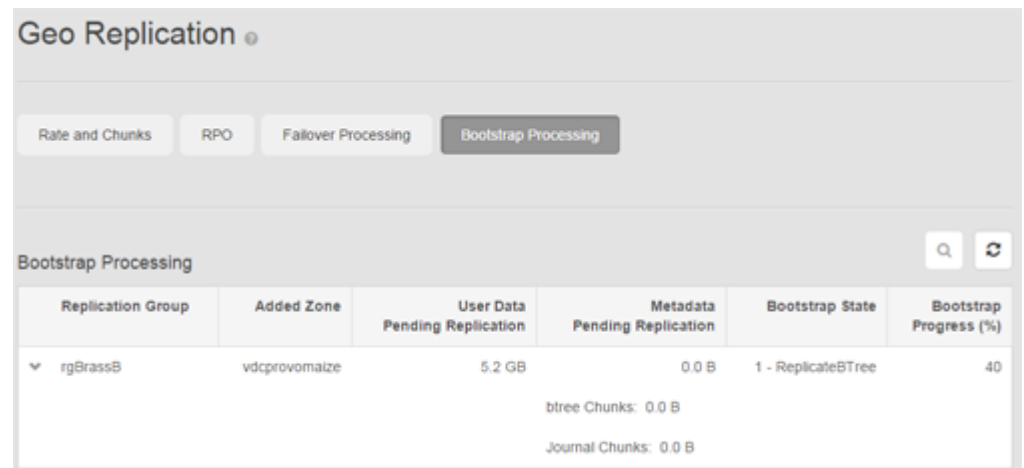
Table 24 Bootstrap Processing columns

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

Table 24 Bootstrap Processing columns (continued)

Column	Description
	<p>chunks are in an acceptable state has completed successfully.</p> <ul style="list-style-type: none"> • 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.
Bootstrap Progress (%)	The completion percent of the entire bootstrap process.

Figure 27 Bootstrap processing



CHAPTER 29

Service logs

- [Service logs](#).....228
- [ECS service log locations](#)..... 228

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.
- `casvc*.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 `casvc`, `objheadsvc` and `hdfssvc` services are combined into `dataheadsvc`.

APPENDIX A

Audit and Alert Messages

- [Audit messages](#)..... 230
- [Alert messages](#).....236

Audit messages

List of the audit messages used by ECS.

Table 25 ECS audit messages

Service	Audit item	Audit message
Alert	sent_alert	Alert \"\${alertMessage}\" with symptom code \${symptomCode} triggered
Auth Provider	new_authentication_provider_added	New authentication provider \${resourceId} added
Auth Provider	authentication_provider_deleted	Authentication provider \${resourceId} deleted
Auth Provider	authentication_provider_updated	Existing Authentication provider \${resourceId} updated
Bucket	bucket_created	Bucket \${resourceId} has been created
Bucket	bucket_deleted	Bucket \${resourceId} has been deleted
Bucket	bucket_updated	Bucket \${resourceId} has been updated
Bucket	bucket_ACL_set	Bucket \${resourceId} ACLs have changed
Bucket	bucket_owner_changed	Owner of \${resourceId} bucket has changed
Bucket	bucket_versioning_set	Versioning has been enabled on \${resourceId} bucket
Bucket	bucket_versioning_unset	Versioning has been suspended on \${resourceId} bucket
Bucket	bucket_versioning_source_set	Bucket \${resourceId} versioning source set
Bucket	bucket_metadata_set	Metadata on \${resourceId} bucket has been changed
Bucket	bucket_head_metadata_set	Bucket \${resourceId} head metadata set
Bucket	bucket_expiration_policy_set	Bucket \${resourceId} expiration policy has updated
Bucket	bucket_expiration_policy_deleted	Bucket \${resourceId} expiration policy has been deleted

Table 25 ECS audit messages (continued)

Service	Audit item	Audit message
Bucket	bucket_cors_config_set	Bucket \${resourceId} CORS rules have been changed
Bucket	bucket_cors_config_deleted	Bucket \${resourceId} CORS rules have been deleted
Bucket	notification_size_exceeded_on_bucket	Notification size has been exceeded on \${resourceId} bucket
Bucket	block_size_exceeded_on_bucket	Block size has been exceeded on \${resourceId} bucket
Bucket	bucket_set_quota	Bucket \${resourceId} quota has been updated with notification size as \${notificationSize} and block size as \${blockSize}
Cluster	cluster_set	Cluster id \${resourceId} has been set
License	user_added_license	License \${resourceId} has been added
License	managed_capacity_exceeded	Managed capacity has exceeded licensed \${resourceId} capacity
License	license_expired	License \${resourceId} has expired
Local user	domain_group_mapping_created	Domain group \${resourceId} to \${roles} role(s) mapping is added
Local user	domain_group_mapping_created_no_roles	Domain group \${resourceId} without role mappings is added
Local user	domain_group_mapping_updated	Domain group \${resourceId} roles mapping is changed to \${roles} role(s)
Local user	domain_group_mapping_updated_no_roles	All roles of domain group \${resourceId} mapping have been removed
Local user	domain_user_mapping_created	Domain user \${resourceId} to \${roles} role(s) mapping is added
Local user	domain_user_mapping_created_no_roles	Domain user \${resourceId} without role mappings is added
Local user	domain_user_mapping_deleted	Domain user \${resourceId} mapping is removed

Table 25 ECS audit messages (continued)

Service	Audit item	Audit message
Local user	domain_user_mapping_updated	Domain user \${resourceId} role mapping is changed to \${roles} role(s)
Local user	domain_user_mapping_updated_no_roles	All roles of domain user \${resourceId} mapping have been removed
Local user	local_user_created	Management user \${resourceId} with \${roles} role(s) has been created
Local user	local_user_created_no_roles	Management user \${resourceId} without roles has been created
Local user	local_user_deleted	Management user \${resourceId} has been deleted
Local user	local_user_password_changed	Credential of management user \${resourceId} has changed
Local user	local_user_updated	Roles of management user \${resourceId} have been changed to \${roles}
Local user	local_user_roles_updated_no_roles	All roles of management user \${resourceId} have been removed
Locked	vdc_lock_successful	VDC lock was successful
Locked	vdc_lock_failed	VDC lock failed
Locked	node_lock_successful	Lock successful for node \${resourceId}
Locked	node_lock_failed	Lock failed for node \${resourceId}
Locked	node_unlock_successful	Unlock successful for node \${resourceId}
Locked	node_unlock_failed	Unlock failed for node \${resourceId}
Login	login_successful	User \${resourceId} logged in successfully
Login	login_failed	User \${resourceId} failed to login
Login	user_token_logout	User logged out token \${resourceId}
Login	user_logout	All user tokens have logged out

Table 25 ECS audit messages (continued)

Service	Audit item	Audit message
Namespace	block_size_exceeded_on_namespace	Block size has been exceeded on \${resourceId} namespace
Namespace	namespace_admin_group_mappings_updated	Namespace \${resourceId} admin group mappings updated to following groups: \${groups}
Namespace	namespace_admin_group_mappings_updated_no_groups	Namespace \${resourceId} admin groups mappings updated to an empty list
Namespace	namespace_admin_user_mappings_updated	Namespace \${resourceId} admin mappings updated to following users: \${admins}
Namespace	namespace_admin_user_mappings_updated_no_admins	Namespace \${resourceId} admin mappings updated to an empty list
Namespace	namespace_created	Namespace \${resourceId} has been created
Namespace	namespace_deleted	Namespace \${resourceId} has been deleted
Namespace	namespace_updated	Namespace \${resourceId} has been updated
Namespace	notification_size_exceeded_on_namespace	Notification size has been exceeded on \${resourceId} namespace
NFS	ugmapping_created	\${type} mapping \${ugMappingName} --> \${resourceId} has been created
NFS	ugmapping_deleted	\${type} mapping \${ugMappingName} --> \${resourceId} has been deleted
NFS	export_created	Export with export path \${exportPath} has been created
NFS	export_deleted	Export with export path \${exportPath} has been deleted
NFS	export_updated	Export with export path \${exportPath} has been updated
Replication Group	replication_group_created	Replication Group \${resourceId} has been created

Table 25 ECS audit messages (continued)

Service	Audit item	Audit message
Replication Group	replication_group_updated	Replication Group \${resourceId} has been updated
Security	command_exec_insufficient_permission	Attempt to execute a command \${command} from \${host} without right permissions
SNMP	snmp_v2_target_created	SNMP target \${snmpTarget} with Community '\${community}' is added
SNMP	snmp_v3_target_created	SNMP target \${snmpTarget} with Username '\${username}', Authentication('\${authProtocol}') and Privacy('\${privProtocol}')
SNMP	snmp_target_deleted	SNMP target \${snmpTarget} is deleted
SNMP	snmp_engineid_updated	SNMP agent EngineID is set to \${engineid}
SNMP	snmp_v2_target_updated	SNMP target \${oldSnmpTarget} is updated as \${newSnmpTarget} with Community string \${community}
SNMP	snmp_v3_target_updated	SNMP target \${oldSnmpTarget} is updated as \${newSnmpTarget} with Username \${username}, Authentication('\${authProtocol}') and Privacy('\${privProtocol}')
Storage Pool	storage_pool_created	Storage Pool \${resourceId} has been created
Storage Pool	storage_pool_deleted	Storage Pool \${resourceId} has been deleted
Storage Pool	storage_pool_updated	Storage Pool \${resourceId} has been updated
Syslog	syslog_server_added	Syslog server \${protocol}://\${host}:\${port} with severity \${severity} is added into the configuration
Syslog	syslog_server_updated	Syslog server \${old_protocol}://\${old_host}:\${old_port} is

Table 25 ECS audit messages (continued)

Service	Audit item	Audit message
		updated to \${protocol}://\${host}:\${port} with severity \${severity} in the configuration
Syslog	syslog_server_deleted	Syslog server \${protocol}://\${host}:\${port} is removed from the configuration
Transformation	transformation_created_message	Transformation created
Transformation	transformation_updated_message	Transformation updated
Transformation	transformation_pre_check_started_message	Transformation precheck started
Transformation	transformation_enumeration_started_message	Transformation enumeration started
Transformation	transformation_indexing_started_message	Transformation indexing started
Transformation	transformation_migration_started_message	Transformation migration started
Transformation	transformation_recovery_migration_started_message	Transformation recovery migration started
Transformation	transformation_reconciliation_started_message	Transformation reconciliation started
Transformation	transformation_sources_updated_message	Transformation sources updated
Transformation	transformation_deleted_message	Transformation deleted
Transformation	transformation_retried_message	Transformation %s retried
Transformation	transformation_canceled_message	Transformation %s canceled
Transformation	transformation_profile_mappings_updated_message	Transformation profile mappings updated
User	user_created	Object user \${resourceId} has been created
User	user_deleted	Object user \${resourceId} has been deleted
User	user_set_password	New password has been set for object user \${resourceId}
User	user_delete_password	Password has been deleted for object user \${resourceId}

Table 25 ECS audit messages (continued)

Service	Audit item	Audit message
User	user_set_metadata	New metadata has been set for object user \${resourceId}
User	user_locked	Object user \${resourceId} has been locked
User	user_unlocked	Object user \${resourceId} has been unlocked

Alert messages

List of the alert messages used by ECS.

Alert message **Severity** labels have the following meanings:

- **Critical:** Messages about conditions that require immediate attention.
- **Error:** Messages about error conditions that report either a physical failure or a software failure.
- **Warning:** Messages about less than optimal conditions.
- **Info:** Routine status messages.

Table 26 ECS Object alert messages

Alert	Severity	Sent to...	Message	Description
Bucket hard quota	Error	Portal, API, SNMP Trap, Syslog	HardQuotaLimitExceeded: bucket {bucket_name}	
Capacity exceeded threshold	WARNING	Portal, API, ESRS, SNMP Trap, Syslog	Used Capacity exceeded configured threshold, current usage is {usage}%	
Chunk not found	Error	Portal, API, ESRS, SNMP Trap, Syslog	chunkId {chunkId} not found	
DT init failure	Error	Portal, API, ESRS, SNMP Trap, Syslog	There are more than {number} DTs failed or DT stats check failed in last {number} rounds of DT status check	DT is a directory table
License expiration	INFO	Portal, API, ESRS, SNMP Trap, Syslog	Expiration event	
License registration	INFO	Portal, API, ESRS, SNMP Trap, Syslog	Registration Event	

Table 26 ECS Object alert messages (continued)

Alert	Severity	Sent to...	Message	Description
Namespace hard quota	Error	Portal, API, SNMP Trap, Syslog	HardQuotaLimitExceeded: Namespace {namespace}	
VDC in TSO	CRITICAL	Portal, API , SNMP Trap, Syslog	Site {vdc} is marked as temporarily unavailable	TSO is a temporary site outage.
Bucket soft quota	Error	Portal, API, SNMP Trap, Syslog	SoftQuotaLimitExceeded: bucket {bucket_name}	

Table 27 ECS Fabric alert messages

Alert	Severity	Sent to...	Message	Description
Disk added	Info	Portal, API, SNMP Trap, Syslog	Disk {diskSerialNumber} on node {fqdn} was added	Disk was added
Disk failure	Critical	Portal, API, SNMP Trap, Syslog , ESRS	Disk {diskSerialNumber} on node {fqdn} has failed	Health of disk changed to BAD
Disk good	Info	Portal, API, SNMP Trap, Syslog	Disk {diskSerialNumber} on node {fqdn} was revived	Disk was revived
Disk removed	Info	Portal, API, SNMP Trap, Syslog	Disk {diskSerialNumber} on node {fqdn} was removed	Disk was removed
Disk suspect	Error	Portal, API, SNMP Trap, Syslog , ESRS	Disk {diskSerialNumber} on node {fqdn} has suspected	Health of disk changed to SUSPECT
Docker container configuration failure	Critical	Portal, API, SNMP Trap, Syslog , ESRS	Container {containerName} configuration has failed on node {fqdn} with exit code {exitCode} {happenedOn}	Configure script returned non-zero exit code
Docker container paused	Warning	Portal, API, SNMP Trap, Syslog	Container {containerName} has paused on node {fqdn}	Container paused

Table 27 ECS Fabric alert messages (continued)

Alert	Severity	Sent to...	Message	Description
Docker container running	Info	Portal, API, SNMP Trap, Syslog	Container {containerName} is up on node {fqdn}	Container moved to running state
Docker container stopped	Error	Portal, API, SNMP Trap, Syslog	Container {containerName} has stopped on node {fqdn}	Container stopped
Fabric agent failure	Critical	Portal, API, SNMP Trap, Syslog , ESRS	FabricAgent has failed on node {fqdn} (FabricAgent Failure Event)	Fabric agent health changed to BAD
Fabric agent suspect	Error	Portal, API, SNMP Trap, Syslog , ESRS	FabricAgent has suspected on node {fqdn} (FabricAgent Suspect Event)	Fabric agent health changed to SUSPECT
Net interface health down	Critical	Portal, API, SNMP Trap, Syslog , ESRS	Net interface {netInterfaceName} with ip address {ipAddress} is down on node {fqdn}	Fabric's net interface is down
Net interface health up	Critical	Portal, API, SNMP Trap, Syslog , ESRS	Net interface {netInterfaceName} with ip address {ipAddress} is up on node {fqdn}	Fabric's net interface is up
Net interface permanent down	Critical	Portal, API, ESRS	Net interface {netInterfaceName} with ip address {ipAddress} is permanently down on node {fqdn}	Net interface is down for at least 10 minutes
Net interface IP address updated	Critical	Portal, API, SNMP Trap, Syslog , ESRS	Net interface's {netInterfaceName} ip address on node {fqdn} was changed from {oldIpAddress} to {newIpAddress}	Fabric's net interface IP address changed
Node failure	Critical	Portal, API, SNMP Trap, Syslog , ESRS	Node {fqdn} has failed (Service Health Failure Event)	Node is not reachable for 30 minutes

Table 27 ECS Fabric alert messages (continued)

Alert	Severity	Sent to...	Message	Description
Node suspect	Error	Portal, API, SNMP Trap, Syslog , ESRS	Node {fqdn} has suspected failure (Node Suspect Event)	Node is not reachable for 15 minutes
Node up	Info	Portal, API, SNMP Trap, Syslog	Node {fqdn} is up	Node moved to 'up' state after it was down for at least 15 minutes
Service failure	Critical	Portal, API, SNMP Trap, Syslog , ESRS	Service {serviceName} has failed on node {fqdn}	Service health (fabric/object) changed to BAD
Service suspect	Error	Portal, API, SNMP Trap, Syslog , ESRS	: Service {serviceName} has suspected on node {fqdn}	Service health (fabric/object) changed to SUSPECT
Slot permanent down	Critical	Portal, API, SNMP Trap, Syslog , ESRS	Container {containerName} is permanently down on node {fqdn}	Container stopped/paused or not started at all for at least 10 minutes

APPENDIX B

ECS Support for SNMP

- [SNMP support in ECS](#).....242
- [SNMP MIBs supported for querying in ECS](#)..... 242
- [ECS-MIB SNMP Object ID hierarchy and MIB definition](#).....242

SNMP support in ECS

The types of SNMP servers supported in ECS.

ECS provides support for Simple Network Management Protocol (SNMP) in the following ways:

- During the Installation process, EMC personnel can configure and start an `snmpd` server to support specific monitoring of ECS node-level metrics. A Network Management Station can query these kernel-level `snmpd` servers to gather information from the ECS nodes directly for Memory and CPU usage, as defined by standard MIBs. For the list of MIBs for which ECS supports SNMP queries, see [SNMP MIBs supported for querying in ECS](#) on page 242.
- The ECS Fabric lifecycle layer includes an `snmp4j` library which acts as an SNMP server to generate SNMP Traps v2 and v3 Traps and send them to as many as ten discreet SNMP Trap recipient Network Management Stations. For details of the MIBs for which ECS supports as SNMP traps, see [ECS-MIB SNMP Object ID hierarchy and MIB definition](#) on page 242. Use the Event Notification UI in the ECS Portal to configure details about the Trap recipient servers. For details, see [Configure Event Notification servers \(SNMP or Syslog\)](#) on page 144.

SNMP MIBs supported for querying in ECS

A list of standard SNMP MIBs supported for querying by the node-level `snmpd` server that can run on each ECS node.

ECS supports basic queries from Network Management Stations for the following Simple Network Management Protocol (SNMP) Management Information Bases (MIBs):

- MIB-2
- DISMAN-EVENT-MIB
- HOST-RESOURCES-MIB
- UCD-SNMP-MIB

Using an SNMP Management Station or equivalent software, you can query ECS nodes for the following basic information:

- CPU usage
- Memory usage
- Number of processes running

ECS-MIB SNMP Object ID hierarchy and MIB definition

This section describes the SNMP OID hierarchy and provides the full SNMP MIB-II definition for the enterprise MIB known as ECS-MIB.

The SNMP enterprise MIB named ECS-MIB defines the objects `trapAlarmNotification`, `notifyTimestamp`, `notifySeverity`, `notifyType`, and `notifyDescription`. The SNMP enterprise includes supported

SNMP traps that are associated with managing ECS appliance hardware. The objects contained in the ECS-MIB have the following hierarchy:

```
emc.....1.3.6.1.4.1.1139
  ecs.....1.3.6.1.4.1.1139.102
    trapAlarmNotification...1.3.6.1.4.1.1139.102.1.1
      notifyTimestamp.....1.3.6.1.4.1.1139.102.0.1.1
      notifySeverity.....1.3.6.1.4.1.1139.102.0.1.2
      notifyType.....1.3.6.1.4.1.1139.102.0.1.3
      notifyDescription...1.3.6.1.4.1.1139.102.0.1.4
```

The following Management Information Base syntax defines the SNMP enterprise MIB named ECS-MIB:

Note

ECS sends traps from the Fabric lifecycle container, using services provided by the `snmp4j` Java library.

You can download the ECS-MIB definition (as the file `ECS-MIB-v2.mib`) from the Support Site in the Downloads section under Add-Ons.

```
ECS-MIB DEFINITIONS ::= BEGIN
  IMPORTS enterprises, Counter32, OBJECT-TYPE,
  MODULE-IDENTITY, NOTIFICATION-TYPE
  FROM SNMPv2-SMI;

  ecs MODULE-IDENTITY
    LAST-UPDATED "201605161234Z"
    ORGANIZATION "EMC ECS"
    CONTACT-INFO "EMC Corporation 176 South Street Hopkinton,
  MA 01748"
    DESCRIPTION "The EMC ECS Manager MIB module"
    ::= { emc 102 }

  emc OBJECT IDENTIFIER ::= { enterprises 1139 }

  -- Top level groups

  notificationData OBJECT IDENTIFIER ::= { ecs 0 }
  notificationTrap OBJECT IDENTIFIER ::= { ecs 1 }

  -- The notificationData group
  -- The members of this group are the OIDs for VarBinds
  -- that contain notification data.

  genericNotify OBJECT IDENTIFIER ::= { notificationData 1 }

  notifyTimestamp OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION "The timestamp of the notification"
    ::= { genericNotify 1 }

  notifySeverity OBJECT-TYPE
    SYNTAX INTEGER {
      informational (1),
      warning (2),
      error (3),
      critical (4)
    }
    MAX-ACCESS read-only
    STATUS current
```

```

        DESCRIPTION "The severity level of the event is indicated
by an integer number in the range from 1 to 3"
        ::= { genericNotify 2 }

notifyType OBJECT-TYPE
    SYNTAX OCTET STRING
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION "A type of the event"
    ::= { genericNotify 3 }

notifyDescription OBJECT-TYPE
    SYNTAX OCTET STRING
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION "A complete description of the event"
    ::= { genericNotify 4 }

-- The SNMP trap
-- The definition of these objects mimics the SNMPv2 convention
for
-- sending traps. The enterprise OID gets appended with a 0
-- and then with the specific trap code.

trapAlarmNotification NOTIFICATION-TYPE
    OBJECTS {
        notifyTimestamp,
        notifySeverity,
        notifyType,
        notifyDescription,
    }
    STATUS current
    DESCRIPTION "This trap identifies a problem on the ECS. The
description can be used to describe the nature of the change"
    ::= { notificationTrap 1 }
END

```

Trap messages that are formulated in response to a Disk Failure Alert are sent to the ECS Portal Monitor > Events > Alerts page in the format Disk {diskSerialNumber} on node {fqdn} has failed:

```

2016-08-12 01:33:22 lviprbig248141.lss.emc.com [UDP:
[10.249.248.141]:39116->[10.249.238.216]]:
iso.3.6.1.6.3.18.1.3.0 = IpAddress: 10.249.238.216 iso.
3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.4.1.1139.102.1.1 iso.
3.6.1.4.1.1139.102.0.1.1 = STRING: "Fri Aug 12 13:48:03 GMT 2016"
iso.3.6.1.4.1.1139.102.0.1.2 = STRING: "Critical" iso.
3.6.1.4.1.1139.102.0.1.3 = STRING: "2002" iso.
3.6.1.4.1.1139.102.0.1.4 = STRING: "Disk 1EGAGMRB on node provo-
mustard.ecs.lab.emc.com has failed"

```

Trap messages that are formulated in response to a Disk Back Up Alert are sent to the ECS Portal Monitor > Events > Alerts page in the format Disk {diskSerialNumber} on node {fqdn} was revived:

```

2016-08-12 04:08:42 lviprbig249231.lss.emc.com [UDP:
[10.249.249.231]:52469->[10.249.238.216]]:
iso.3.6.1.6.3.18.1.3.0 = IpAddress: 10.249.238.216 iso.
3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.4.1.1139.102.1.1 iso.
3.6.1.4.1.1139.102.0.1.1 = STRING: "Fri Aug 12 16:23:23 GMT 2016"
iso.3.6.1.4.1.1139.102.0.1.2 = STRING: "Info" iso.
3.6.1.4.1.1139.102.0.1.3 = STRING: "2025" iso.

```

```
3.6.1.4.1.1139.102.0.1.4 = STRING: "Disk 1EV1H2WB on node provo-  
copper.ecs.lab.emc.com was revived"
```

