

ECS

Version 3.0 and higher

New Features and Changes

302-003-825

A06

Copyright © 2017-2018 Dell Inc. or its subsidiaries. All rights reserved.

Published August 2018

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

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS-IS." DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

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

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 | 5 |
| Chapter 2 | ECS Version 3.2.2 | 7 |
| | Dell EMC Generation 3 (Gen3) Hardware..... | 8 |
| | Sudo required for cs_hal led command on Gen3 hardware..... | 10 |
| | ECS Software..... | 10 |
| Chapter 3 | ECS Version 3.2.1 | 11 |
| | CAS monitoring of unused objects..... | 12 |
| | Large object support improvements for Data Domain Cloud tier..... | 12 |
| | Configuration parameter for NFS directory listing | 12 |
| Chapter 4 | ECS Version 3.2 | 15 |
| | ECS Portal changes..... | 16 |
| | ESRS improvements..... | 16 |
| | Monitoring improvements..... | 16 |
| | More active sites supported for the ECS Passive replication configuration.... | 17 |
| | Log improvements..... | 17 |
| | Additional software alerts..... | 18 |
| | HDP 2.6.2 support..... | 18 |
| | Large object support for Data Domain Cloud Tier..... | 19 |
| | ECS Software installation improvements..... | 19 |
| | Licensing..... | 19 |
| | Centera migration..... | 20 |
| | Ability to enable encryption for migration of Centera data to ECS | 20 |
| | Transform service (transformsvc) is disabled by default..... | 20 |
| | Documentation changes..... | 20 |
| Chapter 5 | ECS Version 3.1 | 21 |
| | Retention and Expiration on Atmos Objects through the Atmos (UMD)..... | 22 |
| | Support for Geo-Passive architecture..... | 22 |
| | Support for hosted sites..... | 22 |
| | S3 bucket policies..... | 22 |
| | Read-only access to buckets during an outage..... | 23 |
| | Metadata search and D@RE..... | 23 |
| | Swift and S3 interoperability..... | 23 |
| | Network support improvements..... | 23 |
| | Ability to separate networks after ECS is installed and running. ... | 24 |
| | Policy based routing for network separation | 24 |
| | Additional network for data traffic for Centera systems..... | 24 |
| | IP address change on ECS nodes..... | 24 |
| | Application registration for CAS API is disabled by default..... | 25 |
| | User tags..... | 25 |
| | Partial garbage collection (space reclamation) is enabled by default..... | 25 |
| | ESRS support for FOB-based passwords..... | 25 |

| | | |
|------------------|---|-----------|
| | ECS Service Console..... | 25 |
| | Changes in the documentation set..... | 26 |
| Chapter 6 | ECS Version 3.0 and 3.0 Hotfixes | 29 |
| | ECS 3.0 HF3 improvements..... | 30 |
| | ECS 3.0 HF2 improvements..... | 30 |
| | ECS 3.0 new features and changes | 31 |
| | S3 Protocol enhancements..... | 31 |
| | Openstack Swift protocol support for Dynamic Large Objects (DLOs) and Static Large Objects (SLOs)..... | 32 |
| | Support for sending SNMP traps from ECS..... | 32 |
| | Support for Remote Syslog Servers..... | 32 |
| | Space Reclamation by Partial Garbage Collection..... | 32 |
| | Platform Locking..... | 33 |
| | Dashboard and Monitoring Improvements..... | 33 |
| | CAS Advanced Retention Management..... | 34 |
| | CAS Behavior Change for Default Retention Period in Objects Written without Object-level Retention in Compliance Namespaces | 34 |
| | ECS Compliance certified for ECS Appliances with ECS 3.0 and ECS Software Only installations on ECS-certified third-party storage hardware..... | 35 |
| | Atmos Support Improvements..... | 35 |
| | CIFS-ECS Support..... | 35 |
| | Network separation..... | 35 |
| | ECS Software..... | 36 |

CHAPTER 1

New Features and Changes Overview

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

CHAPTER 2

ECS Version 3.2.2

- [Dell EMC Generation 3 \(Gen3\) Hardware](#)..... 8
- [Sudo required for cs_hal led command on Gen3 hardware](#)..... 10
- [ECS Software](#)..... 10

Dell EMC Generation 3 (Gen3) Hardware

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

Table 1 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>In addition to the Dell EMC-racked EX300 appliance, the EX300 nodes can be installed in customer-provided racks. For more information on third-party racking requirements, see the <i>ECS EX300 Third-Party Rack Installation Guide</i>.</p> | <p>Customer-provided 40U rack 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 1200mm cabinet Contain Gen3 2U PDUs supporting single phase, three-phase delta, and three-phase WYE. <p>For more information on 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 (OS) 10.3.2</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 (OS) 10.3.2</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 OS 10.3.2, 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 OS 10.3.2, 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>HDD 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>Twelve SATA hard disk drives (HDDs) in each node.</p> <p>480 GB M.2 (BOSS) system disk in each node.</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> <p>One and two-node chassis cannot be mixed together in a rack. Within a single rack, chassis must be of the same one- or two-node configuration (that is, a rack</p> |

Table 1 EX300 and EX3000 appliance hardware components (continued)

| Component | EX300 appliance | EX3000 appliance |
|-----------|---|--|
| | <p>64 GB RAM per node</p> <p>Single 8-core SkyLake CPU per node. Xeon Bronze 3106 8 core/8 thread-11MB L3, 1.7GHz, 85W</p> <p>4x 16GB RDIMM, 2667MT/s, Dual Rank, x4 Data Width</p> <p>Dual 750W Platinum power supply (hot swappable)</p> <p>Each node has 4 x 10 GbE networking</p> | <p>must contain all EX3000S nodes or all EX3000D nodes).</p> <p>Chassis have the following disk configurations (all HDD 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 midplane routing to 90 drive slots. Filler is in the second server sled.</p> <p>EX3000D node has dual server sleds with the midplane 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 2.1GHz 20M cache 85W</p> <p>4x 16GB RDIMM, 2400MT/s, Dual Rank, x8 Data Width</p> <p>Dual 1600W 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 four nodes for performance optimization.

Node upgrade requests for less than four nodes require a Request for Product Qualification (RPQ).

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 four nodes of any drive size (1 TB, 2 TB, 3 TB or 4TB). 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 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 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 execute the cs_hal led command without sudo, you will 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 execute the sudo cs_hal led command to successfully complete the operation, as shown below.

```
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 prior to 3.2.2 ECS software, could be installed on certified, and custom hardware installations, as well as 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 3

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](#) 12
- [Large object support improvements for Data Domain Cloud tier](#) 12
- [Configuration parameter for NFS directory listing](#) 12

CAS monitoring of unused objects

The ECS 3.2.1 release includes the CAS (Content Addressable Storage) 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 completes, the removal of CAS garbage data will have to be enabled by support. You must 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 will not change on the **CAS Processing** page. For example, you will see zero for the unreferenced blob data 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 allow 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 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 allows 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 customer support to increase listing performance.

This operation can only be performed by an ECS customer support representative with `emcservice` credentials.

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

```
hostname:/opt/storageeos/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 customer support representative logged in as the `emcservice` user.
- `<emcservice_password>` is the password provided for the `emcservice` user.

There is no output if the command was run successfully.

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

```
hostname:/opt/storageeos/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 4

ECS Version 3.2

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

- [ECS Portal changes](#)..... 16
- [ESRS improvements](#)..... 16
- [Monitoring improvements](#)..... 16
- [More active sites supported for the ECS Passive replication configuration](#)..... 17
- [Log improvements](#)..... 17
- [Additional software alerts](#)..... 18
- [HDP 2.6.2 support](#)..... 18
- [Large object support for Data Domain Cloud Tier](#)..... 19
- [ECS Software installation improvements](#)..... 19
- [Licensing](#)..... 19
- [Centera migration](#)..... 20
- [Documentation changes](#)..... 20

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 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 allow you to set configurable available capacity thresholds that will 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 of the endpoints in a GEO configuration after the networks have been separated, or after the IP addresses of the nodes have been changed.

ESRS improvements

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

- Automation of ESRS initial setup. Previous releases required manual steps to set up and configure ESRS. In the 3.2 release, there are no manual steps; you can set up and configure ESRS 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 ESRS 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 also be retrieved using the ECS Management REST API. ECS reports whether garbage collection is enabled for user data and system metadata, total garbage detected, the capacity reclaimed (which can be further broken down into user data 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 also 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 also 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 additional 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 log file is located in `/opt/emc/caspian/fabric/agent/services/object/main/log`. Prior to 3.2, the log file was named `datahead-access.log`. In 3.2, the log file name changed from `datahead-access.log` to `dataheadsvc-access.log`. If you upgrade from 3.1.0.x to 3.2, this means that after the upgrade you will 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 customer 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 will trigger storage pool capacity alerts. Capacity alerts can be configured with severities of Critical, Error, and Warning. When a capacity alert is generated, and ESRS is configured a call home alert is also generated that alerts ECS customer 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 will need to 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 allow 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 allows 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 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

Prior to ECS 3.2, when migrating Centera data to ECS, it was not possible to securely store Data at Rest (that is, data 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, allowing Data at Rest (D@RE) to be securely stored. This allows 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 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 will be disabled after upgrade.
- 3.2 to 3.2.0.x, transformsv will be in the state it was left prior to upgrade. For example, if transformsv was enabled in the earlier release, prior to 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* document:

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

CHAPTER 5

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 \(UMD\)](#)22
- [Support for Geo-Passive architecture](#).....22
- [Support for hosted sites](#).....22
- [S3 bucket policies](#)..... 22
- [Read-only access to buckets during an outage](#)..... 23
- [Metadata search and D@RE](#)..... 23
- [Swift and S3 interoperability](#).....23
- [Network support improvements](#).....23
- [Application registration for CAS API is disabled by default](#).....25
- [User tags](#)..... 25
- [Partial garbage collection \(space reclamation\) is enabled by default](#)..... 25
- [ESRS support for FOB-based passwords](#).....25
- [ECS Service Console](#)..... 25
- [Changes in the documentation set](#).....26

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), as well as 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 will be 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 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 refer to 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 additional 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 allows 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 will automatically be selected as the target for a Geo-Passive configuration.

Support for hosted sites

Prior to 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 will always expect 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. In addition, 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. Note that 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 created using S3 protocol support to be accessed by Swift applications and for objects and buckets created using the Swift protocol support to be accessible from S3 applications. This 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](#)

Additionally, the *ECS Networking and Node IP Change* document has been added to the Solve documentation. Contact your ECS customer support representative 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 allows 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 prior to 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       255.255.255.0          10.100.100.1          10.100.100.0          public.mgmt
2       255.255.255.0          10.100.100.1          10.100.100.0          public.mgmt
3       255.255.255.0          10.100.100.1          10.100.100.0          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. Additionally, in ECS 3.1 and higher, the following services will 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 prior to ECS 3.1, during upgrade to ECS 3.1, the file is removed from your configuration and the settings defined in the file will be imported and managed through the NAN.

If you are upgrading to ECS 3.1.0.0, please refer to the *Installation and upgrade* section of the *ECS 3.1.0.0 Release Notes* prior to 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 perform the 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 enabled slows down the namespace update process.

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

User tags

ECS now provides the ability 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 a set of 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.

ESRS support for FOB-based passwords

FOB-based passwords are now supported when configuring ESRS (EMC Secure Remote Support) 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 of the nodes are upgraded in parallel. The OS update is performed using the blackout/offline procedure and then all steps for services upgrade are performed at the same time. 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* (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 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*

CHAPTER 6

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](#)..... 30
- [ECS 3.0 HF2 improvements](#)..... 30
- [ECS 3.0 new features and changes](#) 31

ECS 3.0 HF3 improvements

This section highlights the improvements provided with ECS version 3.0 hot fix 3 (HF3).

ECS 3.0 HF3 includes all patch fixes from ECS 3.0 HF2 on until August 25, 2017. If you received a patch after August 25, contact your ECS customer support representative.

ECS 3.0 HF2 improvements

This section highlights the improvements provided with ECS version 3.0 hot fix 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 customer support representative

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, Storage17678

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 SWIFT small file performance is on par with S3.
- CAS resiliency improvements for large blob sizes (>2MB).

Related fixed issue: Storage-17664, Storage17673, 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 I 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

Centera 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 lifecycle on versioned objects
ECS support for the S3 protocol now includes support for lifecycle on versioned objects so that it is possible to specify when object versions, in version-enabled buckets, will expire. This feature enables old versions of objects to be deleted after a specified period of 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 from 5 to 30. The keys can comprise both system metadata keys and user metadata keys.

Note

In the case of small objects (100KB 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) provided by the ECS SNMP Trap feature may be incompatible with certain SNMP targets (e.g. iReasoning).

For more details about ECS SNMP, refer to 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 syslog 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 data sets 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 a set of 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. This 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 prior to 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 pre-provisioned 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 sorted by user or system

error, 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 currently 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) will be assigned to the C-Clips. The C-Clips can never 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.

This 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 assigned to the C-Clips will be zero (0). Here, the effective retention period for the C-Clips will be 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 can never 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 allows 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 significantly improved in ECS 3.0 by using an index. Listable tags are optional object-level tags that allow for easy 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 normally using CIFS-ECS's virtual drive and files are automatically uploaded asynchronously to ECS.

More information on 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 allows for:

- 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 pre-certified 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.