

ECS

Version 3.0 and higher

New Features and Changes

302-999-906

A11

February 2020

Copyright © 2019-2020 Dell Inc. or its subsidiaries. All rights reserved.

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 Technologies, Dell, EMC, 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.

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

CONTENTS

Chapter 1	New Features and Changes Overview	7
Chapter 2	ECS Version 3.4	9
	SSL ciphers.....	10
	Notify failed password change attempts.....	10
	Limit number of login sessions per management user.....	10
	Security improvements.....	10
	Lock user account after failed login attempts.....	12
	Secure ports 9069, 9099.....	12
	Option to disable HTTP	12
	Block NFS and HDFS ports.....	12
	CAT II STIGs for OS.....	12
	Automatic logout inactivity timer is configurable from server setting.....	12
	Configure User Agreement for UI.....	13
	Configure user interface inactivity timer separate from session timeout.....	13
	Configure management session timeout.....	13
	API for retrieving Geo replication status of an object.....	14
	Ability to remove or fail a VDC.....	14
	Replication groups changes.....	14
	Support and Qualify Gemalto Safenet and IBM SKLM as EKM vendors.....	14
	Allow customer provided LDAPs certificates.....	14
	Support Custom LDAP Certificate.....	14
	ECS Monitoring and Alerting user interface.....	15
	ECS Performance for Requests.....	15
	ECS Performance Monitoring - Request Latency	15
	ECS Performance Monitoring - Error Type.....	15
	ECS Performance - Latency Calculation in percentiles.....	15
	ECS Performance Alerts.....	15
	Metadata Overhead Improvements.....	16
	EX500 Hardware.....	17
Chapter 3	ECS Version 3.3	19
	Product Name Change to ECS.....	20
	ECS support for retention and retention expiration periods for Atmos objects....	20
	File system enabled.....	25
	S3A support.....	28
	Prefix capability in metadata search.....	28
	Capacity forecast.....	29
	ECS Report Export.....	29
	Alert policy.....	30
	Configurable policy alerts.....	30
	New alert policy.....	30
	Send a test SNMP trap.....	31
	Acknowledge all alerts.....	32
	Alert messages.....	32
	Key Management.....	43
	External key management.....	46
	External Key Manager Configuration.....	47

	Add External Key Management Servers to Cluster.....	48
	Add VDC to EKM Server Mapping.....	48
	Key Rotation.....	48
	DISA STIG rules support.....	49
	Garbage Collection changes and improvements.....	49
	Reduce object metadata overhead.....	49
	New Monitoring Stack.....	50
Chapter 4	ECS Version 3.2.2	51
	Dell EMC Generation 3 (Gen3) Hardware.....	52
	Sudo required for cs_hal led command on Gen3 hardware.....	54
	ECS Software.....	54
Chapter 5	ECS Version 3.2.1	55
	CAS monitoring of unused objects.....	56
	Large object support improvements for Data Domain Cloud tier.....	56
	Configuration parameter for NFS directory listing	56
Chapter 6	ECS Version 3.2	59
	ECS Portal changes.....	60
	Secure Remote Services improvements.....	60
	Monitoring improvements.....	60
	More active sites supported for the ECS Passive replication configuration.....	61
	Log improvements.....	61
	Additional software alerts.....	62
	HDP 2.6.2 support.....	62
	Large object support for Data Domain Cloud Tier.....	62
	ECS Software installation improvements.....	62
	Licensing.....	62
	Centera migration.....	63
	Ability to enable encryption for migration of Centera data to ECS	63
	Transform service (transformsvc) is disabled by default.....	63
	Documentation changes.....	64
Chapter 7	ECS Version 3.1	65
	Retention and Expiration on Atmos Objects through the Atmos (UMD).....	66
	Support for Geo-Passive architecture.....	66
	Support for hosted sites.....	66
	S3 bucket policies.....	66
	Read-only access to buckets during an outage.....	67
	Metadata search and D@RE.....	67
	Swift and S3 interoperability.....	67
	Network support improvements.....	67
	Ability to separate networks after ECS is installed and running.	67
	Policy-based routing for network separation	68
	Additional network for data traffic for Centera systems.....	68
	IP address change on ECS nodes.....	68
	Application registration for CAS API is disabled by default.....	69
	User tags.....	69
	Partial garbage collection (space reclamation) is enabled by default.....	69
	Secure Remote Services support for FOB-based passwords.....	69
	ECS Service Console.....	69
	Changes in the documentation set.....	70

Chapter 8	ECS Version 3.0 and 3.0 Hotfixes	71
	ECS 3.0 HF3 improvements.....	72
	ECS 3.0 HF2 improvements.....	72
	ECS 3.0 new features and changes	73
	S3 Protocol enhancements.....	73
	OpenStack Swift protocol support for Dynamic Large Objects (DLOs) and Static Large Objects (SLOs).....	73
	Support for sending SNMP traps from ECS.....	73
	Support for Remote Syslog Servers.....	74
	Space Reclamation by Partial Garbage Collection.....	74
	Platform Locking.....	75
	Dashboard and Monitoring Improvements.....	75
	CAS Advanced Retention Management.....	75
	CAS Behavior Change for Default Retention Period in Objects Written without Object-level Retention in Compliance Namespaces.....	75
	ECS Compliance certified for ECS Appliances with ECS 3.0 and ECS Software Only installations on ECS-certified third-party storage hardware.....	76
	Atmos Support Improvements.....	76
	CIFS-ECS Support.....	76
	Network separation.....	77
	ECS Software.....	77

CHAPTER 1

New Features and Changes Overview

This document lists and describes the new features that are introduced in ECS releases 3.0 and higher, as well as any product changes in the releases.

CHAPTER 2

ECS Version 3.4

This chapter lists the new features, changes, and the improvements that are introduced in version 3.4 of ECS.

• SSL ciphers	10
• Notify failed password change attempts	10
• Limit number of login sessions per management user	10
• Security improvements	10
• Lock user account after failed login attempts	12
• Secure ports 9069, 9099	12
• Option to disable HTTP	12
• Block NFS and HDFS ports	12
• CAT II STIGs for OS	12
• Automatic logout inactivity timer is configurable from server setting	12
• Configure User Agreement for UI	13
• Configure user interface inactivity timer separate from session timeout	13
• Configure management session timeout	13
• API for retrieving Geo replication status of an object	14
• Ability to remove or fail a VDC	14
• Replication groups changes	14
• Support and Qualify Gemalto Safenet and IBM SKLM as EKM vendors	14
• Allow customer provided LDAPs certificates	14
• Support Custom LDAP Certificate	14
• ECS Monitoring and Alerting user interface	15
• ECS Performance for Requests	15
• ECS Performance Monitoring - Request Latency	15
• ECS Performance Monitoring - Error Type	15
• ECS Performance - Latency Calculation in percentiles	15
• ECS Performance Alerts	15
• Metadata Overhead Improvements	16
• EX500 Hardware	17

SSL ciphers

ECS security has been enhanced, where weak ciphers are excluded.

ECS 3.4 and higher requires that ECS clients support the following SSL ciphers:

- DHE-RSA-AES256-GCM-SHA384
- DHE-RSA-AES128-GCM-SHA256

Note: The load balancer must support these ciphers before upgrading to 3.4.0.x. If the load balancer does not support these ciphers, do not proceed with an upgrade. As the load balancer cannot communicate with ECS post upgrade that may cause a DU.

Notify failed password change attempts

ECS logs an audit log event when the password change fails for reasons such as connection disruption, forgetting the password. You can view audit log for failed password changes.

Use the **Monitor > Events > Audit** tab to view and manage audit data. Refer to *ECS Version 3.4 Monitoring Guide* for details.

Limit number of login sessions per management user

The feature provides a mechanism for system admin to limit the number of management sessions per user. This operation requires the System Administrator role in ECS.

Refer to *ECS Version 3.4 Admin Guide* for details and procedures.

Security improvements

You can use the ECS Portal to change your password, set password rules, manage user sessions, and set user agreement text. In ECS 3.4, to change password, you need to provide both the current password and the new password. You must have the System Administrator role in ECS to set password rules, manage user sessions, and set user agreement text.

The Password Rules take effect only when the **Password Rules Switch** is **Enabled**. The **Password Rules Switch** is **Disabled** by default. Click **Settings > Security > Password Rules** to change the **Password Rules Switch** settings.

The password rules are applied for system admin creating and changing other user's password except history rules and common characters rule. Password complexity and expiration is not applicable to the AD/LDAP users.

The following table provides the list of password rules:

Table 1 Password rules

Field	Description
Password Rules Switch	<p>The Password Rules Switch controls the compliance checks for the rules on this page.</p> <ul style="list-style-type: none"> • Enabled: All rules apply when creating/updating passwords • Disabled: No rules apply when creating/updating passwords

Table 1 Password rules (continued)

Field	Description
Minimum number of characters	The minimum number of characters that are required for a password must be greater than or equal to 8, and not to exceed 256. The combination of minimum number of characters (upper case, lower case, numeric, special) must not exceed 256 characters.
Minimum number of upper case characters	The minimum number of upper case characters that are required for a password must be greater than or equal to 1.
Minimum number of lower case characters	The minimum number of lower case characters that are required for a password must be greater than or equal to 1.
Minimum number of numeric characters	The minimum number of numeric characters that are required for a password must be greater than or equal to 1.
Minimum number of special characters	The minimum number of special characters that are required for a password must be greater than or equal to 1.
Minimum character change	The minimum number of characters that must change from previous password must be greater than or equal to 3.
Password expiration(days)	The maximum lifetime of a password until it expires must be greater than or equal to 1. After password expiration, you can log in using the old password, however you need to change the password upon login to UI.
Max login attempts	<p>The maximum number of attempts to log in with invalid password before the user is locked.</p> <ul style="list-style-type: none"> • Must be greater than or equal to 3. • System admin can unlock users through the management user interface.
Max lock day	The number of days to lock a user when password is not changed. Must be greater than or equal to 1.
Min Password duration(hours)	The minimum lifetime for a password. Must be greater than or equal to 1. You cannot change the password more than once in 24-hour period which is configurable except when system admin is changing other user's password.
Passwords repeat	The minimum number of previous passwords that cannot be reused. Must be greater than or equal to 1.

Refer to *ECS Version 3.4 Admin Guide* for the procedure.

Lock user account after failed login attempts

User account is locked after repeated failed login attempts, to prevent unauthorized access to ECS.

Attempting to log in with invalid password for set number times locks the user. The default number of times is three, but it is configurable under **Settings > Security > Password Rules**. System admin can unlock the user through the management user interface. After system admin resets a password, you need to change the password upon next login to the user interface. Refer to *ECS Version 3.4 Admin Guide* for details.

Secure ports 9069, 9099

The 9069 and 9099 are public IP ports protected by Fabric Firewall manager. Port is not available outside of the cluster.

Refer to *ECS Version 3.4 Security Configuration Guide* for details.

Option to disable HTTP

You can now disable HTTP ports on ECS nodes.

Refer to *ECS Version 3.4 Security Configuration Guide* for details.

Block NFS and HDFS ports

With the System admin rights, you can block NFS ports of buckets that are file system enabled. You can now block HDFS ports or disable HDFS service.

Refer to *ECS Version 3.4 Security Configuration Guide* for details.

CAT II STIGs for OS

ECS provides compliance for DISA STIG security standards which addresses the CAT I and CAT II vulnerabilities.

For more information on CAT I and CAT II vulnerabilities see, *ECS Hardening Guide*.

Automatic logout inactivity timer is configurable from server setting

Automatic logout inactivity timer is configurable from the ECS Portal user interface.

The default setting is 15 minutes for inactivity session timeout and inactive user interface session timeout. Both the values are now configurable through the user interface **Settings > Security > Sessions** interface. Refer to *ECS Version 3.4 Admin Guide* for details.

Configure User Agreement for UI

Added support to the security features that allows the customization of user agreement text which can be displayed at the time of login.

You can enter the customized agreement text in the **Security > User agreement** tab. After configuring the user agreement, the login screen displays the following *By clicking Login, I agree to the terms in the Information Systems User Agreement. (User Agreement is a link to pop up the configured user agreement text).* Refer to *ECS Version 3.4 Admin Guide* for details.

Configure user interface inactivity timer separate from session timeout

As part of support for session timeouts, the ECS Portal user interface can be configured to force a timeout sooner than the session timeout.

Refer to *ECS Version 3.4 Admin Guide* for details.

Configure management session timeout

You can use the ECS Portal to manage user sessions.

This operation requires the System Administrator role in ECS.

The following table provides the description of the session timeout fields:

Table 2 Sessions

Field	Description
Management sessions per user	The maximum number of sessions per user. Must be greater than or equal to 2.
Inactive session timeout(mins)	<p>The maximum number of minutes before an inactive session is terminated.</p> <ul style="list-style-type: none"> Must be greater than or equal to 15. Inactive session timeout applies to any logged in session. Use the Inactive UI Session timeout setting to limit ECS UI session inactivity.
Inactive UI session timeout(mins)	<p>The maximum number of minutes before an inactive ECS UI session is terminated.</p> <ul style="list-style-type: none"> Must be greater than or equal to 10. Inactive UI session timeout value may not exceed the Inactive session timeout value. Inactive UI session timeout applies to any session where user logged in through ECS UI.

Refer to *ECS Version 3.4 Admin Guide* for details.

API for retrieving Geo replication status of an object

You can retrieve the Geo replication status of an object using API to confirm that the object is successfully replicated.

The ECS S3 head supports Geo replication status of an object with Get-ObjectInfo. The API retrieves the Geo replication status of an object using Get-ObjectInfo. You can automate the capacity management operations, enable site reliability operations, and avoid accidental deletion of critical data. For more information, see *ECS 3.4 Data Access Guide*.

Ability to remove or fail a VDC

You can remove a VDC from a replication group (RG) in a multi VDC federation without affecting the VDC or other RGs associated with the VDC. Removing VDC from RG no longer initiates PSO. Removing a VDC from RG initiates recovery.

The ability to fail a VDC is also added. Failing a VDC or permanent site outage (PSO) is done from the VDC level replication group (RG) initiates recovery, but failed VDC does not initiate recovery. Refer to *ECS Version 3.4 Admin Guide* for details.

Replication groups changes

You can remove a virtual data center (VDC) from a replication group for a site failover using ECS user interface.

When you remove a VDC from the replication group, it initiates recovery operations for the VDC for only the replication group. The VDC continues to work for other replication groups.

Support and Qualify Gemalto Safenet and IBM SKLM as EKM vendors

The new version of ECS Supports the Gemalto Safenet and IBM SKLM (Security Key Lifecycle Manager) key managers. The key manager supported versions are determined by Dell EMC's Key-Trust-Platform (KTP) client.

Refer *ECS Version 3.4 Security Configuration Guide* for details.

Allow customer provided LDAPs certificates

ECS allows customer provided LDAPs certificates for management user access.

Refer to *ECS Version 3.4 Admin Guide* for details.

Support Custom LDAP Certificate

You can customize the LDAP certificate according to your security standards.

Refer to *ECS Version 3.4 Admin Guide* for details.

ECS Monitoring and Alerting user interface

ECS user interface(UI) is enhanced to display a series of new metrics, NATIVE ECS API, and customer-facing content. With the UI improvements, you can visualize the new and existing datasets with the ability to establish thresholds for alarming and export the data.

In ECS 3.4, Advanced Monitoring from Grafana is integrated into ECS user interface. Refer to the Advanced Monitoring section of *ECS Version 3.4 Monitoring Guide* for details.

ECS Performance for Requests

You can manage the system performance and capacity and troubleshoot issues, as you have visibility on All, GET, PUT, POST, DELETE, and HEAD Requests in the Advanced Monitoring dashboard.

Refer to *ECS Version 3.4 Monitoring Guide* for details.

ECS Performance Monitoring - Request Latency

Advanced Monitoring dashboards provide critical information about the ECS processes on the VDC you are logged in to. The advanced monitoring dashboards are based on time series database.

The following items are added as part of advance monitoring:

- Request Latency (First Byte Read, and Last Byte Write)
- Median (p50)
- 99th percentile latency (p99) numbers

Refer to *ECS Version 3.4 Monitoring Guide* for details.

ECS Performance Monitoring - Error Type

You can view the status of your REQUESTS, troubleshoot, and visualize when requests are errored out. The feature enables the onboarding of CLOUDIQ in the future. The dashboard is enhanced to display the error type by requests, protocols, User, and top 10 contributors (IP addresses) of each type of error.

You can see the enhancements under **Advanced Monitoring > Related dashboards > Data Access Performance - by Nodes**. Refer to *ECS Version 3.4 Monitoring Guide* for details.

ECS Performance - Latency Calculation in percentiles

The advanced monitoring shows percentile latency Median and 99th percentile latency. In the user interface the **Read Latency** is changed to **First Byte Latency for Read** and **Write Latency** is changed to **Last Byte Write Latency**.

You can see Latency under the **Advanced Monitoring > Data Access Performance - Overview** tab. Refer to *ECS Version 3.4 Monitoring Guide* for details.

ECS Performance Alerts

New configurable performance alerts are added for ECS 3.4 release.

The following table lists the performance alerts introduced in 3.4:

Table 3 Alerts

Alert	Severity	Symptom code	Message	Description
First Byte Latency For Read	Warning Error Critical	4009 4010 4011	First Byte Latency for Read is \$ {inspectorValue}ms crosses threshold \$ {thresholdValue}ms	If TTFB for read latency crosses the threshold that is specified, then the alert is triggered.
Last Byte Latency For Write	Warning Error Critical	4003 4014 4015	Last Byte Latency for Write is \$ {inspectorValue}ms crosses threshold \$ {thresholdValue}ms	If TTLB for write latency crosses the threshold that is specified, then the alert is triggered.
CPU Usage Percent	Warning Error Critical	4001 4002 4003	CPU usage is \$ {inspectorValue}% crosses threshold \$ {thresholdValue}%	If CPU usage percent crosses the threshold that is specified, then the alert is triggered.
Space Usage Percent	Warning Error Critical	4005 4006 4007	Disk space usage is \$ {inspectorValue}% crosses threshold \$ {thresholdValue}%	If Disk usage percent crosses the threshold that is specified, then the alert is triggered.
Monitoring Health	Critical	4016 4017 4018	Data that are recorded in TSDB is lagging by {thresholdValue} mins on node x.x.x.x.	Data that are recorded in TSDB is lagging by {thresholdValue} mins on node x.x.x.x.
Root File System filling on node	WARNING ERROR INFO	2039 2042 2043	Root File System is {percent} % full on node.{node}	Filling percent 85 to 95. Filling percent greater 95. Filling percent less 83.
Firewall health is BAD or SUSPECT	BAD SUSPECT	2051 2052	Firewall health is BAD! {reason} Firewall health is SUSPECT! {reason}	Rules or IP sets do not exist, system firewall is off, IP tables or IP set utilities do not exist Rules or IP sets do not exist, trying to recover.

Refer to *ECS Version 3.4 Monitoring Guide* for details.

Metadata Overhead Improvements

This feature enables ECS to consume less storage for small objects.

This feature has no user interface or REST API changes. You will realize the benefits of this feature by being able to store more data on a given ECS system compared to ECS 3.3.

EX500 Hardware

EX500 is a new Gen3 EX series hardware that is introduced in ECS 3.4.

The EX500 series are available in Dell EMC or third-party rack configurations. The EX500 rack can accommodate EX300 nodes. The features of EX500 are:

- Hyper-converged nodes
 - Up to 16 nodes per rack
 - 12 to 24 drives per node
 - 8 TB and 12 TB disk options
- More flexibility: 5 to 16 nodes per rack.
- Linear scalability: Every node has same performance characteristics.
- Easy serviceability: Drives are front accessible and hot-pluggable.

CHAPTER 3

ECS Version 3.3

This chapter lists the new features, changes, and the improvements that are introduced in version 3.3 of ECS.

• Product Name Change to ECS	20
• ECS support for retention and retention expiration periods for Atmos objects	20
• File system enabled	25
• S3A support	28
• Prefix capability in metadata search	28
• Capacity forecast	29
• ECS Report Export	29
• Alert policy	30
• Key Management	43
• DISA STIG rules support	49
• Garbage Collection changes and improvements	49
• Reduce object metadata overhead	49
• New Monitoring Stack	50

Product Name Change to ECS

The product is no longer called Elastic Cloud Storage, it is renamed to ECS and this change is implemented in:

- GUI
- Product documentation including release notes
- White papers and any other external facing documents
- Any other tools or software components such as GeoDrive tool and xDoctor

ECS support for retention and retention expiration periods for Atmos objects

ECS supports setting retention periods, and retention expiration periods on Atmos objects.

Retention periods

Retention periods define how long ECS retains an object before it can be edited or deleted. During the retention period, the object cannot be edited or deleted from the system until the retention period has expired.

While creating an Atmos object in ECS, the object retention can be:

- Defined directly on the object
- Inherited from the retention period set on the ECS bucket in which the object is created

When a retention policy is set on the ECS namespace, set the retention period directly on the object. The object does not inherit the retention policy in the namespace.

The table shows the Atmos retention periods

Table 4 Atmos retention periods

Retention set on the	Using the	Notes
Object	<p>Atmos API through the</p> <ul style="list-style-type: none"> • Header retention period in seconds: 'x-emc-retention-period:60' • User meta data (UMD), end date: 'x-emc-meta:user.maui.retentionEnable=true,user.maui.retentionEnd=2016-10-21T10:00Z' • Both header, and UMD: 'x-emc-meta:user.maui.retentionEnable=true,user.maui.retentionEnd=2016-10-21T18:14:30Z' -header 'x-emc-retention-period:60' 	<ul style="list-style-type: none"> • Retention can be set on the object while creating, or updating the object settings. • Header retention period is defined in seconds. • End date defines the UMD retention. • If retention period is set from both the header and the UMD, the UMD attribute is checked first and takes precedence over the setting in the header. • You cannot modify the retention period after it has been set on the object until the period has expired. • When using the x-emc header to set retention <ul style="list-style-type: none"> ▪ If one is defined, -1 sets an infinite retention period and disable the expiration period. ▪ -2 disables the retention period set on the object.

Table 4 Atmos retention periods (continued)

Retention set on the	Using the	Notes
ECS namespace	ECS Portal from the New Namespace or Edit Namespace page.	<ul style="list-style-type: none"> If you want to set a retention period for an object, and a retention policy has been defined on the object user's namespace, you must still define a retention period directly on the object as described earlier. If a retention policy is set on the ECS namespace, and/or a retention period is set on a bucket within the namespace, and an object is created within the bucket, ECS retains the namespace, bucket, and object for the longest retention periods set for either the namespace, or bucket.
	ECS REST API <code>POST /object/namespaces/namespace/{namespace}/retention</code>	
ECS bucket	ECS Portal from the New Bucket , or Edit Bucket page.	<ul style="list-style-type: none"> If a retention period has been set on the object itself through the object header, ECS retains the object for the longest time set on the namespace, bucket, or object. If a retention end date is defined on an object through the Atmos API, ECS uses the Atmos API retention end date set on the object, and ignores the namespace retention policy, and bucket retention periods when creating the object. While applying a retention policy on a subtenant (bucket) containing Atmos objects, the retention policy is applied to both objects created in the subtenant after the retention policy was set, and objects that were created in the subtenant before the retention policy was set.
	ECS REST API <code>PUT /object/bucket/{bucketName}/retention</code>	

Note: For further details about Namespace Retention Policies and Bucket Retention Periods, see the ECS Administration Guide that is available on [ECS Product Documentation page](#).

Example: Request and response to create an object with retention set:

```
POST /rest/namespace/file1 HTTP/1.1
User-Agent: curl/7.37.0
Host: 10.247.179.228:9022
Accept: */*
x-emc-date:Thu, 16 Feb 2017 19:28:13 GMT
x-emc-meta:user.maui.retentionEnable=true,user.maui.retentionEnd=2017-06-30T06%3A38%3A44Z
x-emc-uid:f082110e13f249649340e172fb7b4956/u1
x-emc-utf8:true
Content-Type:plain/text
x-emc-signature:2Gz51WT+jQdMjlobDV0mz7obsXM=
Content-Length: 774
```

Response

```
HTTP/1.1 201 Created
Date: Thu, 16 Feb 2017 19:28:17 GMT
x-emc-policy: default
x-emc-utf8: true
x-emc-request-id: 0af7b3e4:15a4849d95e:37c:0
x-emc-delta: 774
Location: /rest/objects/
0a40bd045f7373d367639f095d1db0d15acadb82d5d2cd108e2142f4be04635c-59bdb9b6-20c0-4f55-
bc91-9db728a58854
```

```
x-emc-mtime: 1487273295379
Content-Length: 0
Server: ViPR/1.0
```

Example: Request and response to get object metadata:

```
curl --head -H "x-emc-date:Mon, 30 Jan 2017 16:56:35 GMT"
-H "x-emc-uid:7a2593be81374744adbf8e3983e7bd84/u1"
-H "x-emc-signature:CQgfoiIQ/DCif7TafcIskWyVpME="
http://10.247.179.228:9022/rest/objects/
dlbced53f2ebbc51af1d84747bd198d123d3b8585293a5bf0d32bb73c6cf4b-365f4482-c24a-4eca-
b24a-070efe29bf63

Response

HTTP/1.1 200 OK
Date: Mon, 30 Jan 2017 16:56:35 GMT
x-emc-mtime: 1485795387838
x-emc-retention-period: 21798212
x-emc-meta: user.maui.retentionEnd=2017-10-10T00:00:00Z,user.maui.retentionEnable=true,allow-
inline-update=false,atime=2017-01-30T16:45:48Z,ctime=2017-01-30T16:56:27Z,ctype=plain/
text,data-range=CAAQgFA=,dek=kq/WlRg/
7qbmaCcLF8pFvqlDJ8+suPTdVddBBZFwZA86muG3P0Pb7w==,dekAlgo=AESKeyWrapRFC5649,etag=0-,fs-mtime-
millisec=1485795387838,itime=2017-01-30T16:45:48Z,kekId=s3.7a2593be81374744adbf8e3983e7bd843cd
da755061bac6c12c06eb02800a7fee4b11ac2e03f62bb01eee02995068e56,keypoolid=s3.7a2593be81374744adb
f8e3983e7bd84,keypoolname=7a2593be81374744adbf8e3983e7bd84,keyversion=0,mtime=2017-01-30T16:56
:27Z,namespace=s3,nlink=1,object-
name=,objectid=dlbced53f2ebbc51af1d84747bd198d123d3b8585293a5bf0d32bb73c6cf4b-365f4482-
c24a-4eca-
b24a-070efe29bf63,objname=file,parentOid=53ae036bfcfb46f5580b912222f3026835e3ef972c7e3e532ba4a
5de30b1946e,parentZone=urn:storageos:VirtualDataCenterData:365f4482-c24a-4eca-
b24a-070efe29bf63,policynamespace=default,retention=CgYIoKOZmlE=,size=0,type=regular,uid=u1,parent=
apache,gid=apache
x-emc-useracl: ul=FULL_CONTROL
x-emc-groupacl: other=READ
x-emc-policy: default
x-emc-request-id: 0af7b3e4:159f0185cf7:957:4
Content-Type: plain/text
Content-Length: 0
Server: ViPR/1.0
```

Example: Update an object with retention values.

```
POST /rest/namespace/file2?metadata/user HTTP/1.1
User-Agent: curl/7.37.0
Host: 10.247.179.228:9022
Accept: */*
x-emc-date:Thu, 16 Feb 2017 19:37:15 GMT
x-emc-meta:user.maui.retentionEnable=true,user.maui.retentionEnd=2017-07-30T06%3A38%3A44Z
x-emc-uid:f082110e13f249649340e172fb7b4956/u1
x-emc-utf8:true
Content-Type:plain/text
x-emc-signature:5UPpZcCfO0vtxMTW62fa2/2SmLg=

Response

HTTP/1.1 200 OK

Date: Thu, 16 Feb 2017 19:37:16 GMT
x-emc-policy: _int
x-emc-utf8: true
x-emc-request-id: 0af7b3e4:15a4849d95e:582:0
```

```
Content-Length: 0
Server: ViPR/1.0
```

Expiration period

When a retention period end date is defined for an Atmos object, and the expiration period is also set on the object, ECS automatically deletes the object at the date that is defined in the expiration period. The expiration period:

- Can be set on objects using the Atmos API, or the `x-emc` header.
- The expiration period must be later than the retention end date.
- The expiration period is disabled by default.
- When using the `x-emc` header to set retention and expiration, a -1 value disables the expiration period.

Example: Set the expiration period using the `x-emc` header:

```
POST /rest/namespace/file2 HTTP/1.1
User-Agent: curl/7.37.0
Host: 10.247.179.228:9022
Accept: */*
x-emc-date:Tue, 31 Jan 2017 19:38:00 GMT
x-emc-expiration-period:300
x-emc-uid:a2b85977fd08488b80e646ea875e990b/u1
Content-Type:plain/text
x-emc-signature:krhYBfKSiM3mFOT6FtRB+2/xZnw=
Content-Length: 10240
Expect: 100-continue
```

Example: Request and response using the Atmos API:

```
POST /rest/namespace/file2 HTTP/1.1
User-Agent: curl/7.37.0
Host: 10.247.179.228:9022
Accept: */*
x-emc-date:Thu, 02 Feb 2017 02:47:32 GMT
x-emc-meta:user.maui.expirationEnable=true,user.maui.expirationEnd=2017-03-30T20:20:00Z
x-emc-uid:239e20dec7a54301a0b02f6090edcace/u1
Content-Type:plain/text
x-emc-signature:5tGEyK/9qUZCPSnQ9OPodktN+Zo=
Content-Length: 10240
Expect: 100-continue

Response

HTTP/1.1 100 Continue
HTTP/1.1 201 Created
Date: Thu, 02 Feb 2017 02:47:33 GMT
x-emc-policy: default
x-emc-request-id: 0af7b3e4:159fb81ddae:345e:0
x-emc-delta: 10240
Location: /rest/objects/5c3abaf60e0e207abec96baf0618c0461b7cd716898f8a12ee236aed1ec94bea-
c86ee0e9-8709-4897-898e-c3d1895e1d93
x-emc-mtime: 1486003652813
Content-Length: 0
Server ViPR/1.0 is not blacklisted
Server: ViPR/1.0
```

Example: Request and response for update meta data with Atmos API:

```
POST /rest/namespace/file?metadata/user HTTP/1.1
User-Agent: curl/7.37.0
Host: 10.247.179.228:9022
Accept: */*
x-emc-date:Thu, 02 Feb 2017 02:44:13 GMT
x-emc-meta:user.maui.expirationEnable=true,user.maui.expirationEnd=2017-03-30T20:20:00Z
x-emc-uid:239e20dec7a54301a0b02f6090edcace/u1
Content-Type:plain/text
x-emc-signature:9pzcc/Ce4Lq3k52QKdfWLY1Z1Yc=
```

Response

```
HTTP/1.1 200 OK
Date: Thu, 02 Feb 2017 02:44:14 GMT
x-emc-policy: _int
x-emc-request-id: 0af7b3e4:159fb81ddae:339e:0
Content-Length: 0
Server ViPR/1.0 is not blacklisted
Server: ViPR/1.0
```

Retention start delay window

Atmos enables you to specify a start delay window when creating a retention period, which enables you to migrate to ECS. Also, this feature prevents the objects from getting into retention after initial upload of an object.

Atmos creates subtenant request header, `x-emc-retention-start-delay` that captures the autocommit interval.

```
./atmoscurl.pl -user USER1 -action PUT -pmode TID -path / -header "x-emc-retention-period:300" -header "x-emc-retention-start-delay:120" -include
```

Retention start delay applied on object mtime

In Atmos object creation, if retention start delay is set on the bucket (`x-emc-retention-start-delay`), the start delay for the object is calculated based on `time-since-mtime` of the object.

Note: The `time-since-mtime` is considered to calculate the start delay as it does not give an exact time to complete an upload and `x-emc-retention-start-delay` could be shorter even as a few minutes.

Override bucket-level retention for migrated objects

- If the user decides to migrate data through Atmos API to an ECS bucket in a compliant namespace with maui retention headers and if there are any conflicting retentions, the longest retention wins.
- On noncompliant buckets, for Atmos migrated objects, the `user.maui*headers` specifies the final retention value on an object. If there are no `user.maui*headers` available, the longest retention wins.
- On object creation in ECS through Atmos API, the `user.maui*headers` cannot be combined with any of `x-emc-retention` headers.

Atmos API supports GeoDrive

Atmos API supports GeoDrive on ECS. GeoDrive is a windows application that enables Atmos data to be mirrored to the local Windows file system, and it is the same as CIFS-ECS.

File system enabled

S3 buckets can be File System (FS) enabled so that the files that are written using the S3 protocol can be read using the file protocols, such as Network File system (NFS) and Hadoop Distributed File System (HDFS), and the opposite way.

Enabling FS access

You can enable FS access using the `x-emc-file-system-access-enabled` header when creating a bucket using the S3 protocol. File system access can also be enabled when creating a bucket from the ECS Portal (or using the ECS Management REST API).

Limitation on FS support

When a bucket is FS enabled S3 life cycle management cannot be enabled.

Cross-head support for FS

Cross-head support is accessing objects written using one protocol using a different, ECS-supported protocol. Objects written using the S3 head can be read and written using NFS and HDFS file system protocols.

An important aspect of cross-head support is how object and file permissions translate between protocols and for file system access how user and group concepts translate between object and file protocols.

You can find more information about the cross-head support with file systems in the ECS Administration Guide which is available from the [ECS Product Documentation page](#).

NFS WORM (Write Once, Read Many)

NFS data become Write Once Read Many (WORM) compliant when autocommit is implemented on it.

In detail, creating files through NFS is a multi step process. To write to a new file, NFS client first sends the CREATE request with no payload to NFS server. After receiving a response, the server issues a WRITE request. It is a problem for FS enabled buckets under retention as the file created with 0 bytes blocks any writes to it. Due to this reason, until ECS 3.3, retention on FS enabled bucket makes the whole mounted file-system read-only. There is no End of File (EOF) concept in NFS. Setting a retention for files, on the FS enabled buckets, after writing to them does not work as expected.

To remove the constraints that are placed on NFS files in a retention enabled bucket, the autocommit period is implemented on NFS data. For this reason, it is decided to introduce the autocommit period during which certain types of updates (for now identified as writes, Acl updates and deletes that are required for rsync, and rename that is required for Vim editor) are allowed, which removes the retention constraints for that period alone.

Note:

- The autocommit and the Atmos retention start delay are the same. See [Retention start delay window](#) on page 24.
- Autocommit period is a bucket property like retention period.
- Autocommit period is:
 - Applicable only for the file system enabled buckets with retention period
 - Applicable to the buckets in noncompliant namespace
 - Applies to only requests from NFS and Atmos

Seal file

The seal file functionality helps to commit the file to WORM state when the file is written ignoring the remaining autocommit period. The seal function is performed through the command: `chmod ugo-w <file>` on the file.

Note: The seal functionality does not have any effect outside the retention period.


High level overview

This table describes the Autocommit terms

Table 5 Autocommit terms

Term	Description
Autocommit period	Time interval relative to the object's last modified time during which certain retention constraints (example: file modifications, file deletions, and so on) are not applied. It does not have any effect outside of the retention period.
Retention Start Delay	Atmos head uses the start delay to indicate the autocommit period.

The following diagram provides an overview of the autocommit period behavior.



Autocommit configuration

The autocommit period can be set from the user interface or bucket REST API or S3 head or Atmos subtenant API.

User Interface

The user interface has the following support during bucket create and edit:

- When the File System is not enabled, no autocommit option is displayed.
- When the File System is enabled /no retention value that is specified, autocommit is displayed but disabled.
- When the File System is enabled/retention value selected/autocommit is displayed and enabled for selection.

Note: Maximum autocommit period is limited to the smaller of the Bucket Retention period or the default maximum period of one day.

REST API

Create bucket REST API is modified with the new header, `x-emc-autocommit-period`.

```
lgrou063:~ # curl -i -k -T /tmp/bucket -X POST https://10.247.99.11:4443/object/bucket -H "$token" -H "Content-Type: application/xml" -v
```

The contents of /tmp/bucket

```
<object_bucket_create>
  <name>bucket2</name>
  <namespace>s3</namespace>
  <filesystem_enabled>true</filesystem_enabled>
  <autocommit_period>300</autocommit_period>
  <retention>1500</retention>
</object_bucket_create>
```

S3 head

Bucket creation

Bucket creation flow through s3 head can make use of optional request header, `x-emc-autocommit-period:seconds` to set the autocommit period. The following checks are made in this flow:

- Allow only positive integers
- Settable only for file system buckets
- Settable only when the retention value is present

```
./s3curl.pl --ord --id=naveen --key=+1Zh4YC2r2puuUaj3Lbnj3u0G9qgPRj0RIWJhPxH
--createbucket -- -H 'x-emc-autocommit-period:600' -H 'x-emc-file-system-
access-enabled:true' -H 'x-emc-namespace:ns1' http://10.249.245.187:9020/
bucket5 -v
```

Atmos

Atmos creates a subtenant request header, `x-emc-retention-start-delay`, captures the autocommit interval.

```
./atmoscurl.pl -user USER1 -action PUT -pmode TID -path / -header "x-emc-
retention-period:300" -header "x-emc-retention-start-delay:120" -include
```

Behavior of file operations

This table describes the behavior of file operations

Table 6 Behavior of file operations

File Operation	Expected within autocommit period	Expected within retention period (after autocommit period)
Change permission of file	Allowed	Denied
Change ownership of file	Allowed	Denied
Write to existing file	Allowed	Denied
Create empty file	Allowed	Allowed
Create non-empty file	Allowed	Denied
Remove file	Allowed	Denied
Move file	Allowed	Denied
Rename file	Allowed	Denied
Make dir	Allowed	Allowed
Remove directory	Denied	Denied

Table 6 Behavior of file operations (continued)

File Operation	Expected within autocommit period	Expected within retention period (after autocommit period)
Move directory	Denied	Denied
Rename directory	Denied	Denied
Change permission on directory	Denied	Denied
list	Allowed	Allowed
Read file	Allowed	Allowed
Truncate file	Allowed	Denied
Copy of local read-only files to NFS share	Allowed	Allowed
Copy of read-only files from NFS share to NFS share	Allowed	Allowed
Change atime/mtime of file/ directory	Allowed	Denied

S3A support

The AWS S3A client is a connector for AWS S3, which enables you to run MapReduce jobs with ECS S3.

Note:

- ECS does not enable you to run S3A client on FS enabled buckets.
- S3A support is available on Hadoop 2.7 or later version.

Prefix capability in metadata search

S3 API metadata search supports the prefix and delimiter parameters. It follows the standard S3 definition of these parameters. Prefix capability transforms every metadata query into a multi query request with AND operation between prefix and the query string. In other words, it is possible to combine the AND and OR predicates in the queries.

S3 API metadata is modified to support prefix and delimiter parameters as described here:

```
GET /bucketName/?prefix={prefix}&delimiter={delimiter}&query={queryString}
```

Limitations

- A prefix is always applied before the query.

- Custom sorting is not supported with prefixes. If sorting is specified together with a prefix, the API returns 400 Bad Request.
- Objects are returned in lexicographical order.
- Using `ObjectName` in a query string together with a prefix is not allowed. It creates ambiguity as both filter objects based on name. If both are specified, the API returns 400 Bad Request.

Capacity forecast

You can use the **Capacity** tab to monitor when the capacity is expected to reach 50% and 80%. Capacity forecast is based on the current usage pattern that is shown on 1 day, 7 days, and 30-days usage trend. Capacity Forecast data is shown either for the entire VDC, for an individual storage pool or for nodes.

Note: The capacity ETA shown as N/A could be due to the following reasons:

1. There is not enough historical data for forecast. At least two data points (1 hour apart) are required. It could happen when the ECS system is deployed. Click the **History** button at VDC, storage pool, or node levels to verify.
2. If capacity passed intended target, the ETA is set to 0.
3. The used capacity shows a down trend for the specified time (for example, 7 days). Click the **History** button or get the history through dashboard API to verify.

To see the capacity forecast data from the ECS Portal, select **Monitor > Capacity Utilization > Capacity**. **Capacity** tab is the default.

To see the data about total capacity, used capacity, and available capacity, click **History**.


Capacity Forecast is calculated based on the total capacity and used capacity.

ECS Report Export

Export icon enables you to export all capacity data from all the tables and graphs to pdf, doc, excel, and .csv formats for later consumption. To select the format, and export the data, use the export icon in the upper right of the menu bar on each table and graph.

The exported data can be used to get a longer term view on capacity usage and consumption trends.

Figure 1 Export icons



Alert policy

Alert policies are created to alert about metrics, and are triggered when the specified conditions are met. Alert policies are created per VDC.

You can use the **Settings > Alerts Policy** page to view alert policies.

There are two types of alert policy:

System alert policies

- System alert policies are precreated and exist in ECS during deployment.
- All the metrics have an associated system alert policy.
- System alert policies cannot be updated or deleted.
- System alert policies can be enabled/disabled.
- Alert is sent to the UI and all channels (SNMP, SYSLOG, and Secure Remote Services).

User-defined alert policies

- You can create User-defined alert policies for the required metrics.
- Alert is sent to the UI and customer channels (SNMP and SYSLOG).

Configurable policy alerts

The ECS 3.3 release enables you to construct or update some of the Alert Policies. There is a new Landing Page with table of existing policies and a row action to Edit or Delete an alert policy. Alerts that are configurable by this page:

- PO
- Performance
- Capacity
- GC
- RPO
- P1
- All other existing alerts

New alert policy

You can use the **Settings > Alerts Policy > New Alert Policy** tab to create user-defined alert policies.

Procedure

1. Select **New Alert Policy**.
2. Give a unique policy name.
3. Use the metric type drop-down menu to select a metric type.

Metric Type is a grouping of statistics. It consists of:

- Btree Statistics
- CAS GC Statistics

- Geo Replication Statistics
 - Metering Statistics
 - Garbage Collection Statistics
 - EKM
4. Use the metric name drop-down menu to select a metric name.
 5. Select level.
 - a. To inspect metrics at the node level, select Node.
 - b. To inspect metrics at the VDC level, select VDC.
 6. Select polling interval.

Polling Interval determines how frequently data should be checked. Each polling interval gives one data point which is compared against the specified condition and when the condition is met, alert is triggered.
 7. Select instances.

Instances describe how many data points to check and how many should match the specified conditions to trigger an alert.

For metrics where historical data is not available only the latest data is used.
 8. Select conditions.

You can set the threshold values and alert type with Conditions.

The alerts can be either a Warning Alert, Error Alert, or Critical Alert.
 9. To add more conditions with multiple thresholds and with different alert levels, select **Add Condition**.
 10. Click **Save**.

Send a test SNMP trap

You can send a test SNMP trap to validate configuration, and test traps without having to do a real failure.

Before you begin

This operation requires the System Administrator role in ECS.

Procedure

1. In the ECS Portal, select **Settings > Event Notification**.

On the **Event Notification** page, the **SNMP** tab displays by default and lists the SNMP servers that have been added to ECS.
2. To send a test SNMP trap, click the drop-down under Actions, and select **Send Test Trap**.

A success message is displayed for a successful SNMP trap.
3. Confirm that the test trap has reached the SNMP Target Host.

Current test functionality is to ensure that ECS can send out a trap, but there is no verification that it was reached.

Acknowledge all alerts

Alerts can be acknowledged individually or by bulk using the Acknowledge All Alerts button. You can choose to acknowledge all the alerts or acknowledge a subset of the alerts using filters.

About this task

You can use the **Monitor > Events > Alerts** tab to acknowledge alerts.

Procedure

1. To acknowledge all alerts, click the **Acknowledge All Alerts** button.
 - a. To acknowledge a subset of all alerts, use the table filter to filter by a combination of date and time, severity, type, or namespace, and then click **Acknowledge All Alerts**.

The bulk alert acknowledgment process runs in the background and may take a few minutes to complete. Only one bulk alert acknowledgment can be processed at a time.

2. On the confirmation pop-up screen, to initiate acknowledgment, click **OK** or to exit without acknowledgment click **Cancel**.

Clicking the **Acknowledge All Alerts** initiates a background task to acknowledge all the matching alerts. The response either shows successfully initiated or fails.

To keep a record of the acknowledge all alerts request, a new informational alert of type **Bulk Alert Ack** will be generated after the acknowledgment completes. Clear the filter and manually refresh the table.

Alert messages

List of the alert messages that ECS uses.

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 7 ECS Object alert messages

Alert	Severity	Symptom code	Sent to...	Message	Description	Action
Btree chunk level GC	Warning	1321	Portal, API, Secure Remote Services, SNMP Trap, Syslog	System metadata garbage reclamation throughput is too slow to catch up with garbage detection.	Event trigger source <ul style="list-style-type: none"> • Example: Reclaimed Btree Garbage is less than 10% of the remaining BTree garbage as BTree GC is slow at Chunk reclamation. • This condition has persisted for last 7 	Contact ECS Remote Support

Table 7 ECS Object alert messages (continued)

Alert	Severity	Symptom code	Sent to...	Message	Description	Action
					<p>days, leading to creation of this alert.</p> <ul style="list-style-type: none"> Derived it from formula: Full_Garbage > 1TB, and Garbage_Detected_Rate - Garbage_Chunk_Reclaim_Rate > 100GB 	
Btree disk level GC	Warning	1325	Portal, API, Secure Remote Services, SNMP Trap, Syslog	Capacity free-up throughput is too slow to catch up with system metadata garbage reclamation.	<p>Event trigger source</p> <ul style="list-style-type: none"> Example: Reclaimed Btree Garbage is less than 10% of the Full garbage, as BTree GC is slow at disk level reclamation. This condition has persisted for last 7 days, leading to creation of this alert. Derived from formula: if Garbage_Pending_Delete > 1TB, and Garbage_Chunk_Reclaim_Rate - Garbage_Capacity_Reclaim_Rate > 100GB 	Contact ECS Remote Support.
Btree partial GC	Warning	1329	Portal, API, Secure Remote Services, SNMP Trap, Syslog	Partial GC for system metadata is too slow.	<p>Event trigger source</p> <ul style="list-style-type: none"> Example: Rate of Btree Partial GC conversion to full Garbage is less than 10% of the Partial GC eligible for Conversion. Btree partial GC works too slow to convert partial 	Contact ECS Remote Support.

Table 7 ECS Object alert messages (continued)


Alert	Severity	Symptom code	Sent to...	Message	Description	Action
					<p>garbage into full garbage.</p> <ul style="list-style-type: none"> This condition has persisted for last 7 days, leading to creation of this alert. Derived from formula : If Partial_Eligible_Garbage > 1TB, and Partial_To_Full_Convert_Rate < 100GB 	
Bucket hard quota	Error	1006	Portal, API, SNMP Trap, Syslog	HardQuotaLimitExceeded: bucket {bucket_name}		
Bucket soft quota	Warning	1008	Portal, API, SNMP Trap, Syslog	SoftQuotaLimitExceeded: bucket {bucket_name}		
Capacity alerting	Warning Error Critical	1111 1112 1113	Portal, API, SNMP Trap, Syslog	Storage pool {Storage pool} has {id}% remaining capacity meeting threshold of {id}%.	The severity of the alert depends on how close the remaining storage pool capacity is to reaching the configured threshold. Capacity alerting is not set by default: set capacity alerts to receive them. You can set them by editing an existing storage pool or when you create a storage pool.	
Capacity exceeded threshold	Warning	1100	Portal, API, Secure Remote Services, SNMP Trap, Syslog	Used Capacity of the VDC exceeded configured threshold, current usage is {usage}%.	<p>The configured threshold is set at 80% of the Used Capacity of the VDC by default.</p> <p> CAUTION If the used capacity reaches 90%, you cannot write or modify object data.</p>	Contact ECS Remote Support representative to determine the appropriate solution.

Table 7 ECS Object alert messages (continued)

Alert	Severity	Symptom code	Sent to...	Message	Description	Action
Capacity license threshold	Error	997	Portal, API, Secure Remote Services, Trap, Syslog	Licensed Capacity Entitlement Exceeded Event	The capacity of the system is greater than was licensed.	
Chunk not found	Error	1004	Portal, API, Secure Remote Services, SNMP Trap, Syslog	chunkId {chunkId} not found		
CPU Usage Percent	Warning Error Critical	4001 4002 4003	Portal, API, SNMP Trap, Syslog	CPU usage is \$ {inspectorValue} % crosses threshold \$ {thresholdValue} %	If CPU usage percent crosses the threshold specified then the alert is triggered.	
Disabled CAS GC	Info Warning Error Critical	1316 1317 1318 1319	Portal, API, Secure Remote Services, SNMP, Trap, Syslog	CAS Processing is paused.	<ul style="list-style-type: none"> CAS GC is Content Addressable Storage Garbage Collection. CAS GC is disabled. 	Contact ECS Remote Support to ensure that it should stay enabled.
DT init failure	Error	3001	Portal, API, Secure Remote Services, SNMP Trap, Syslog	There are more than {numbers} DTs failed or DT stats check failed in last {number} rounds of DT status check.	<ul style="list-style-type: none"> DT is a directory table. The default value is set at 8 DTs for this alert to trigger. 	
EKM Server Certificate Expiry	Warning Error	1361 1362	Portal, API, Secure Remote Services, SNMP Trap, Syslog	<ul style="list-style-type: none"> The server certificate for EKM server expires in 30 days. Renew the certificate. The server certificate for EKM server expires in 7 days. Renew the certificate. 		

Table 7 ECS Object alert messages (continued)

Alert	Severity	Symptom code	Sent to...	Message	Description	Action
EKM Server Connection Status	Warning Error	1369 1370	Portal, API, Secure Remote Services, SNMP Trap, Syslog	The EKM server is not responding. Ensure that the server is connected.		
First Byte Latency For Read	Warning Error	4009 4010 4011	Portal, API, SNMP Trap, Syslog	First Byte Latency for Read is \$ {inspectorValue}ms crosses threshold \$ {thresholdValue}ms	If TTFB for read latency crosses the threshold specified then the alert is triggered.	
Last Byte Latency For Write	Warning Error Critical	4003 4014 4015	Portal, API, SNMP Trap, Syslog	Last Byte Latency for Write is \$ {inspectorValue}ms crosses threshold \$ {thresholdValue}ms	If TTLB for write latency crosses the threshold specified then the alert is triggered.	
License expiration	Info	998	Portal, API, Secure Remote Services, SNMP Trap, Syslog	Expiration event		
License registration	Info	100	Portal, API, Secure Remote Services, SNMP Trap, Syslog	Registration Event		
Memory outside Btree writes cache	Warning	1349	Portal, API, Secure Remote Services, SNMP Trap, Syslog	For cm process memory of X bytes is allocated outside Btree write cache on node <Node IP>.		
Metering read latency	Warning Error Critical	1205 1206 1207	Portal, API, Secure Remote Services, SNMP Trap, Syslog	Read latency is 300 millisecond, crosses threshold 250 millisecond. Read latency is 505 millisecond,		Contact ECS Remote Support.

Table 7 ECS Object alert messages (continued)

Alert	Severity	Symptom code	Sent to...	Message	Description	Action
				crosses threshold 500 millisecond. Read latency is 1050 millisecond, crosses threshold 1000 millisecond.		
Metering write latency	Warning Error Critical	1205 1206 1207	Portal, API, Secure Remote Services, SNMP Trap, Syslog	Write latency is 300 millisecond, crosses threshold 250 millisecond. Write latency is 555 millisecond, crosses threshold 500 millisecond. Write latency is 1500 millisecond, crosses threshold 1000 millisecond.		Contact ECS Remote Support.
Monitoring Health	Critical	4016 4017 4018	Portal, API, Secure Remote Services, SNMP Trap, Syslog	Data recorded in TSDB is lagging by {thresholdValue} mins on node x.x.x.x		
Namespace hard quota	Error	1005	Portal, API, SNMP Trap, Syslog	HardQuotaLimitExceeded: Namespace {namespace}		
Namespace soft quota	Warning	1009	Portal, API, SNMP Trap, Syslog	SoftQuotaLimitExceeded: Namespace {namespace}		
Notification	Any		Any	User-defined message.	Custom message that is defined and provided by the user.	
Process memory table free space percent	Error	1354	Portal, API, Secure Remote Services, SNMP Trap, Syslog	Memory table size for blob process is X % less than the specified threshold of Y % on <node IP>.		Contact ECS Remote Support.
Repo chunk level GC	Warning	1333	Portal, API, Secure	User garbage collection	Event trigger source	Contact ECS Remote Support.

Table 7 ECS Object alert messages (continued)

Alert	Severity	Symptom code	Sent to...	Message	Description	Action
			Remote Services, SNMP Trap, Syslog	throughput is too slow to catch up with garbage detection.	<ul style="list-style-type: none"> Example: Repo Chunk reclamation rate is less than 10% of the remaining garbage. This condition has persisted for last 7 days, leading to creation of this alert. Derived from formula: $\text{Full_Garbage} > 10\text{TB}$, and $\text{Garbage_Detected_Rate} - \text{Garbage_Chunk_Reclaim_Rate} > 100\text{GB}$ 	
Repo disk level GC	Warning	1337	Portal, API, Secure Remote Services, SNMP Trap, Syslog	Capacity free-up throughput is too slow to catch up with user garbage collection.	<p>Event trigger source</p> <ul style="list-style-type: none"> Example: Repo disk level GC reclamation rate is less than 10 % of Garbage pending delete at disk level. This condition has persisted for last 7 days, leading to creation of this alert. Derived from formula: If $\text{Garbage_Pending_Delete} > 10\text{TB}$, and $\text{Garbage_Chunk_Reclaim_Rate} - \text{Garbage_Capacity_Reclaim_Rate} > 100\text{GB}$ 	Contact ECS Remote Support.
Repo partial GC	Warning	1341	Portal, API, Secure Remote	Partial GC for user garbage is too slow.	<p>Event trigger source</p> <ul style="list-style-type: none"> Example: Repo Partial repo GC 	Contact ECS Remote Support.

Table 7 ECS Object alert messages (continued)

Alert	Severity	Symptom code	Sent to...	Message	Description	Action
			Services, SNMP Trap, Syslog		<p>works too slow to convert partial garbage into full garbage.</p> <ul style="list-style-type: none"> This condition has persisted for last 7 days, leading to creation of this alert. Derived from formula: If Partial_Eligible_Garbage > 10TB, and Partial_To_Full_Convert_Rate < 100GB 	
RPO	Warning	1012	Portal, API, Secure Remote Services, Trap, Syslog	RPO for replication group {RG} is {HH} hour {SS} seconds greater than {HH} hour threshold set.	The recovery point objective (RPO) is greater than the RPO threshold. The default value is one hour.	
Slow CAS GC Object Cleanup	Info Warning Error Critical	1312 1313 1314 1315	Portal, API, Secure Remote Services, SNMP, Trap, Syslog	CAS Processing object cleanup speed is slow.	CAS GC cleanup tasks are lagging.	
Slow CAS GC Reference Collection	Info Warning Error Critical	1308 1309 1310 1311	Portal, API, Secure Remote Services, SNMP, Trap, Syslog	CAS Processing reference collection speed is slow.	CAS GC reference collection tasks are lagging.	
Slow Journal Parsing	Info Warning Error Critical	1304 1305 1306 1307	Portal, API, Secure Remote Services, SNMP, Trap, Syslog	Journal parsing speed is slow.	Journal parsing speed is slow.	
Space Usage Percent	Warning Error	4005 4006	Portal, API, SNMP, Trap, Syslog	Disk space usage is \$ {inspectorValue}	If Disk usage percent crosses the threshold	

Table 7 ECS Object alert messages (continued)

Alert	Severity	Symptom code	Sent to...	Message	Description	Action
	Critical	4007		% crosses threshold \$ {thresholdValue} %	specified then the alert is triggered.	
GC Status	Warning	1345	Portal, API, Secure Remote Services, SNMP Trap, Syslog	Space reclamation for user data/system metadata is disabled. Make sure it is disabled for temporary purpose, and re-enable it when ready.		Contact ECS Remote Support.
VDC in TSO	Critical	1007	Portal, API , SNMP Trap, Syslog	Site {vdc} is marked as temporarily unavailable.	TSO is a temporary site outage.	

Table 8 ECS fabric alert messages

Alert	Severity	Symptom code	Sent to...	Message	Description	Action
Disk added	Info	2019	Portal, API, SNMP Trap, Syslog	Disk {diskSerialNumbe r} on node {fqdn} was added.	Disk was added.	
Disk failure	Critical	2002	Portal, API, SNMP Trap, Syslog, Secure Remote Services	Disk {diskSerialNumbe r} on node {fqdn} has failed.	Health of disk that is changed to BAD.	
Disk good	Info	2025	Portal, API, SNMP Trap, Syslog	Disk {diskSerialNumbe r} on node {fqdn} was revived.	Disk was revived.	
Disk mounted	Info	2035	Portal, API, SNMP Trap, Syslog	Disk {diskSerialNumbe r} on node {fqdn} has mounted.	Disk was mounted.	
Disk removed	Info	2020	Portal, API, SNMP Trap, Syslog	Disk {diskSerialNumbe r} on node {fqdn} was removed.	Disk was removed.	

Table 8 ECS fabric alert messages (continued)

Alert	Severity	Symptom code	Sent to...	Message	Description	Action
				r} on node {fqdn} was removed.		
Disk suspect	Error	2003	Portal, API, SNMP Trap, Syslog, Secure Remote Services	Disk {diskSerialNumbe r} on node {fqdn} has suspected.	Health of disk that is changed to SUSPECT.	
Disk unmounted	Warning	2036	Portal, API, SNMP Trap, Syslog	Disk {diskSerialNumbe r} on node {fqdn} has unmounted.	Disk was unmounted.	
Docker container configuration failure	Critical	2022	Portal, API, SNMP Trap, Syslog, Secure Remote Services	Container {containerName} configuration has failed on node {fqdn} with exit code {exitCode} {happenedOn}.	Configure script returned nonzero exit code. The configure script is provided by object and called by fabric on object container start-up. It is only applicable for the object container.	
Docker container paused	Warning	2017	Portal, API, SNMP Trap, Syslog	Container {containerName} has paused on node {fqdn}.	Container paused	
Docker container running	Info	2016	Portal, API, SNMP Trap, Syslog	Container {containerName} is up on node {fqdn}.	Container moved to running state.	
Docker container stopped	Error	2015	Portal, API, SNMP Trap, Syslog	Container {containerName} has stopped on node {fqdn}.	Container stopped	
Events cannot be delivered.	Error	2038	Portal, API, Secure Remote Services, SNMP Trap, Syslog	Events cannot be delivered through {SMTP ESRS} and lost.	Verify configuration of the channel for which the alert is.	
Firewall health is BAD or SUSPECT	Bad Suspect	2051 2052	Portal, API, Secure Remote Services,	Firewall health is BAD! {reason}	Rules or ip sets do not exist, system firewall is off, ip tables or ip set utils do not exist.	

Table 8 ECS fabric alert messages (continued)

Alert	Severity	Symptom code	Sent to...	Message	Description	Action
			SNMP Trap, Syslog	Firewall health is SUSPECT! {reason}	Rules or ip sets do not exist, trying to recover.	
Fabric agent failure	Critical	2013	Portal, API, SNMP Trap, Syslog	FabricAgent has failed on node {fqdn}.	Fabric agent health is bad.	
Fabric agent suspect	Error	2014	Portal, API, SNMP Trap, Syslog	FabricAgent has suspected on node {fqdn}.	Fabric agent health is suspect.	
Net interface health down	Critical	2023	Portal, API, SNMP Trap, Syslog, Secure Remote Services	Net interface {\$netInterfaceName}[on node \$FQDN] is down[with IP address \$IP]".	Fabric's net interface is down.	
Net interface health up	Info	2024	Portal, API, SNMP Trap, Syslog, Secure Remote Services	Net interface {\$netInterfaceName}[on node \$FQDN] is up[with IP address \$IP]".	Fabric's net interface is up.	
Net interface permanent down	Critical	2026	Portal, API, Secure Remote Services	Net interface {\$netInterfaceName}[on node \$FQDN] is permanently down[with IP address \$IP].	Net interface is down for at least 10 minutes.	
Net interface IP address updated	Info	2027	Portal, API, SNMP Trap, Syslog	Net interface's {netInterfaceName} IP address on node {fqdn} was changed to {newIpAddress}.	Fabric's net interface IP address changed	
Node failure	Critical	2006	Portal, API, SNMP Trap, Syslog, Secure Remote Services	Node {fqdn} has failed.	Node is not reachable for 30 minutes.	
Node suspect	Error	2007	Portal, API, SNMP Trap, Syslog, Secure	Node {fqdn} has suspected.	Node is not reachable for 15 minutes.	

Table 8 ECS fabric alert messages (continued)

Alert	Severity	Symptom code	Sent to...	Message	Description	Action
			Remote Services			
Node up	Info	2018	Portal, API, SNMP Trap, Syslog	Node {fqdn} is up.	Node moved to 'up' state after it was down for at least 15 minutes.	
Root file system filling on node	Warning Critical	2039 2042	Portal, API, SNMP Trap, Syslog, Secure Remote Services	Thresholds exceeded, usable space on root fs <BYTES> are less than threshold for <LEVEL> level on node <NODE>	<ul style="list-style-type: none"> Threshold between 15G and 10G triggers warning. Threshold Less than 10G of free space results in Critical alert. 	
Slot permanent down	Critical	2021	Portal, API, SNMP Trap, Syslog, Secure Remote Services	Container {containerName} is permanently down on node {fqdn}.	Container stopped/paused or not started at all for at least 10 minutes	
Service failure	Critical	2011	Portal, API, Syslog, Secure Remote Services	Service Health Failure Event	Service failed	
Service suspect	Error	2012	Portal, API, Syslog, Secure Remote Services	Service Health Suspect event	Service health is suspect.	

Table 9 Secure Remote Services alert messages

Alert	Severity	Symptom code	Sent to...	Description
TestDialHome	N/A	TestDialHome	Secure Remote Services	Tests that Secure Remote Services connections can be established and that the call home functionality works.

Key Management

To support Data at Rest Encryption (D@RE), ECS maintains a hierarchy of encryption keys where a parent key in the hierarchy is used to protect a child key. Before ECS 3.3, these keys were natively managed by ECS across the geo-federated environment. From ECS 3.3, support for

certain External Key Management solutions that are Key Management Interoperability Protocol (KMIP) compliant have been added. Further to support industry standard practices, ECS 3.3 now supports user initiated key rotation to limit amount of data that is protected by any given key. Note that key rotation is available for both native and external key management.

Key Management uses centralized key servers to store top-level Key Encrypting Keys (KEKs). Key management servers provide HSM-based key protection and latest encryption technology.

Key Management protects against the loss of an entire appliance. The top-level key information is stored outside the appliance. If an appliance is compromised, data remains inaccessible.


 **Note:** Key Management in Settings is visible only when an encryption license is installed.

There are two types of keys:

- User-supplied keys with the S3 API headers
- Randomly generated, hierarchically structured system keys
- External key management

With the S3 API, encryption keys can be specified in the header. If the encryption header has the key that is provided, and then encryption of the object is done with the user-supplied key. ECS validates that the key provided for update, appends, and reads is the same as the key used for object creation. If the encryption header does not provide a key and the object still requires encryption, and then ECS uses a system-generated key to encrypt the data.

System-generated keys at each level are autogenerated and encrypted using the key of its immediate parent. The master key is encrypted by using asymmetric public-private encryption in which both the public and private keys are used to encrypt and decrypt the master key.

 **Note:** D@RE does not restrict caching of unencrypted encryption keys in memory, but ensures that plain text keys are never persisted on disk.

Keys are generated and encrypted in a hierarchical order:

- Public-Private Key pair – a pair of public-private keys generated for each ECS system. The private key is stored in the system root disk of the node and the public key is stored with the master key on the ECS commodity disks.
- Master Key – randomly generated key encrypted using the public key.
- Namespace Key – randomly generated namespace key encrypted using the master key.
- Bucket Key – randomly generated bucket key encrypted using the namespace key.
- Object Key – randomly generated object key encrypted using the bucket key and object id.
- Rotation Key - modified per rotation request. It is used to alter the object key encryption hierarchy to protect new data. When external key management is in use, the Rotation Key is maintained (and protected) by the external key manager. When native key management is in use, the rotation key has a protection hierarchy (consider each "->" as representing the statement "is protected by"): `Rotation Key -> Rotation Master Key -> Public-Private Key`.
- Virtual Key - represents the combination of a rotation key and a parent key. The two types of virtual keys are:
 - Virtual Master Key = Rotation Key + Master Key
 - Virtual Bucket Key = Rotation Key + Bucket Key

The hierarchy looks similar to 3.2 and earlier versions of ECS, with some modifications to use these new virtual keys.

When native key management is being used (consider each "->" as representing the statement "is protected by"):

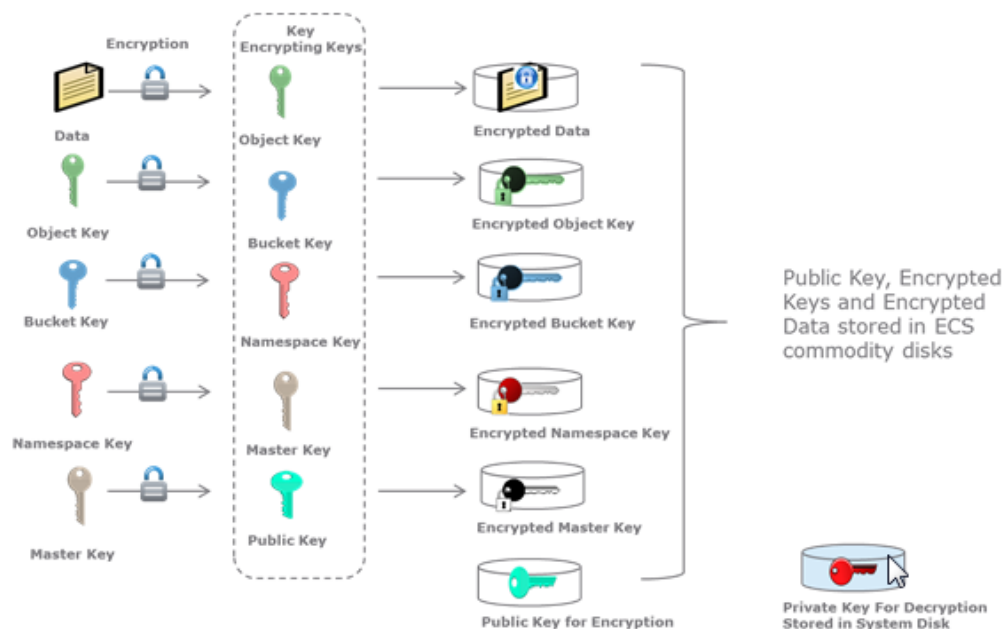
- Object key -> Virtual Bucket Key
- Bucket key -> Namespace Key -> Virtual Master Key -> Public-Private Key pair
- Rotation Key -> Rotation Master Key -> Public-Private Key

When external key management is being used (consider each "->" as representing the statement "is protected by"):

- Object key -> Virtual Bucket Key
- Bucket key -> Namespace Key -> Virtual Master Key
- Rotation Key maintained and protected by external key manager
- Master Key maintained and protected by external key manager

The private key for decryption of master key is stored on system disks (managed by vNest) and the other encrypted keys are stored in logical tables and in chunks, similar to data. They are also triple-mirrored, like data. When an object read or write request comes in, the node servicing the request traverses the key hierarchy to encrypt or decrypt the object. It uses the private key that is common to all nodes to decrypt the master key. Figure 24 provides a pictorial view of the key hierarchy.

Figure 2 Data encryption using system-generated keys



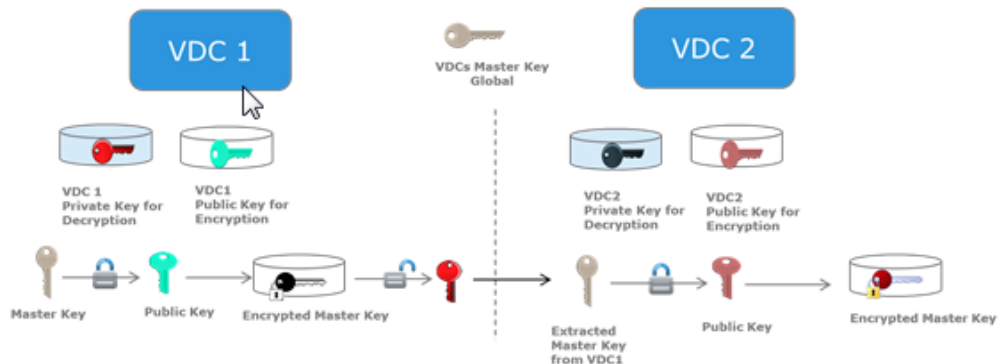
In a geo-replicated environment:

- A system being added to form or extend a federation generates public/private keys to use for encryption/decryption of the federation's master key.
- Upon federation, ECS extracts the master key on the original system using the private key of that system.
- ECS encrypts the master key using the public key that is generated by the new system and shares it with the new system.
- The master key is now global and known to both systems within the federation.

In a geo-replicated environment, when a new ECS system joins an existing system (referred to as a federation), the master key is extracted using the public-private key of the existing system and encrypted using the new public-private key pair that is generated from the new system that joined the federation. From this point on, the master key is global and known to both systems within the federation. In Figure 25, the ECS system that is labeled VDC 2 joins the federation and the master

key of VDC 1 (the existing system) is extracted and passed to VDC 2 for encryption with the public-private key randomly generated by VDC 2.

Figure 3 Encryption of the master key in a geo-replicated environment



External key management

As a part of Data at Rest Encryption (D@RE), ECS supports centralized external key managers. The centralized external key managers are compliant with the Key Management Interoperability Protocol (KMIP) which enhance the enterprise grade security in the system. Also, it enables the customers to use the centralized key servers to store top-level Key Encrypting Keys (KEKs) to provide the following benefits:

- Helps in obtaining benefits from the Hardware Security Module (HSM) based key production and the latest encryption technology that is provided by the specialized key management servers.
- Provides production against loss of the entire appliance by storing top-level key information outside of the appliance.

ECS incorporates the KMIP standard for integration with external key managers and serves as a KMIP client, and supports the following:

- Supports the Gemalto Safenet and IBM SKLM (Security Key Lifecycle Manager) key managers.
- **Note:** The key manager supported versions are determined by Dell EMC's Key-Trust-Platform (KTP) client.
- Supports the use of top-level KEK (master key) supplied by an external key manager.
- Supports rotation of top-level KEK (master key) supplied by an external key manager.

Activating external key management

Before you begin

- Ensure that you are the security administrator or have the credentials to log in as an administrator.
- Ensure that you complete the following steps before the activation.

Procedure

1. Create an EKM Cluster representing the key management cluster. For more information, see [New cluster](#)
2. Create EKM Servers, each representing a member of the external key management cluster. For more information, see [External key servers](#)

3. Map a set of the EKMServers to each VDC. For more information, see [VDC EKM Mapping](#)
 - By default, ECS requires that there are at least two EKM Servers that are mapped per VDC.
 - When mapping, the first EKMServer in the list is considered the primary, which is the server that is expected to handle key creations and retrieval.
 - The other EKMServers are considered secondaries and are used as a backup for key retrieval in case the primary is unreachable or unavailable.

Once the EKMServers have been mapped, a background process is run to validate connectivity from each VDC. If all mapped primary EKMServers are reachable, activate the EKMCluster, either through the UI or the API.

Results

Invoking the EKMCluster activation operation, triggers a background task to run the activation steps. These steps can be observed through the UI. Upon completion of the activation steps, you complete the following:

- Master key created on external key manager
- Master key retrieval validated against all external key members
- Internal reference to the new Master Key is updated
- Rotation key created on external key manager
- Rotation key retrieval validated against all external key members.
- Internal reference to the new Rotation key is updated
- All namespace keys are re-protected using the virtual master key, which would now reference the new rotation key)

External Key Manager Configuration

System Administrators can add a cluster, view VDC EKM-mapping information, and rotate keys on the **Settings > Key Management** page in the ECS Portal.

Create cluster

From the Key Management **External Key Servers** you can create a cluster and then create external key servers.

About this task

Table 10 Create cluster

Field	Description
Cluster	Name of the cluster
Type	Vendor Type
Server Count	Total number of servers that have been created for the cluster
Status	Indicates the status of the cluster. When first created, it is in the 'UNACTIVATED' status. When activation is performed, the status changes to match the step in the activation process.
FQDN/IP	FQDN or IP address of the EKM Server


Table 10 Create cluster (continued)

Field	Description
Server Host	Server host is provided in the certificate that is used to identify the client associated with the identity store.
Port	Port number that is associated with the KMIP server. The port number is used for communicating between ECS and the external key server. Default is 5696.
Actions	<ul style="list-style-type: none"> • Edit - Edit the cluster name • Delete - Delete inactive cluster • Add Server - Add External Key Server to the cluster • Activate - Activate the cluster

Procedure

1. Select **Settings > Key Management > External Key Servers > New Cluster**.
2. In the **Cluster Name** field, type a unique name for the cluster.
3. In the **External Key Management Type** field, select the vendor type from the drop-down menu.
4. Click **Save**.

After creating a cluster and before VDC EKM-Mapping, add key servers.

 **Note:** Only one cluster can be created and used per ECS federation.

Add External Key Management Servers to Cluster

An external key management cluster identifies a set of external key servers that are configured as part of the cluster. External key servers are the entities that ECS nodes contact to create/retrieve cryptographic keys. After a cluster has been created, servers must be added, and then mapped to a VDC before activating the cluster.

Add VDC to EKM Server Mapping

VDC EKM mapping assigns a subset of a cluster member server to a VDC so that nodes in the VDC can use them to access cryptographic keys.

Key Rotation

This section provides information about ECS Key rotation and the limitations.

ECS supports rotation of keys, a practice of changing keys to limit the amount of data that is protected by any given key to support industry standard practices. It can be performed on demand both through API and user interface, and is designed to minimize the risk from compromised keys.

During key rotation, the system does the following:

- Create rotation key natively or on EKM (if activated).
- Activate new rotation key across all sites in the federation.
- Once activated, the new rotation key is used to generate new virtual bucket and master keys.
- The new virtual master key is used to rewrap all namespace keys.

- The new virtual bucket key is used to protect all new object keys and associated new data.
- Rewrapped namespace keys are instrumental in protecting existing data.
- Data is not reencrypted as a result of key rotation.

To initiate key rotation, select **Settings > Key Management > Key Rotation > Rotate Keys**.

Note: Rotation is an asynchronous operation, and the latest status of current operation can be seen in the table. The Rotate Keys table also lists the status of previous rotation operations.

Limitations

- Key rotation does not rotate master, namespace, and bucket keys.
- Only one key rotation request can be active anytime and any other new request fails.
- Scope of the key rotation is at cluster level so all the new system encrypted objects are affected.
- Master, namespace, or bucket level rotation is not supported.

DISA STIG rules support

There are few organizations producing security hardening standards for IT systems. The most proficient of these organizations is the Defense Information Systems Agency (DISA). They have produced Security Technical Implementation Guides (STIG) which are generic hardening procedures for today's most widely used applications, operating systems and technologies. They have also produced Security Readiness Review (SRR) scripts to test systems against known vulnerabilities. While only select Federal agencies and the Department of Defense rigidly adhere to the DISA standards, other agencies are basing their security practices around the DISA standards due to their reputation and rigid security practices.

Dell EMC has fully embraced the DISA STIGs and has incorporated them into the standard product development and life cycle policies. ECS has adopted the SLES 12 STIG Guide by SuSE as the basis for enforcing the STIG rules on the system.

For more information, see the Dell EMC ECS 3.3 Hardening guide available at <http://www.support.dellmc.com>.

Garbage Collection changes and improvements

CAS GC Disable Dry-Run mode by default

In the previous versions of ECS, only the CAS garbage identification was enabled by default and not the garbage removal. In ECS 3.3, both CAS garbage identification and removal are enabled by default.

Improvements in Garbage Collection

With improvements such as some full reclaimable chunks skip GC verification and auto scale of partial GC task count, the rate of capacity reclamation is improved in 3.3 for both fresh install environments and upgraded clusters.

Reduce object metadata overhead

In ECS 3.3, the minimum metadata overhead is reduced from 30 KB to 7 KB per object.

New Monitoring Stack

New Monitoring stack is introduced in ECS 3.3. It is used to report ECS metrics through InfluxDB interface.

In this version of ECS, this feature is intended to be used by Dell EMC personnel only, for advanced troubleshooting. The stack uses the following containers:

- `object-teleggraf`, runs on each cluster node.
- `object-fluxd` , runs on last three nodes of each VDC.
- `object-influxdb`, runs on three nodes of the cluster.
- `object-throttler`, runs on three nodes of the cluster.

CHAPTER 4

ECS Version 3.2.2

This chapter lists the new features, changes, and the improvements that are introduced in version 3.2.2 of ECS.

- [Dell EMC Generation 3 \(Gen3\) Hardware](#).....52
- [Sudo required for cs_hal led command on Gen3 hardware](#)..... 54
- [ECS Software](#)..... 54

Dell EMC Generation 3 (Gen3) Hardware

The EX3000 and EX3000 appliances running 3.2.2 software include the following hardware components.

Table 11 EX300 and EX3000 appliance hardware components

Component	EX300 appliance	EX3000 appliance
40U rack	<p>Dell EMC Titan D racks from the factory that include:</p> <ul style="list-style-type: none"> Gen3 0U PDUs supporting single phase, three-phase delta, and three-phase WYE. Front and rear doors Racking by Dell EMC manufacturing <p>Also to the Dell EMC-racked EX300 appliance, the EX300 nodes can be installed in customer-provided racks. For more information about third-party racking requirements, see the <i>ECS EX300 Third-Party Rack Installation Guide</i>.</p>	<p>Customer-provided 40U rack, which is also available in a Dell EMC racked version must meet the following minimum requirements:</p> <ul style="list-style-type: none"> Accommodate the 1200 mm+ depth of the EX3000 4U chassis 43-mm front protrusion to mounting ears Cable management arms protrude ~4" past rear in 1200-mm cabinet Contain Gen3 2U PDUs supporting single phase, three-phase delta, and three-phase WYE. <p>For more information about third-party racking requirements, see the <i>ECS EX3000 Third-Party Rack Installation Guide</i>.</p>
Back-end (BE) switches for private network connection	<p>Two Dell EMC S5148F 25 GbE 1U Ethernet switches with 48 x 25 GbE SFP ports and 6 x 100 GbE uplink ports</p> <p>Runs network operating system 10.</p> <p>2 x 100 GbE LAG cables per HA pair</p>	<p>Two Dell EMC S5148F 25 GbE 1U Ethernet switches with 48 x 25 GbE SFP ports and 6 x 100 GbE uplink ports</p> <p>Runs network operating system 10.</p> <p>2 x 100 GbE LAG cables per HA pair</p>
Front-end (FE) switches for customer public network connection	<p>Two optional Dell EMC S5148F 25 GbE 1U Ethernet switches can be obtained for network connection or the customer can provide their own 10 GbE or 25 GbE HA pair for the front end.</p> <p>If the customer provides their own front-end switches, they must supply all LAG cables, SFPs, or external connection cables.</p> <p>If Dell EMC S5148F 25 GbE front-end switches are used, 10 GbE ports connect to the EX300 nodes, the switch runs operating system 10, and 2 x 100 GbE LAG cables are provided.</p>	<p>Two optional Dell EMC S5148F 25 GbE 1U Ethernet switches can be obtained for network connection or the customer can provide their own 25 GbE HA pair for the front end.</p> <p>If the customer provides their own front-end switches, they must supply all LAG cables, SFPs, or external connection cables.</p> <p>If Dell EMC S5148F 25 GbE front-end switches are used, 25 GbE ports connect to the EX300 nodes, the switch runs operating system 10, and 2 x 100 GbE LAG cables are provided.</p>
Nodes	<p>Minimum number of nodes per rack is 5 with increments of 1 node up to a maximum of 16 nodes.</p> <p>Hard drive sizes can be 1 TB, 2 TB, 4 TB, or 8 TB. (All drive sizes are the same in the node.)</p> <p>There is no mixing of disk sizes within a rack; for any single system, all nodes must be of the exact same drive size.</p>	<p>Up to eight server chassis in a rack.</p> <p>Chassis are in one- and two-node configurations. (Each server chassis contains either one or two nodes.) One-node chassis configuration is referred to as EX3000S and dual-node chassis configuration is referred to as EX3000D.</p>

Table 11 EX300 and EX3000 appliance hardware components (continued)

Component	EX300 appliance	EX3000 appliance
	<p>Twelve SATA hard disk drives (HDDs) in each node.</p> <p>480 GB M.2 (BOSS) system disk in each node.</p> <p>64 GB RAM per node</p> <p>Single 8-core SkyLake CPU per node. Xeon Bronze 3106 8 core/8 thread-11MB L3, 1.7 GHz, 85 W</p> <p>4x 16 GB RDIMM, 2667MT/s, Dual Rank, x4 Data Width</p> <p>Dual 750-W Platinum power supply (hot swappable)</p> <p>Each node has 4 x 10 GbE networking</p>	<p>One and two-node chassis cannot be mixed in a rack. Within a single rack, chassis must be of the same one- or two-node configuration (that is, a rack must contain all EX3000S nodes or all EX3000D nodes).</p> <p>Chassis have the following disk configurations (all hard drive are 12 TB):</p> <ul style="list-style-type: none"> EX3000S node with 45, 60, and 90 disks EX3000D node with 30 and 45 disks <p>EX3000S node has a single-server sled with the mid-plane routing to 90 drive slots. Filler is in the second server sled.</p> <p>EX3000D node has dual server sleds with the mid-plane routing each to 45 drive slots.</p> <p>Single 480 GB SSD sysdisk per node (hot swappable)</p> <p>64 GB RAM per node</p> <p>Dual 8-core Broadwell CPU per node. E5-2620v4 8-core/16-thread 2.1-GHz 20M cache 85 W</p> <p>4x 16 GB RDIMM, 2400MT/s, Dual Rank, x8 Data Width</p> <p>Dual 1600 W Gold PS per node (hot swappable)</p> <p>LSI 9361-8i SAS Controller</p> <p>Each node has 4 x 25 GbE networking</p>

EX300 upgrade paths

There are single node upgrade kits, but upgrades require a minimum order of five nodes for performance optimization.

You can add nodes in one-node increments if they are of the same configuration. Upgrades of a different capacity require a minimum order of five nodes for performance optimization.

Node capacity upgrade requests do not have to match the existing configuration. For example, if you have an EX300 appliance with five nodes containing 1 TB drives, you can add five nodes of any drive size (1 TB, 2 TB, 3 TB, or 4 TB). All drives within a node must be of the same drive size, but there can be nodes of differing drive sizes within a rack.

There are no drive upgrades.

EX3000 upgrade paths

You can add EX3000S nodes in one-node increments to an existing system. The nodes must match the existing drive configuration. For example, if your system contains EX3000S nodes with 90 disk drives you can only add EX3000S nodes that are configured with 90 disk drives.

You can add EX3000D nodes in two-node increments to add to an existing system. The nodes must match the existing drive configuration. For example, if your system contains EX3000D nodes

with 30 disk drives you can only add EX3000D nodes that are configured with 30 disk drives (60 drives per chassis).

You can add 15 drives at a time. The following 15-drive upgrades are supported:

- Convert EX3000S-45 to EX3000S-60
- Convert EX3000S-60 to EX3000S-90
- Convert EX3000D-30 to EX3000D-45

Sudo required for cs_hal led command on Gen3 hardware

For Gen3 hardware, `sudo` is required for the `cs_hal led` command.

If you run the `cs_hal led` command without `sudo`, you receive an error message similar to the following:

```
admin@dallas-artichoke:~> cs_hal led ZC14HM56 blink
cs_hal: setting LED state of disk ZC14HM56 to 'BLINK'
cs_hal: requested operation requires root privileges; retry with sudo!
```

On Gen3 hardware, you must run the `sudo cs_hal led` command to successfully complete the operation, as shown here:

```
admin@dallas-artichoke:~> sudo cs_hal led ZC14HM56 on
cs_hal: setting LED state of disk ZC14HM56 to 'ON'
admin@dallas-artichoke:~>
admin@dallas-artichoke:~> sudo cs_hal led ZC14HM56 blink
cs_hal: setting LED state of disk ZC14HM56 to 'BLINK'
admin@dallas-artichoke:~>
admin@dallas-artichoke:~> sudo cs_hal led ZC14HM56 off
cs_hal: setting LED state of disk ZC14HM56 to 'OFF'
```

ECS Software

ECS Software is a software-only solution for users seeking to deploy on an ECS appliance.

In releases earlier than 3.2.2 ECS software, could be installed on certified, custom hardware installations, and ECS appliances. In ECS 3.2.2:

- **Certified** — There is no longer support for certified hardware.
- **Custom** — To install ECS Software on custom hardware make an RPQ request for technical qualification (see: <https://inside.dell.com/docs/DOC-305519>).

CHAPTER 5

ECS Version 3.2.1

This chapter lists and describes the features and changes that were introduced in version 3.2.1 of ECS.

- [CAS monitoring of unused objects](#)..... 56
- [Large object support improvements for Data Domain Cloud tier](#)..... 56
- [Configuration parameter for NFS directory listing](#)56

CAS monitoring of unused objects

The ECS 3.2.1 release includes the Content Addressable Storage (CAS) processing feature which monitors unused CAS objects (CAS garbage data). This feature is enabled by default and is targeted for deployments migrating data from Centera storage systems into ECS. In the ECS Portal, the **Monitor > Capacity Utilization > CAS Processing** tab monitors the CAS processing data collection metrics for CAS data in buckets within a selected namespace over a specified time range. CAS processing data collection metrics include number and size of unreferenced blobs and expired reflections. Unreferenced blobs and expired reflections are unused CAS objects.

New 3.2.1 installations

The monitoring of CAS garbage data runs by default, but to enable the removal of CAS garbage data from your ECS system, you must open a Service Request with support.

3.2.1 upgrades

For ECS systems with existing CAS data that upgrade to 3.2.1, there is a CAS data bootstrap process that is automatically triggered post upgrade. After the bootstrap process is completed, the removal of CAS garbage data will have to be enabled by support. Open a Service Request to upgrade to 3.2.1 and to have the CAS garbage data removal feature enabled in your environment.

The bootstrap process builds necessary references over the existing CAS data and can require a significant amount of time depending on the amount of existing CAS data. During the bootstrap process, the unreferenced blob and reflection values do not change on the **CAS Processing** page. For example, you see zero for the unreferenced blob data that is detected and unreferenced blobs detected values. The values will not change until after the bootstrap process is complete and the CAS garbage data removal feature is enabled by support.

Large object support improvements for Data Domain Cloud tier

The ECS 3.2 release supported S3 API extensions to enable Data Domain 6.1.2 to store large objects (4 MiB or larger) on ECS. The ECS 3.2.1 release includes the following large object performance and storage efficiency improvements for data that is tiered with Data Domain 6.1.2:

- There is a new `x-emc-index-granularity` header in S3 PUT commands.
- Clients can set indexing at very fine (64 KB) granular offsets.

The finer index offset enables for faster byte range copy access which improves Data Domain tiering read performance.

Configuration parameter for NFS directory listing

NFS large directory listing operations may be slow and could result in a `BAD_COOKIE` error preventing the listing operation from completing. If NFS listing failures of large directories are observed, the configuration parameter to sort the NFS listing order can be disabled by ECS technical support to increase listing performance.

This operation can only be performed by an ECS technical support professional with `emcservice` credentials.

Use the following command on each VDC from which the mounts of NFS exports are made.

```
hostname:/opt/storageos/tools # ./cf_client --user <emcservice> --password
<emcservice_password> --set --name com.emc.ecs.blobsvc.listing.sortedorder --value false --
reason NFS Performance
```

Where:

- **<emc_service_username>** is the ECS technical support professional logged in as the **emcservice user**.
- **<emcservice_password>** is the password that is provided for the **emcservice user**.

There is no output if the command was run successfully.

To validate it, the sort order has been disabled use the following command:

```
hostname:/opt/storageos/tools # ./cf_client --user emcservice --password ChangeMe --list --
name com.emc.ecs.blobsvc.listing.sortedorder
{
  "config":
  [
    {
      "name": "com.emc.ecs.blobsvc.listing.sortedorder",
      "description": "Enable listing of LS entries in sorted order for fs-buckets.",
      "configured_value": "false",
      "default_value": "true",
      "audit": "NFS Performance",
      "modified": "1528734866625"
    }
  ]
}
```


CHAPTER 6

ECS Version 3.2

This chapter lists the features and changes that were introduced in version 3.2 of ECS.

• ECS Portal changes	60
• Secure Remote Services improvements	60
• Monitoring improvements	60
• More active sites supported for the ECS Passive replication configuration	61
• Log improvements	61
• Additional software alerts	62
• HDP 2.6.2 support	62
• Large object support for Data Domain Cloud Tier	62
• ECS Software installation improvements	62
• Licensing	62
• Centera migration	63
• Documentation changes	64

ECS Portal changes

Changes in the ECS Portal for the 3.2 release include the following:

- The **Monitoring > Erasure Coding** page has been moved under the **Monitoring > Capacity Utilization** page. There is now a new **Erasure Coding** tab and a new **Garbage Collection** tab on the **Monitoring > Capacity Utilization** page. The **Garbage Collection** tab shows the amount of garbage data that is detected, reclaimed, and pending reclamation in a local storage pool. Garbage data is blocks of data that are no longer referenced or used.
- There is a new **Node Rebalancing** tab on the **Monitoring > System Health** page that shows the erasure coding data rebalance process when a new node is added to the ECS system.
- In the **New Storage Pool** and **Edit Storage Pool** pages, there are new **Available Capacity Alerting** fields that enable you to set configurable available capacity thresholds that trigger storage pool capacity alerts.
- Storage capacities in the ECS Portal are now reported in units of GiB, TiB, and PiB. In previous releases, storage capacities were reported in units of GB, TB, and PB.
- The new **Update All VDC Endpoints** page can be accessed from the **Manage > Virtual Data Center** page. The **Update All VDC Endpoints** page can be used to update all the endpoints in a GEO configuration after the networks have been separated, or after the IP addresses of the nodes have been changed.

Secure Remote Services improvements

Improvements to Secure Remote Services (ESRS) in the 3.2 release include the following:

- Automation of Secure Remote Services initial setup. Previous releases required manual steps to set up and configure Secure Remote Services. In the 3.2 release, there are no manual steps; you can set up and configure Secure Remote Services using the **Settings > ESRS** page in the ECS Portal.
- On the **ESRS** page in the ECS Portal, you can now test the dial home feature and disable call home alerts.
You can temporarily disable call home alerts during planned maintenance activities or during troubleshooting scenarios that require taking nodes offline to prevent flooding Secure Remote Services with unnecessary alerts.

Monitoring improvements

Monitoring improvements for the 3.2 release include the following:

- Garbage collection metrics are visible in the ECS Portal on the **Monitor > Capacity Utilization > Garbage Collection** tab and can be retrieved using the ECS Management REST API. ECS reports whether garbage collection is enabled for user data and system metadata, total garbage that is detected, the capacity reclaimed (which can be further broken down into user data that is reclaimed and system metadata reclaimed), the capacity pending reclamation, and the unreclaimable garbage by virtual data center and by storage pool.
- The reserved capacity metric is visible in the ECS Portal on the **Monitor > Capacity Utilization > Capacity** tab and can be retrieved using the ECS Management REST API. Reserved capacity is the 10 percent of the total capacity that is reserved for failure handling and performing erasure encoding/XOR operations. It is not available to write new content.

- Node rebalancing metrics are visible in the ECS Portal on the **Monitor > System Health > Node Rebalancing** tab and can be retrieved using the ECS Management REST API. When nodes are added to an ECS cluster, this tab shows the progress of the erasure encoding rebalance process in the background. Statistics include whether node rebalancing is enabled, the data rebalanced, pending rebalancing, and the rate of rebalance (per day).

More active sites supported for the ECS Passive replication configuration

In ECS 3.1, the Passive configuration included exactly three sites. The Passive configuration consisted of two active sites with a third passive site, the replication target (backup site).

ECS 3.2 supports more than two active sites in a Passive configuration. The Passive configuration can now include two, three, or four active sites with an extra passive site that is the replication target. The minimum number of sites for a Passive configuration is three (two active, one passive) and the maximum number of sites is five (four active, one passive).

Log improvements

The `dataheadsvc-access.log` file records the aspects of the object heads (S3, Swift, and Atmos) supported by the object service, the file service supported by HDFS, and the CAS service. The logfile is located in `/opt/emc/caspian/fabric/agent/services/object/main/log`. Before ECS 3.2, the logfile was named `datahead-access.log`. In ECS 3.2, the logfile name changed from `datahead-access.log` to `dataheadsvc-access.log`. If you upgrade from 3.1.0.x to 3.2, it means that after the upgrade you can see both files, but the access requests will be logged to the `dataheadsvc-access.log` file only.

In 3.2, the ECS data head access log content in the `dataheadsvc-access.log` file is enhanced to include more information so that you can parse the log data by user name, namespace, bucket name, and object name. ECS data head access logs now include fields such as `LOCAL_IP`, `REMOTE_IP`, `HTTP_METHOD`, `NAMESPACE`, `BUCKET`, `USER_NAME`, `OBJECT_NAME`, `STATUS_CODE`, `CONTENT_COUNT`, and `TOTAL_TIME` (total transaction duration). The logs also include the new `STORAGE_PROCESSING_TIME` field, which is the time from when the request is received by ECS until the first byte of the response body is sent to the user. The log format is shown in the following examples.

S3 create log entry:

```
TIMESTAMP: 2018-02-19T22:57:37,809, REQUEST_ID: 0af5897d:15f31d2664d:145:19, LOCAL_IP:
10.245.137.119:9020, REMOTE_IP: 10.200.210.125:54655, USER_NAME: siyuan, HTTP_METHOD: PUT,
NAMESPACE : ns1, BUCKET:bucket1, OBJECT_NAME: logging%20life%20cycle%20test%20result.xlsx,
QUERY_STRING: -, PROTOCOL: HTTP/1.1, STATUS_CODE: 200, TOTAL_TIME: 7184, CONTENT_READ:
4619304, CONTENT_COUNT: -, STORAGE_PROCESSING_TIME: 7165
```

S3 read log entry:

```
TIMESTAMP: 2017-02-19T23:32:05,262, REQUEST_ID: 0af5897d:15f31d2664d:264:2, LOCAL_IP:
10.245.137.119:9020, REMOTE_IP: 10.200.210.125:57682, USER_NAME: siyuan, HTTP_METHOD: GET,
NAMESPACE: ns1, BUCKET: bucket1, OBJECT_NAME: logging%20life%20cycle%20test%20result.xlsx,
QUERY_STRING: acl=, PROTOCOL: HTTP/1.1, STATUS_CODE: 200, TOTAL_TIME: 29, CONTENT_READ: -,
CONTENT_COUNT: 446, STORAGE_PROCESSING_TIME: 25
```

Audit logs have been added to ECS when the following actions occur:

- The **Show Secret Key** checkbox is selected in the ECS Portal for an object user (when an S3/Atmos object user is created or edited).
- The ECS Management REST API is used to create, edit, or retrieve the secret key/password of an object user.

You can view the audit messages:

- In the ECS Portal on the **Monitor > Events > Audit** tab.
- Via the ECS Management REST API using the `GET /vdc/events` call.
- Via Syslog servers, if they have been added to ECS.

Additional software alerts

The following new alerts have been added in this release.

- Overall RPO for a replication group alert
This RPO alert is generated when the RPO exceeds its threshold which is set at 60 minutes by default. This alert can only be configured by ECS technical support. RPO refers to the point in time in the past to which you can recover.
- Capacity alerts
A System Administrator can now set configurable available capacity alerts when creating and editing storage pools. The System Administrator can set capacity thresholds that trigger storage pool capacity alerts. Capacity alerts can be configured with severities of Critical, Error, and Warning. When a capacity alert is generated, and Secure Remote Services is configured a call home alert is also generated that alerts ECS technical support that the ECS system is reaching its capacity limit.



Note: The capacity alerts are not set by default. If you are upgrading to ECS 3.2, you must configure the capacity alerts on your storage pools as described in the *ECS Administration Guide* which is available on the [ECS Product Documentation page](#).

HDP 2.6.2 support

ECS 3.2 supports the Hortonworks Data Platform (HDP) 2.6.2 distribution. ECS 3.1 supported the HDP 2.5 distribution.

Large object support for Data Domain Cloud Tier

ECS supports S3 API extensions to enable Data Domain 6.1.2 to store large objects (4 MiB or larger) on ECS, thereby providing improved storage efficiency for data tiered from Data Domain.

ECS Software installation improvements

The `install_all` command enables you install the ECS software as one package rather than having to install the Hardware Abstraction Layer (HAL) or the Fabric services separately. Refer to the *ECS Installation Guide for Appliance Only*, which is available in SolVe for complete instructions.

Licensing

In the ECS Portal on the **Settings > Licensing** page, additional information is shown, such as the maximum storage licensed for each feature. The page displays the following information:

- **Feature:** ViPR Unstructured (base feature) and may include ECS Server-Side Encryption (free software add-on feature)
- **Type:** Permanent or Temporary
- **Status:** Licensed or Expired
- **Entitlement:** Describes the maximum storage that is licensed for the ViPR Unstructured feature in TB
- **VDC Serial Number:** The Software ID of the VDC
- **PSNT:** The quantity of PSNTs (racks) in the VDC. Each rack in a VDC has a product serial number tag (PSNT). In a VDC with multiple racks, multiple PSNTs map to the serial number of the VDC. Click the right-facing arrow > next to the Feature name in the licensing table to expand and display the PSNT values.
- **Activated Site:** The license site number for the physical site where ECS is installed
- **Expiration:** If the license is temporary, the license expiration date displays. If the license is permanent, the date the license was issued displays

The ECS license file is a single file that contains base and add-on software features. A license is capacity-based, and applies to a single VDC. It is equal to the total raw capacity of the VDC and is measured in Terabytes. In geo-federated systems, each VDC requires a license. In a VDC configuration with multiple racks, the license file includes the total capacity for all racks in the VDC. There is a single ECS license file for new ECS 3.2 installations.

This new licensing scheme is only applicable for new ECS customers. There is no impact on existing customers. If you are upgrading from ECS 3.1.0.x to 3.2, you do not need to take any action.

Centera migration

The following changes have been made to Centera migration processes.

- [Ability to enable encryption for migration.](#)
- [Transform service \(transformsvc\) disabled by default.](#)

Ability to enable encryption for migration of Centera data to ECS

Before ECS 3.2, when migrating Centera data to ECS, it was not possible to securely store Data at Rest (data that is saved on disks) by encrypting the contents. In 3.2, you can enable encryption (if licensed) at either the namespace or bucket level when initiating the migration. The ECS transformation engine now implements encryption, enabling Data at Rest (D@RE) to be securely stored. It enables you to safeguard against the exposure of sensitive data in the scenario where disks are stolen from the data center. The ECS transformation engine performs two functions:

- Transformation - ECS serves data that is stored on Centera to applications, and transforms the data to native ECS objects.
- Migration - Data is physically copied to ECS and reconciled for accuracy.

Transform service (transformsvc) is disabled by default

In 3.2, transformsvc is disabled by default.

After upgrading from:

- 3.1, to 3.2.0.x, transformsvc is disabled after upgrade.

- 3.2 to 3.2.0.x, transformsv is in the state it was before upgrade. For example, if transformsv was enabled in the earlier release, before upgrade, transformsvc will remain enabled after upgrade.

Documentation changes

The following documents have been added to the documents available for Dell EMC internal personnel, and replace the previously released *ECS 3.2 Networks and Node IP Change Guide* document:

- *ECS 3.2 Networks Guide*
- *ECS 3.2 Changing the ECS Node IP Addresses Guide*

CHAPTER 7

ECS Version 3.1

This chapter lists the features and changes that were introduced in version 3.1 of ECS.

• Retention and Expiration on Atmos Objects through the Atmos (UMD)	66
• Support for Geo-Passive architecture	66
• Support for hosted sites	66
• S3 bucket policies	66
• Read-only access to buckets during an outage	67
• Metadata search and D@RE	67
• Swift and S3 interoperability	67
• Network support improvements	67
• Application registration for CAS API is disabled by default	69
• User tags	69
• Partial garbage collection (space reclamation) is enabled by default	69
• Secure Remote Services support for FOB-based passwords	69
• ECS Service Console	69
• Changes in the documentation set	70

Retention and Expiration on Atmos Objects through the Atmos (UMD)

ECS supports setting retention periods on an Atmos object through the User Meta Data (UMD), and setting an expiration date on an Atmos object using either the ECS Header (x-emc), or Atmos UMD (user.maui).

Retention periods define how long an object is retained by ECS before it can be edited or deleted. When a retention period end date is defined for an Atmos object, and an expiration period is also set on the object, ECS automatically deletes the object at the date that is defined in the expiration period.

In previous ECS releases, ECS did not support setting an expiration time on an Atmos object, and the retention periods could only be set on Atmos objects using the ECS header. For details see the *ECS Data Access Guide*.

Support for Geo-Passive architecture

ECS supports an Active-Active architecture in which replication occurs to the active sites. For ECS 3.1 an extra architecture, Geo-Passive is available. The Geo-Passive architecture always comprises three sites; two active sites and one replication-only site.

Geo-Passive replication enables customers to use an ECS site as a pure backup site (also called the replication target) and enables other sites (called replication sources) in the ECS federation to back up data into the replication target. The improved storage efficiency provided by ECS when using three or more sites is maintained. In this configuration, data can be read from and written to the replication source sites, however, data cannot be written to the replication target site.

Where a hosted site is present, it is automatically selected as the target for a Geo-Passive configuration.

Support for hosted sites

Before ECS 3.1, ECS supported the ability to use a hosted site as part of an ECS federated system, however, the hosted site would not be recognized as any different to an on-premise VDC.

In 3.1, ECS software identifies a site that is hosted and marks sites as Hosted or On-Premise sites. Hosted sites can be part of any of the three main replication architectures. However, where a Geo-Passive architecture is chosen, ECS always expects the replication site to be the hosted site.

S3 bucket policies

ECS now supports S3 bucket policies which provide fine grained control over the operations on buckets and objects that can be performed by an ECS object user. The ECS Portal provides an editor that enables a policy file to be associated with a bucket. Also, support for adding a policy file to a bucket resource is provided from the ECS Management REST API or using the S3 object protocol.

Read-only access to buckets during an outage

ECS now supports the ability to define the type of access you have to the objects in a bucket during a temporary site outage (TSO) by enabling or disabling the **Read-Only Access During Outage** property on a bucket. When you create a bucket and enable the **Access During Outage** property, you now also have the option of enabling the **Read-Only Access During Outage** property on the bucket. You can only set the **Read-Only Access During Outage** property while creating the bucket; you cannot change this property after the bucket has been created. When you enable the **Read-Only Access During Outage** property, the following occurs during a TSO:

- Creation of new objects in the bucket is restricted.
- Access to file systems is not impacted since they are automatically put into read-only mode when **Access During Outage** is set on the file system buckets.

You can set the **Access During Outage** and **Read-Only Access During Outage** properties when creating a bucket from the following interfaces:

- ECS Portal
- ECS Management REST API
- ECS CLI
- Object API REST interfaces such as S3, Swift, and Atmos

Metadata search and D@RE

Metadata search can now be configured on buckets that are D@RE encryption enabled.

Swift and S3 interoperability

ECS now supports the ability for objects and buckets that are created using S3 protocol support to be accessed by Swift applications and for objects and buckets that are created using the Swift protocol support to be accessible from S3 applications. It is sometimes referred to as cross-head operation.

Network support improvements

ECS has improved network support as follows:

- [Ability to separate networks after ECS is installed and running](#)
- [Policy based routing for network separation](#)
- [Additional data network for Centera systems](#)
- [IP address change on ECS nodes](#)

Also, the *ECS Networking and Node IP Change Guide* document has been added to the Solve documentation. Contact your ECS technical support professional if you would like to perform any of these networking operations.

Ability to separate networks after ECS is installed and running.

The Management, Replication, and Data networks can be separated in existing ECS environments.

Policy-based routing for network separation

Policy-based routing has been implemented for network separation. Policy-based routing enables a default route to automatically be configured for each interface when one or more of the ECS networks has been separated from the public network.

In releases earlier than 3.1, static routes were manually configured through the `/etc/sysconfig/network/ifroute-public.<interface>` file as demonstrated in the following example where the management interface is separated from the public interface.

```
cat /etc/sysconfig/network/ifroute-public.mgmt
10.100.100.0    10.100.100.1    255.255.255.0    public.mgmt

/etc/sysconfig/network> getrackinfo --static-route-list
Static route list
=====
NodeID  Network          Netmask          Gateway          Interface
1       10.100.100.0     255.255.255.0    10.100.100.1    public.mgmt
2       255.255.255.0    10.100.100.1     10.100.100.0    public.mgmt
3       10.100.100.0     255.255.255.0    10.100.100.1     public.mgmt
4       255.255.255.0    10.100.100.1     10.100.100.0    public.mgmt
```

Implementing policy-based routing for network separation has removed the need to configure multiple static routes for each interface through the `ifroute-public.<network>` file. Also, in ECS 3.1 and higher, the following services no longer require a static route in a network separated environment: SMTP, DNS, NTP, LDAP, sLDAP, and ECS Geo Replication.

If you have configured separate networks in versions earlier than ECS 3.1, during upgrade to ECS 3.1, the file is removed from your configuration and the settings that are defined in the file are imported and managed through the NAN.

If you are upgrading to ECS 3.1.0.0, see the "Installation and upgrade" section of the *ECS 3.1.0.0 Release Notes* before performing the upgrade.

Additional network for data traffic for Centera systems

ECS supports a second network for data traffic for Centera systems when network separation is configured.

When there is no network separation, or the data traffic is not configured on separate networks, all data traffic is run through the same network.

IP address change on ECS nodes

You can change the IP addresses of ECS nodes.

The following, however, is not supported with node IP change:

- You cannot change the hostname or FQDN of a node.
- Changing private IP addresses (192.168.219.xxx) is not supported.
- You cannot change node IPs while upgrading ECS. You must upgrade, and then change the node IP addresses after the upgrade is complete.

Application registration for CAS API is disabled by default

In ECS, application registration information is stored in the namespace record. Per ECS design, the namespace record is updated every time there is a request for the pool information. Having the application registration information that is enabled slows down the namespace update process.

If you would like to have the application registration enabled, contact your ECS technical support professional.

User tags

ECS now enables you to add tags in the form of name=value pairs to an ECS object user. The tags can be used to associate information, such as project or cost-center membership, with the user. Tags must be written and read using the ECS Management REST API.

Partial garbage collection (space reclamation) is enabled by default

Partial garbage collection (GC) offers higher storage efficiency for use cases involving random deletion or modification of data. This feature complements the existing full chunk garbage collection in ECS. Partial GC is achieved by copying over valid data from partially empty source chunks to a new chunk so that the source chunks can be freed up. The new chunk is protected and replicated as usual. This technique is known as partial GC by merging. Starting with the ECS 3.1 release, partial GC by merging is enabled by default on all deployments.

Secure Remote Services support for FOB-based passwords

FOB-based passwords are now supported when configuring Secure Remote Services for ECS.

ECS Service Console

The ECS Service Console is a new command-line tool that simplifies and automates various ECS service procedures.

For ECS 3.1.0.x, the Service Console automates health checks and upgrade procedures depending on the following method:

- Sequential rolling upgrade with Service Console: In this process, the full stack (OS, fabric, and object) is upgraded on one node before proceeding to the next node (node by node). The Service Console is used to upgrade both the OS and services. The *Rolling Upgrade Guide* provides complete instructions for a sequential rolling upgrade with the Service Console.
- Blackout upgrade with Service Console (sequential mode): In this process, the offline OS update of all nodes follows the same manual process that was used in prior releases. The Service Console is used for health checks and for automating the ECS services upgrade. The *Blackout Upgrade Guide* provides complete instructions for both the services upgrade using the Service Console and the offline OS update.
- Blackout upgrade with Service Console (parallel mode): In this process, the services (fabric and object) on all nodes are upgraded in parallel. The OS update is performed using the blackout/offline procedure and then all steps for services upgrade are performed

simultaneously. The *Blackout Upgrade Guide* provides complete instructions for the parallel upgrade.

The ECS Service Console upgrades ECS 3.0.0.x to ECS 3.1.0.x. Also, you can use the Service Console to perform health checks before and after any service procedure.

The ECS Service Console software is provided as a separate .tgz file which you install on Node 1 of each site.

The ECS Service Console does not replace the Compatibility Checker which is a prerequisite for ECS installations.

For information about installing the Service Console and performing upgrades with it, consult the *Blackout Upgrade Guide* or *Rolling Upgrade Guide*, depending on your upgrade method. Both guides are available from the SolVe Desktop.

Obtain the ECS Service Console installer .tgz file from https://support.emc.com/downloads/37236_ECS-Appliance-Software.

Changes in the documentation set

The ECS 3.1 documentation set has been reorganized and differs from the 3.0 documentation set as follows.

New documents

- *ECS 3.1 Monitoring Guide*
The monitoring content in the *ECS 3.0 Administrator's Guide* has been moved into a separate guide for the 3.1 release.
- *ECS 3.1 New Features and Changes Guide*
- *ECS Gen1 and Gen2 Configuring the RMM Interface on the Private Arista Switch Guide* (Dell EMC internal resource)
- *Rolling Upgrade to ECS 3.1 Guide* (Dell EMC internal resource)
- *Blackout Upgrade to ECS 3.1 Guide* (Dell EMC internal resource)
The three ECS 3.0 upgrade guides (*OS Online Update Guide*, *OS Offline Update Guide*, and *Fabric and Object Services Upgrade Guide*) have been consolidated into the ECS 3.1 Rolling Upgrade and Blackout Upgrade Guides.

Retired documents

The ECS 3.1 documentation set does not include a planning guide. The content from the *ECS 3.0 Planning Guide* has been redistributed into the *ECS 3.1 Administration Guide* and the *ECS 3.1 Networks Guide*.

Changed documents

After the ECS 3.1.0.x releases, the content from the *ECS 3.1 Networks and Node IP Change Guide* has been redistributed into the following two documents. Both documents are Dell EMC internal resources.

- *ECS 3.1 Networks Guide*
- *ECS 3.1 Changing the Node IP Addresses Guide*

CHAPTER 8

ECS Version 3.0 and 3.0 Hotfixes

This chapter lists the features and changes that were introduced in ECS version 3.0 and 3.0 hotfix versions.

- [ECS 3.0 HF3 improvements](#)..... 72
- [ECS 3.0 HF2 improvements](#)..... 72
- [ECS 3.0 new features and changes](#) 73

ECS 3.0 HF3 improvements

This section describes the improvements that are provided with ECS version 3.0 hotfix 3 (HF3).

ECS 3.0 HF3 includes all patch fixes from ECS 3.0 HF2 on until Aug 25, 2017. If you received a patch after Aug 25, contact your ECS technical support professional.

ECS 3.0 HF2 improvements

This section describes the improvements that are provided with ECS version 3.0 hotfix 2 (HF2).

ECS 3.0 HF2 includes all patch fixes from ECS 3.0 HF1 on until May 14, 2017. If you received a patch after May 14, contact your ECS technical support professional.

Storage engine resiliency improvements

- Manage Race condition in index update to avoid corruption.
- Manage Race condition for encryption key management.
- Manage capacity allocation delay on node restart.

Related fixed issues: Storage-17625, Storage-17578, Storage-17678

Garbage collection resiliency improvements

Prevent slow disks from having ripple effects on read/write performance of entire node.

Related fixed issues: Storage-17598, Storage-17599, Storage-17590, Storage-17479

Resource utilization improvements

- Prevent slow disks from having ripple effects on read/write performance of the entire node.
- Fine-tuning memory utilization limits for object container services.

Related fixed issues: Storage-17463, Storage-17779, Storage-17865

Access heads related fixes and enhancements

- Atmos API fixes for access during TSO and after PSO
- Efficient S3 listing for FS enabled buckets
- S3 versioning fix for resiliency during TSO
- NFS-write performance improvements by batching based on file size
- Ensure that SWIFT small file performance is on par with S3
- CAS resiliency improvements for large blob sizes (>2 MB)

Related fixed issue: Storage-17664, Storage-17673, Storage-17608, Storage-17870, Storage-17579, Storage-17560, Storage-17617

ECS portal fixes

- Address case sensitivity for AD user group mapping
- Fix for removing a failed VDC from a replication group

Related fixed issues: Storage-17604, Storage-17783

Install and upgrade fixes

- Support parallel All-Nodes-At-Once upgrade mode
- VDC bootstrap resiliency improvements

- Support for CoreOS 1235.6.0 and Docker versions 1.12.6

Related fixed issues: Fabric-4862, Storage-17558, Fabric-4499, Fabric-4863

Centra transformation fixes

- Support interoperability of TSO and migration
- Improve stability and indexing performance during transformation

Related fixed issues: Storage-17681, Storage-17612, Storage-17656, Storage-17669

Storage server statistics

ECS can now collect the number of successful, and failed requests storage server metrics.



Related fixed issue: Storage-17586

ECS 3.0 new features and changes

The following new features and changes were introduced with the ECS 3.0 release.

S3 Protocol enhancements

Support for S3 protocol includes the following enhancements.

- S3 protocol support for V4 authentication
- S3 protocol support for life cycle on versioned objects
ECS support for the S3 protocol now includes support for life cycle on versioned objects so that it is possible to specify when object versions, in version-enabled buckets, expire. This feature enables old versions of objects to be deleted after a specified time.
 **Note:** Lifecycle cannot be applied to filesystem-enabled buckets.
- S3 protocol metadata search enhancements
For ECS 3.0, the number of keys that can be indexed has been increased 5–30. The keys can comprise both system metadata keys and user metadata keys.
 **Note:** For small objects (100 KB and below), the ingest rate for data slightly reduces on increasing the number of index keys. Performance testing data showing the impact of using metadata indexes for smaller objects is available in the ECS Performance white paper.

OpenStack Swift protocol support for Dynamic Large Objects (DLOs) and Static Large Objects (SLOs)

ECS OpenStack Swift support now provides support for SLOs and DLOs. SLO support enables many objects to be uploaded, associated with the same SLO using a manifest file, and downloaded as a single object.

DLO support enables multiple objects to be uploaded and assembled into a DLO based on order and key prefixes and for the object to be downloaded as a single object.

Both DLOs and SLOs accept range requests that enable byte ranges of large objects to be retrieved.

Support for sending SNMP traps from ECS

ECS supports reporting of node-level statistics using SNMP queries. With ECS 3.0, SNMP support has been extended to enable sending SNMP v2 and v3 traps to up to 10 Network Management

Station clients. The SNMP traps report ECS events. The ECS Portal enables administrators to configure the trap recipients.

Note: The SNMP query server and trap server are separate servers. The trap server does not respond to SNMP queries.

You can download the ECS-MIB definition (as the file `ECS-MIB.mib`) from the ECS area on support.emc.com in the Downloads section under Add-ons.

ECS supports these possible SNMPv3 cryptographic hash functions:

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

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) that the ECS SNMP Trap feature provides, might be incompatible with certain SNMP targets (for example iReasoning).

For more details about ECS SNMP, see the Event Notification chapter in the *ECS Administrator's Guide*. The *ECS Installation Guide* provides details on configuring SNMP servers to support queries.

Support for Remote Syslog Servers

ECS supports forwarding of alerts and audit messages to remote system log servers, and supports operations using these application protocols:

- BSD Syslog
- Structured Syslog

Space Reclamation by Partial Garbage Collection

Partial garbage collection (GC) offers higher storage efficiency for active datasets and use cases involving random deletion or modification of data. This feature complements the existing full chunk garbage collection in ECS, which supports bulk deletion use cases common for archival data under retention.

Two new techniques have been introduced to reclaim space from partially deleted chunks:

- Partial Garbage Collection By Merging - Valid data, from partially empty source chunks, is copied over to a new chunk such that the source chunks can be freed up. This technique is employed when valid data occupies a small portion of the chunks.
- Partial Garbage Collection by Compaction - Disk ranges containing garbage data are extracted from a chunk and the remaining valid data is treated as a smaller chunk. This technique is employed when valid data occupies a large portion of the chunk.

Global Propagation: Space reclamation-related changes from partial GC are also propagated to replicated copies of the data. In environments with geo replication, as and when new data is ingested, it is replicated by encoding it with garbage ranges so that newly replicated data overwrites the corresponding garbage ranges on the remote sites. It ensures that there is no WAN traffic overhead from partial garbage collection. Reclaimed space becomes available for reuse in the local site, only after newly ingested data is encoded with associated garbage data ranges.

The partial garbage collection feature is released with limited availability in the ECS 3.0 release and can be enabled on demand. This feature is targeted for deployments with active workloads and high utilization. To have partial GC enabled, log a service request after an upgrade to ECS 3.0.

Limited availability of partial GC is in line with the best practice of rolling out features that entail significant architectural changes with limited availability for at least one release before the GA.

Platform Locking

Allowing remote access to nodes from privileged accounts (root, admin, emc) may not be desirable. Using the ECS Portal or the ECS Management REST API, you can lock specific nodes in a cluster or all the nodes in the cluster. Locking a node only affects the ability to remotely SSH to the locked nodes. The lock does not change the way the ECS Portal and REST APIs operate and it does not affect the ability to directly connect to a node and log in using a privileged account.

The new Lock Admin user is a preprovisioned local user called "emcsecurity". Lock Admins can only change their passwords and lock and unlock nodes. System Admins and System Monitors can view the lock status of the nodes. The Lock Admin role cannot be assigned to another user.

Dashboard and Monitoring Improvements

New features include:

- Traffic Metrics panel on the ECS Portal Dashboard: Traffic Metrics displays total requests and breaks that down into successful requests and failed requests by user error and failed requests by system error. User errors are typically known errors from the various object heads (HTTP 400-level errors) and system errors are ECS errors (HTTP 500-level errors).
- Error reporting on the Traffic Metrics monitoring page: Click **History** and select the **Failures by type** panel. ECS lists all failed requests that are sorted by user or system error, and then by the head (object type) with the more frequent error codes sorted to the top.
- The Capacity Utilization page has been redesigned to make it easier to distinguish between online capacity, offline capacity, and total capacity.
- The Hardware page now includes more disk states: `missing` and `removed`. A missing disk is a known disk that is unreachable. The disk may be transitioning between states, disconnected, or pulled. A removed disk is one that the system has completed recovery on and removed from the storage engine's list of valid disks.

CAS Advanced Retention Management

ECS now supports the advanced retention management (ARM) features available through the CAS API. These features do not require a separate license in ECS. ARM features include:

- Event-based retention: The ability to configure an object through its C-Clip to apply (trigger) a retention period or retention class when the CAS application receives a specified event.
- Litigation hold: The ability to prevent deletion of an object if the CAS application has applied a litigation hold to the object through its C-Clip. The CAS application can apply up to 100 litigation holds to an object by creating and applying unique litigation hold IDs.
- Min/Max governor: The ability for an administrator to set a bucket-level fixed retention period or variable retention period. A variable retention period is one that is set in response to an event. In ECS, System or Namespace Admins can set the values with the ECS Portal. Programmers can use the ECS Management API to set the values.

CAS Behavior Change for Default Retention Period in Objects Written without Object-level Retention in Compliance Namespaces

Starting with ECS 3.0, when an application writes C-Clips with no object retention to an ECS CAS bucket in a Compliance namespace, and the bucket has a retention value (6 months, for example), the default retention period of infinite (-1) is assigned to the C-Clips. The C-Clips cannot be

deleted because their effective retention period is the longest one between the two: the bucket-level retention period and the default object-level retention.

It is a change from ECS 2.2.1 behavior which brings ECS in line with Centera behavior, where default pool retention in CE+ Compliance mode is always infinite (-1).

In ECS 2.2.1, when an application writes C-Clips with no object retention to an ECS CAS bucket in a Compliance namespace, and the bucket has a retention value (6 months, for example), the retention period that is assigned to the C-Clips are zero (0). Here, the effective retention period for the C-Clips is the bucket retention value (6 months). The C-Clips can be deleted in 6 months.

After upgrading from ECS 2.2.1 to ECS 3.0 or any later version, applications that rely on the ECS 2.2.1 behavior will be writing C-Clips that cannot be deleted.

Halt and reconfigure your applications to assign appropriate object-level retention before they interact with ECS 3.0 or later.

In the example above, the application should assign 6-month object-level retention to the C-Clips.

ECS Compliance certified for ECS Appliances with ECS 3.0 and ECS Software Only installations on ECS-certified third-party storage hardware

ECS meets the storage requirements of the following standards, as certified by Cohasset Associates Inc:

- Securities and Exchange Commission (SEC) in regulation 17 C.F.R. § 240.17a-4(f)
- Commodity Futures Trading Commission (CFTC) in regulation 17 C.F.R. § 1.31(b)-(c)

Compliance is certified on ECS Appliances with ECS version 2.2.1 software and later. Installations of ECS Software Only version 3.0 and later running on ECS-certified third-party hardware are also certified.

Atmos Support Improvements

ECS now supports the following Atmos features:

- Shareable URL: The CIFS-ECS tool creates a shareable URL that enables a user to retrieve a file directly from ECS. ECS now supports the Atmos API Shareable URLs feature enabling the use of shared URLs in the CIFS-ECS tool.
- Checksums: The x-emc-wschecksum header is now supported.
- Subtenant IDs: These IDs are now preserved in ECS after migration: The header is `x-emc-subtenant-id: {original_subt_id}`.
- Read, Write, and Delete ACLs: These ACLs for Atmos data now work on ECS the same as Atmos.
- Indexed listable tags: The performance of the Atmos listable tags feature has been improved in ECS 3.0 by using an index. Listable tags are optional object-level tags that enable for retrieving and sorting of all tagged objects.

CIFS-ECS Support

ECS now provides a lightweight application that enables applications and users to have access to content held in ECS from a local Windows drive.

Files are maintained in a local disk cache on the Windows machine and uploaded and downloaded to and from an ECS bucket using the S3 API. Applications can create, modify, and read files using CIFS-ECS's virtual drive and files are automatically uploaded asynchronously to ECS.

More information about the CIFS-ECS is provided in the *CIFS-ECS Architecture, Performance, and Best Practices White Paper*, which can be accessed from the [ECS 3.0 Product Documentation Index](#).

Network separation

In a standard configuration the ECS management, replication, and data network traffic is configured on the same public 10-GbE interface. Optionally, the traffic can be separated to run on dedicated networks.

Note: In the ECS 3.0 release network, separation can only be performed during a new installation.

Network separation enables:

- Data, Management and/or Replication traffic (to other geographies) to be separately identified on different networks
- Traffic to be separated between the traffic generated by clients that want to access ECS using both ECS Software (ECS Management REST API calls and ECS Portal) and the client data access protocols. Similarly, replication traffic between sites can be identified separately from the data and control plane traffic.

Note: Most ECS installations are configured with the data, management, and replication traffic on a single network. Network separation should only be configured when there is an explicit requirement to separate one or more of the ECS networks. Physical network separation support, where unique uplinks are configured for each separated network, is subject to the EMC Request for Product Qualification (RPQ) process.

ECS Software

ECS Software (ECS SW) is a new offering that enables ECS software to be installed on precertified Dell or HP hardware.

The following hardware has been certified for use with ECS SW and Reference Architecture papers for each certified platform can be obtained from support.emc.com.

- HP Proliant SL4540 Gen8
- Dell DSS7000
- Dell R730xd 13G

Customer Support can use the SolVe desktop to obtain configuration check and health check tools to support customer engagements.

Customers that want to use custom hardware can contact their ECS representative to discuss their requirements.

