elastifile
Cross-Cloud Data Fabric

# ECFS (Elastifile Cloud File System) 3.1.X

## Management Console

## Google Cloud Platform (GCP)

User Guide

**May 2019**

**Document Revision: 0.2**

**Important Notice**

This document is delivered subject to the following conditions and restrictions:

- This guide contains proprietary information belonging to Elastifile Inc. Such information is supplied solely for the purpose of assisting explicitly and properly authorized users of Elastifile Inc. products.
- No part of contents may be used for any other purpose, disclosed to any person or firm, translated or reproduced by any means, electronic and mechanical, without the express prior written permission of Elastifile Inc..
- The text and graphics are for the purpose of illustration and reference only, based on the current version of the product(s) described in this document.
- The software described in this document is furnished under a license agreement. The software may be used or copied only in accordance with the terms of that agreement.
- Information in this document is subject to change without notice. Corporate and individual names and data used in examples herein are fictitious unless otherwise noted.
- Elastifile Inc. makes no warranty of any kind with regard to this printed material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Elastifile Inc. shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.
- Brand or product names are trademarks of their respective companies or organizations.

# Table of Contents

**1. Introduction** .......... **6**

1.1 Document Scope .......... 6

1.2 Related Documentation .......... 6

1.3 Customer Support .......... 6

1.4 ECFS's Call Home Feature .......... 6

**2. Logging in to ECFS** .......... **8**

**3. Getting to Know the Management Console Interface** .......... **9**

3.1 Main Interface .......... 9

3.2 Main Entities .......... 10

3.2.1 Assets Menu .......... 10

3.2.2 Header .......... 10

3.2.3 FAB .......... 10

3.3 Working with the Management Console Interface .......... 11

**4. Provisioning** .......... **12**

4.1 Understanding Provisioning using ECFS .......... 12

4.2 Provisioning Data Containers .......... 12

4.2.1 About Data Containers .......... 12

4.2.2 Creating a Data Container .......... 12

4.2.3 Mounting Your Storage .......... 15

4.2.4 Modifying Data Container Properties .......... 15

4.2.4.1 Adding Data Container Shares .......... 16

4.2.4.2 Editing Data Container Shares .......... 18

4.2.5 Deleting Data Containers .......... 19

4.2.6 Creating Data Container Snapshots and Schedules .......... 19

4.2.6.1 Creating Data Container Snapshots .......... 20

4.2.6.2 Moving a Snapshot to the Object Tier .......... 21

4.2.6.3 Creating a Data Container Snapshot Schedule .......... 22

4.2.6.4 Enabling and Disabling the Snapshot Schedule .......... 24

4.2.6.5 Adding Snapshot Shares ........ 26
4.2.6.6 Editing Snapshot Shares ........ 28
4.2.6.7 Deleting Snapshot Shares ........ 29
4.2.6.8 Deleting Snapshots ........ 30
4.2.7 Associating a Data Container with an Object Tier ........ 31
4.3 Disaster Recovery for Data Containers ........ 33
4.3.1 Replicating Data Containers for Disaster Recovery ........ 33
4.3.2 Swapping Data Container Status between Active and Passive ........ 36
4.3.2.1 Swapping Data Container Status when the Local Site is Accessible ........ 36
4.3.2.2 Swapping Data Container Status when the Local Site is Inaccessible ........ 37
4.3.3 Creating Data Container Test Images ........ 40
4.3.4 Stopping Replication for a Data Container ........ 41
4.4 Upscaling Your System ........ 41
**5. Monitoring System Performance ........ 42**
5.1 Monitoring through the Dashboard ........ 42
5.2 Capacity ........ 43
5.3 Performance ........ 44
5.3.1 Performance Summary Pane ........ 44
5.3.2 Detailed Performance Window ........ 45
5.3.3 Filters and Grouping Pane ........ 46
5.3.4 Performance Graph Pane ........ 47
5.4 Monitoring Assets ........ 47
5.4.1 Top Performance Pane ........ 48
5.5 Monitoring Data Containers ........ 48
5.6 Monitoring Events ........ 52
5.7 Monitoring through System View ........ 52
5.7.1 Host Information ........ 53
5.8 Alerts ........ 53
5.9 Notifiers ........ 54
5.10 System Health Indicator ........ 54
**6. Managing Asynchronous Disaster Recovery ........ 55**

6.1 Enabling an Initial Replication Service for Asynchronous Replication ........ 55
6.2 Pairing Local and Remote Sites ........ 57
6.3 Managing Remote Protection Sites ........ 59
6.3.1 Viewing Remote Protection Sites ........ 59
6.3.2 Editing Local and Remote Site Pairings ........ 60
6.3.3 Removing Remote Site Pairing ........ 61
6.4 Adding Additional Replication Services ........ 62
**7. Managing the ClearTier Object Tier ........ 65**
7.1 Enabling the Object Tier ........ 65
7.1.1 Enabling the Object Tier Automatically ........ 65
7.1.2 Enabling the Object Tier Manually ........ 67
7.1.3 Creating Replication Services for an Object Tier with External Networks ........ 69
7.2 Managing Your Data Tiering Policy ........ 70
7.3 Managing Object Tiers ........ 71
7.3.1 Deleting Object Tiers ........ 71
**8. Administering Your System ........ 73**
8.1 Changing System Settings ........ 73
8.2 Performing System Operations ........ 74
8.3 Generating and Viewing Logs ........ 75
8.4 Viewing REST API Documentation ........ 75
8.5 Licensing Information and System Properties ........ 76
8.5.1 Updating the License ........ 76
8.6 Upgrading the ECFS System ........ 76
8.7 User Profile ........ 77
8.7.1 Changing Your Password ........ 77
8.7.2 Logging Out from ECFS ........ 78
**9. Troubleshooting ........ 79**
9.1 System Lockdown ........ 79
9.2 Resolving Failed Devices ........ 79
9.3 Events List ........ 80

# 1. Introduction

## 1.1 Document Scope

This document contains the explanations and procedures you require to efficiently use the ECFS (Elastifile Cloud File System) through the Management Console.

The document contains the following sections:

- Introduction
- Logging in to ECFS
- Getting to Know Your GUI
- Provisioning
- Monitoring System Performance
- Administering Your System

## 1.2 Related Documentation

Refer to the following documents for additional information:

- ECFS Installation Guide

## 1.3 Customer Support

You can contact Elastifile customer support at support@elastifile.com.

## 1.4 ECFS's Call Home Feature

For system level support and debug, ECFS provides a call-home feature that sends general system details and major events to Elastifile on a daily basis. No user data is transferred to Elastifile.

This feature provides Elastifile's customer support with a clear view of field issues and assists with resolution and root cause analysis.

The sent data includes:

- System Name
- Version
- Number of Hosts in the ECFS system
- Number of Data Containers
- Number of Shares
- Total Capacity
- Utilized Capacity
- Daily Utilization

- Hosts added to system
- Hosts removed from system
- Hosts dropping from system due to an error
- Clients' mounts
- Clients' unmount
- Data Rebuild
- Data Rebuild size
- Writes in bytes
- Reads in bytes
- Number of created files
- Number of removed files
- List of events
- License details

# 2. Logging in to ECFS

**To login to ECFS:**

1. In your browser, type the ECFS Management URL (IP address or DNS you set during installation) and press **Enter**. The login window appears.
2. Enter the password you set during installation.
3. Click **LOGIN**.

# 3. Getting to Know the Management Console Interface

## 3.1 Main Interface

The ECFS user interface is divided into the following main elements:

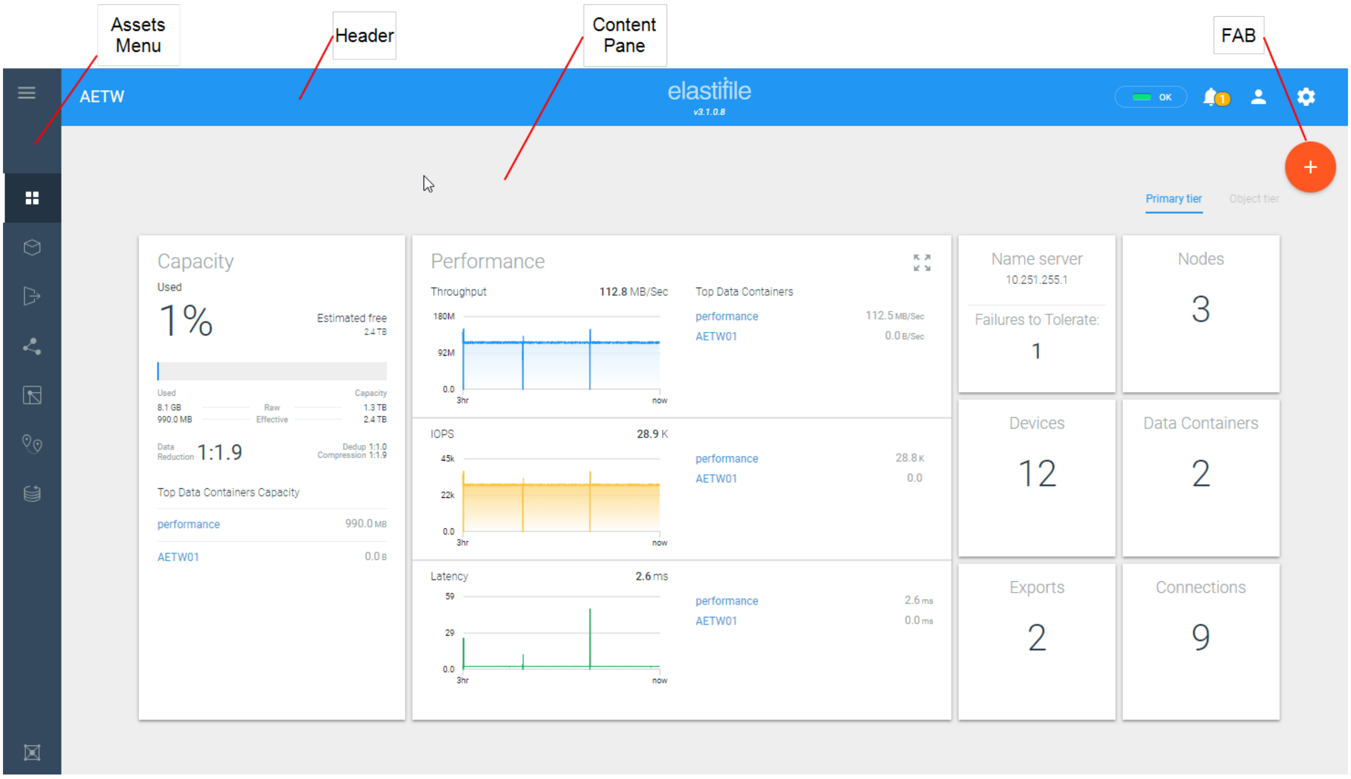

| Element | Description |
|---|---|
| **Assets Menu** | Selects the type of asset for which you want to display information in the content pane. |
| **Header** | Displays the system name, version and provides access to various system features as detailed in Section 3.2.2 - Header. |
| **Content Pane** | Serves as the main working area where you view and interact with the system. |
| **FAB** | Floating Action Button - enables you to perform different tasks according to your current location or workflow. |

## 3.2 Main Entities

### 3.2.1 Assets Menu

You can display information in the content pane for the following assets:

- (**ALL**) - displays the dashboard.
- (**DATA CONTAINERS**), (**SHARES**) , (**CLIENTS**) , (**SERVICES**), - displays the inventory for the asset you select.
- **REMOTE SITES** - opens the remote site window for managing remote sites.
- **OBJECT TIERS** - opens the data object tiering window for managing object tiers and object tiering policy.
- **SYSTEM VIEW** - displays the operational status of the elements comprising the ECFS system.

### 3.2.2 Header

The Header contains the following elements:

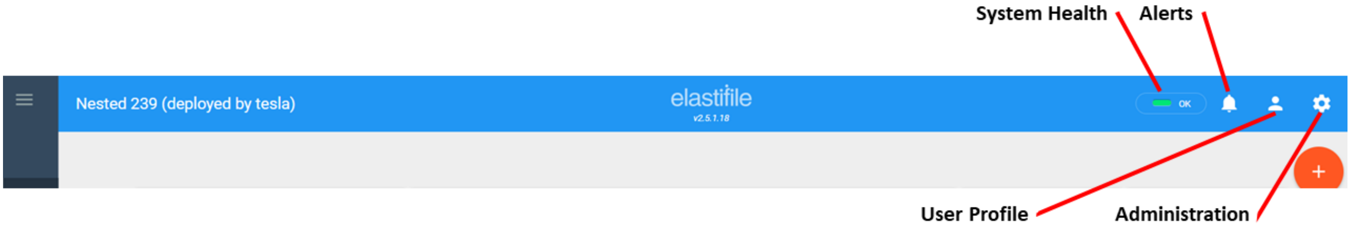

| Element | Description |
|---|---|
| System Health Indicator | Provides real-time health assessment for the ECFS system. |
| Alerts | Displays the number of critical events, and opens the critical events list. |
| User Profile | Opens the User menu items. |
| Administering Your System | Opens the Administration menu items. |

### 3.2.3 FAB

The (**FAB**) (floating action button) is a context-sensitive tool that enables you to perform different tasks according to your current task, such as add a data container or add a host.

## 3.3 Working with the Management Console Interface

When hovering over icons in the Assets menu, each icon expands to show the icon name.

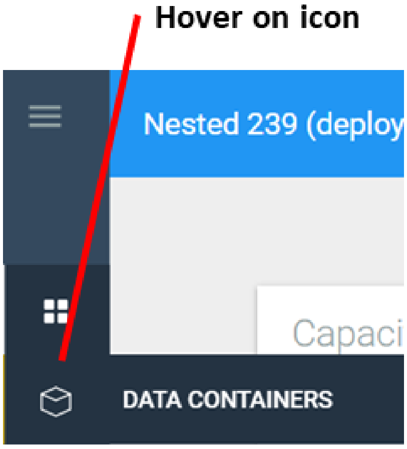

# 4. Provisioning

## 4.1 Understanding Provisioning using ECFS

Provisioning refers to the act of creating and configuring a data container for storing and retrieving data by authorized users. Data containers must contain a default root share that allows authorized users to store and retrieve data. For each share, you can define access rules for relevant client types.

You can also create local snapshots (logical backup points) of a data container and assign shares to the snapshot to allow user access.

## 4.2 Provisioning Data Containers

### 4.2.1 About Data Containers

Data containers are logical groupings of files or objects that are managed as a collection.

> The Management Console supports the creation of the following types of data containers:
>
> - public - a shared data container. Open to all clients unless deauthorized (black-listed). A typical use case is a shared file system.
> - application - a dedicated data container to be accessed by a specific list of clients/applications. Unreachable by all clients unless authorized (white-listed). A typical use case is shared storage for applications or databases.

When you create the data container, a default root share is created without client access rules and with root mapped to "Everyone". This prevents allows access for everyone to the data container. It is advised to create new shares tailored for your requirements.

### 4.2.2 Creating a Data Container

**To create a data container:**

1. Hover over + (**FAB**), then click ⬡ (**Add Data Container**).

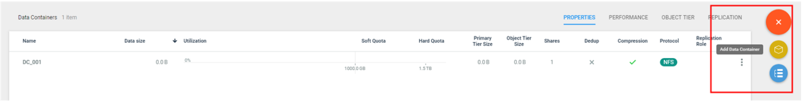

| Name | Data size | Utilization | Soft Quota | Hard Quota | Primary Tier Size | Object Tier Size | Shares | Dedup | Compression | Protocol | Replication Role |
|---|---|---|---|---|---|---|---|---|---|---|---|
| DC_001 | 0.0 B | 0% | 1000.0 GB | 1.5 TB | 0.0 B | 0.0 B | 1 | ✕ | ✓ | NFS | |

2. In the **Select data container type** window, click the type of data container you want to add, and click **Select**.

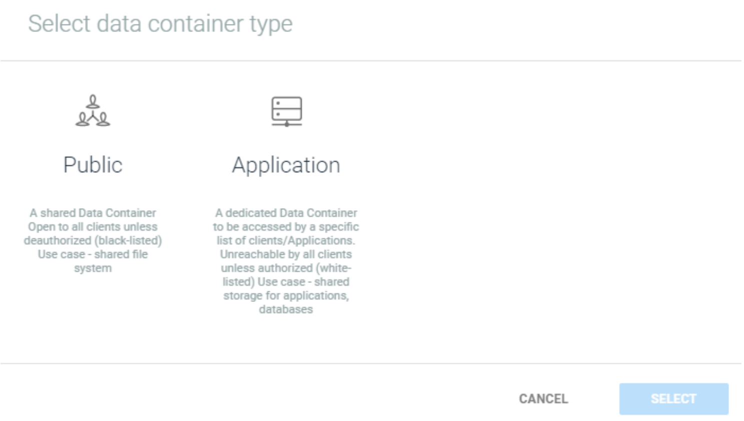

Select data container type

Public

A shared Data Container Open to all clients unless deauthorized (black-listed) Use case - shared file system

Application

A dedicated Data Container to be accessed by a specific list of clients/Applications. Unreachable by all clients unless authorized (white-listed) Use case - shared storage for applications, databases

3. The New data container window for the type of data container you selected appears as shown below.

- New Public data container:

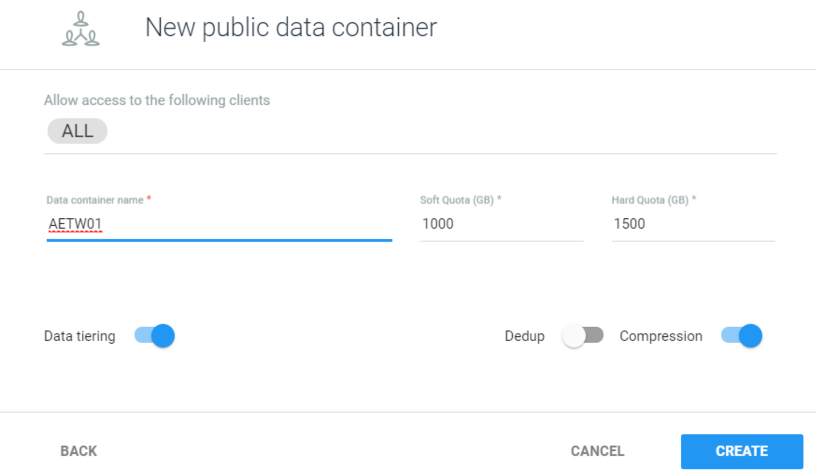

New public data container

Allow access to the following clients

ALL

Data container name * AETW01

Soft Quota (GB) * 1000

Hard Quota (GB) * 1500

Data tiering: on

Dedup: off

Compression: on

- New application data container:

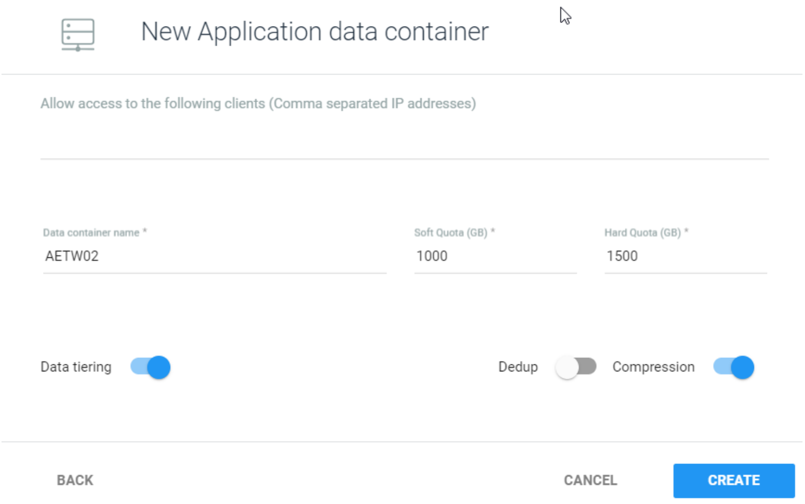

New Application data container

Allow access to the following clients (Comma separated IP addresses)

Data container name * AETW02

Soft Quota (GB) * 1000

Hard Quota (GB) * 1500

Data tiering: on

Dedup: off

Compression: on

4. Configure the fields (all fields are mandatory) as described in the following table:

| Field | Description |
|---|---|
| **Allow access to the following clients** | • Public data containers - all clients are automatically authorized.<br>• Application data containers - set the clients authorized (white-listed) for access. As you type the clients list will automatically open with the names of existing clients with the same or similar names. |
| **Data container name** | A unique name for your data container. As you type the **Used names** list will automatically open with the names of existing data containers with the same or similar name. You cannot use a name from this list. An error message will appear in the event the name you set is invalid. |
| **Soft Quota** | Storage (in GB) which, when passed, raises a notification. |
| **Hard Quota** | Storage (in GB) limit which cannot be surpassed. |
| **Data tiering** | Enable ClearTier to automatically move data to a cheaper object store while keeping the files available via the same NFS mount point (see Section 7 - Managing the ClearTier Object Tier). |
| **Dedup** and **Compression** | Select a data policy with corresponding dedup and compression settings. 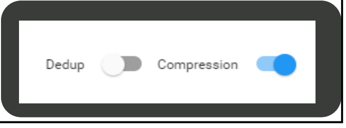 Dedup ○ Compression ● |

5. Click **Create**.

6. The data container is created. Note the mount command to use on your client.
To configure client access to the data container, click **EDIT DATA CONTAINER**.

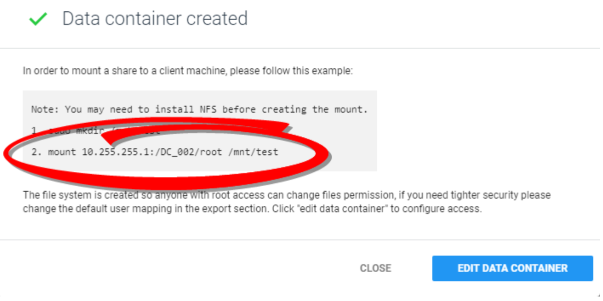

✓ Data container created

In order to mount a share to a client machine, please follow this example:

```
Note: You may need to install NFS before creating the mount.
1. sudo mkdir /mnt/test
2. mount 10.255.255.1:/DC_002/root /mnt/test
```

The file system is created so anyone with root access can change files permission, if you need tighter security please change the default user mapping in the export section. Click "edit data container" to configure access.

The default mapping for a public data container data enables full read/write access by all users as follows:

- **User Mapping**: Map Everyone to
- **User ID**: root
- **Default Access**: Read/Write

If you want to apply more restricted access, see Section 4.2.4.1 - Adding Data Container Shares.

7. Define/modify the data container features as described in the following sections:

- Section 4.2.3 - Mounting Your Storage
- Section 4.2.4 - Modifying Data Container Properties
- Section 4.2.5 - Deleting Data Containers

- Section 4.2.6 - Creating Data Container Snapshots and Schedules
- Section 4.2.7 - Associating a Data Container with an Object Tier
- Section 4.3 - Disaster Recovery for Data Containers
- Section 4.4 - Upscaling Your System
- Section 5.5 - Monitoring Data Containers

### 4.2.3 Mounting Your Storage

Data container shares are mounted as standard NFS shares. Refer to your OS guides for mounting procedures.

The mount path is defined in the following format:

[IP/hostname]:[data container name]/[share name].

**Example:**

```
10.1.1.10:my_fs0/my_share_name
```

### 4.2.4 Modifying Data Container Properties

**To modify a data container's properties:**

1. In the Assets menu, click ⬡ (**DATA CONTAINERS**).

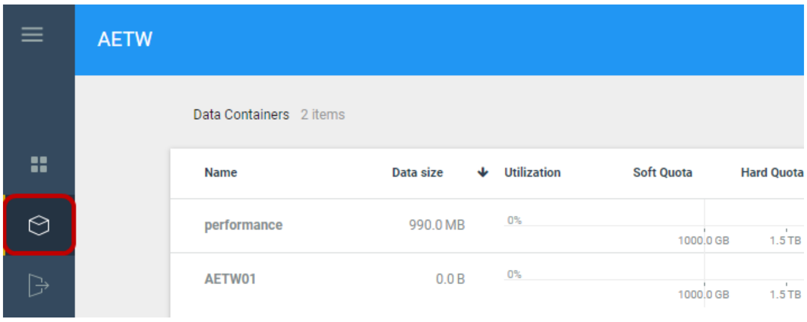

AETW

Data Containers 2 items

| Name | Data size | Utilization | Soft Quota | Hard Quota |
|---|---|---|---|---|
| performance | 990.0 MB | 0% | 1000.0 GB | 1.5 TB |
| AETW01 | 0.0 B | 0% | 1000.0 GB | 1.5 TB |

2. Click the required data container. Modify the fields at the top of the window as required, then click **Update**.

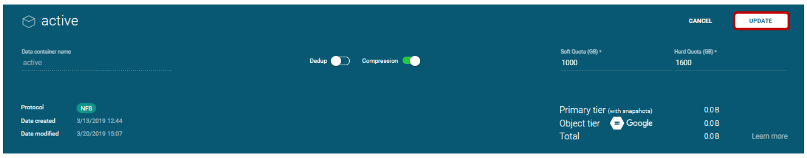

active

Data container name: active

Dedup ○ Compression ●

Soft Quota (GB) *: 1000

Hard Quota (GB) *: 1600

Protocol: NFS

Date created: 3/13/2019 12:44

Date modified: 3/20/2019 15:07

Primary tier (with snapshots) 0.0 B

Object tier Google 0.0 B

Total 0.0 B

### 4.2.4.1 Adding Data Container Shares

**To add a data container share:**

1. In the Assets menu, click (**DATA CONTAINERS**) and click the data container to which you want to add a share.

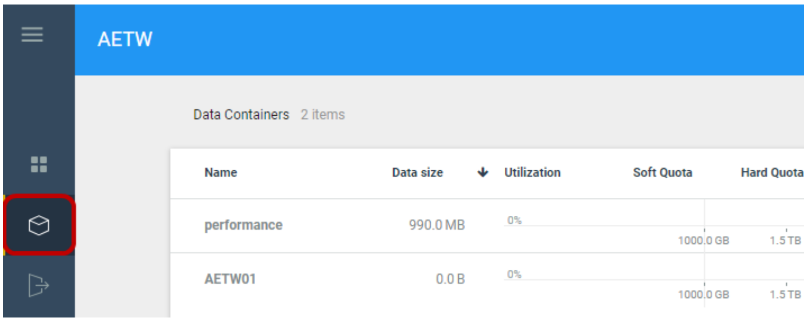

AETW

Data Containers 2 items

| Name | Data size | Utilization | Soft Quota | Hard Quota |
|---|---|---|---|---|
| performance | 990.0 MB | 0% | 1000.0 GB | 1.5 TB |
| AETW01 | 0.0 B | 0% | 1000.0 GB | 1.5 TB |

2. Hover over (**FAB**), then click **Add Share.**

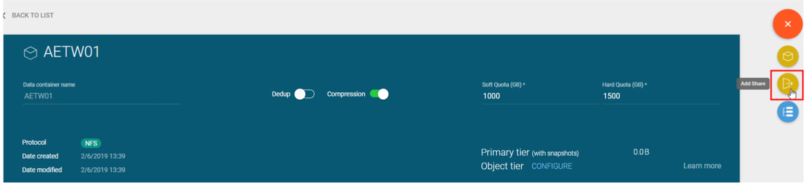

BACK TO LIST

AETW01

Data container name: AETW01

Dedup

Compression

Soft Quota (GB) *: 1000

Hard Quota (GB) *: 1500

Protocol: NFS

Date created: 2/6/2019 13:39

Date modified: 2/6/2019 13:39

Primary tier (with snapshots): 0.0B

Object tier CONFIGURE

Learn more

3. Configure the fields as described in the table below:

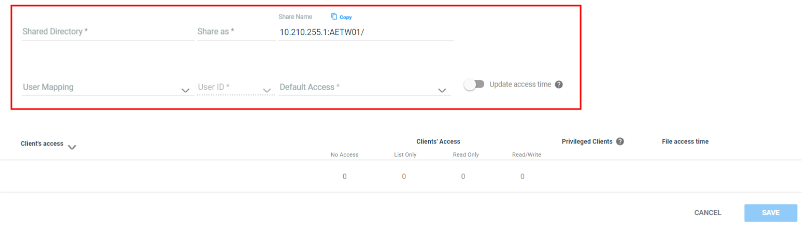

Shared Directory *

Share as *

Share Name: 10.210.255.1:AETW01/

User Mapping

User ID *

Default Access *

Update access time

| Client's access | Clients' Access |  |  |  | Privileged Clients | File access time |
|---|---|---|---|---|---|---|
|  | No Access | List Only | Read Only | Read/Write |  |  |
|  | 0 | 0 | 0 | 0 |  |  |

| Field | Description |
|---|---|
| **Shared Directory** | The path of the data container that will be shared with clients.<br>If the path does not exist, the following message appears: **New directory will be created**. In this case:<br>1. Select the check box **New directory will be created**.<br>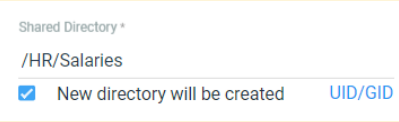<br>Shared Directory *<br>/HR/Salaries<br>☑ New directory will be created UID/GID<br>2. Click **UID/GID**. The following window appears:<br>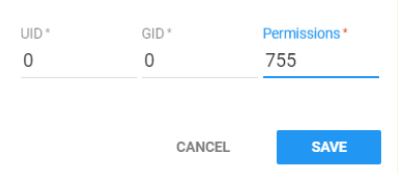<br>UID * 0<br>GID * 0<br>Permissions * 755<br>CANCEL SAVE<br>3. Set the UID, GID and Permissions for the directory that will be created.<br>4. Click **Save**. |
| **Share As** | Optional. A directory alias that will be added to the Share Name. If you don't define an alias, the directory path name is used. |
| **Share Name** | Read only. Comprised from the data container name and the directory path/alias.<br>To copy the **Share Name** to the clipboard, click Copy above the field. |
| **User Mapping** | The directory user mapping rule (squash). |
| **User ID** | Target user for **User Mapping**. |
| **Default Access** | Set the default access permissions for all clients as follows:<br>- No access<br>- List only<br>- Read only<br>- Read/Write |
| **Update time access** | For data containers using the object tier (see Section 7 - Managing the ClearTier Object Tier):<br>- On - clients that mount this data container modify the file access timestamp.<br>- Off - prevents clients from modifying the file access timestamp. |

5. Click **Client's Access**.

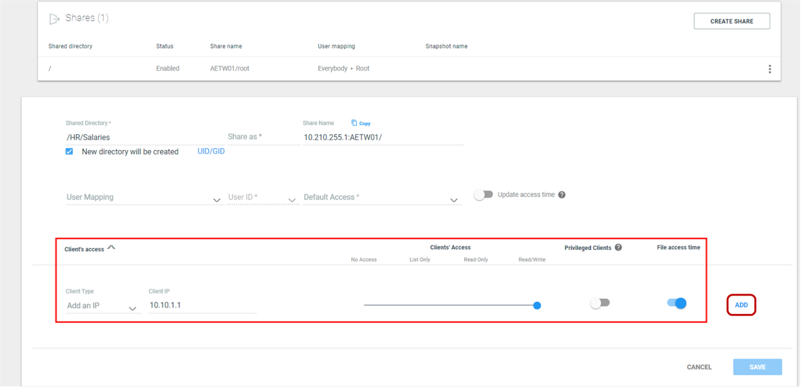

Shares (1)

CREATE SHARE

| Shared directory | Status | Share name | User mapping | Snapshot name |
|---|---|---|---|---|
| / | Enabled | AETW01/root | Everybody ▸ Root | |

Shared Directory * /HR/Salaries

☑ New directory will be created UID/GID

Share as *

Share Name Copy 10.210.255.1:AETW01/

User Mapping

User ID *

Default Access *

Update access time

Client's access

| Client Type | Client IP | Clients' Access | Privileged Clients | File access time |
|---|---|---|---|---|
| Add an IP | 10.10.1.1 | Read/Write | | |

6. In the Client IP field, type either the IP address or the Subnet details of the client.

7. Use the default access rule (above) or drag the slider to the required access rule option:

   - No Access
   - List Only
   - Read Only
   - Read/Write

8. If the data container user mapping (above) is set to **Map root to nobody**, toggle **Privileged Clients** and add at least one client as a privileged client in order to allow root access to create the file system.

9. For data containers using the object tier (see Section 7 - Managing the ClearTier Object Tier):

   - On - clients that mount this data container modify the file access timestamp.
   - Off - prevents clients from modifying the file access timestamp.

10. Click **Add**. The client access rule appears.

11. Repeat Steps 6 to 10 for each additional client for which you need to define access rules.

12. Click **Save**.

### 4.2.4.2 Editing Data Container Shares

**To edit a data container share:**

1. In the Assets menu, click ⬡ (**DATA CONTAINERS**) and click the required data container.

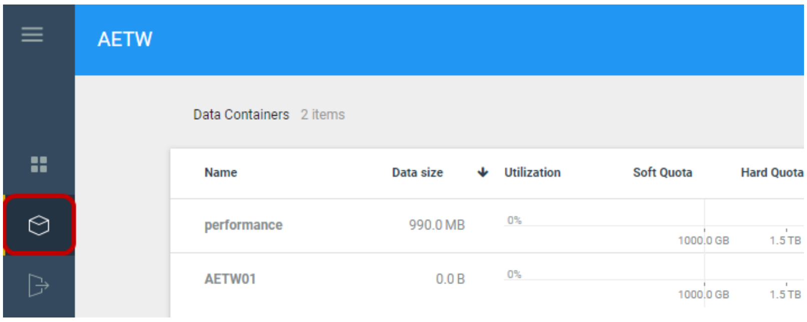

AETW

Data Containers 2 items

| Name | Data size | Utilization | Soft Quota | Hard Quota |
|---|---|---|---|---|
| performance | 990.0 MB | 0% | 1000.0 GB | 1.5 TB |
| AETW01 | 0.0 B | 0% | 1000.0 GB | 1.5 TB |

2. Click the required share.

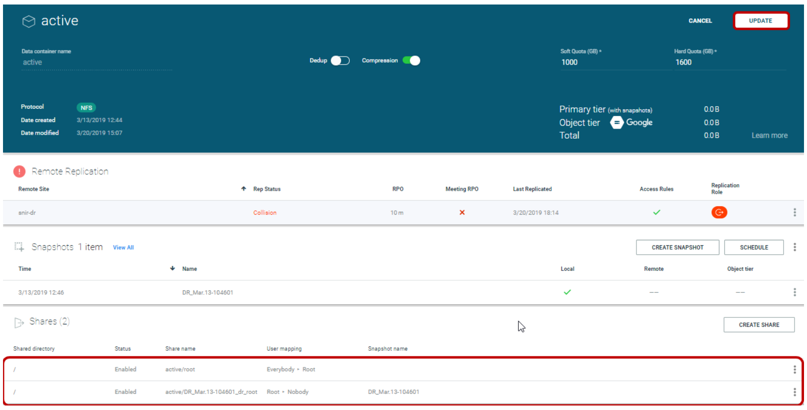

active

Data container name: active

Dedup: off; Compression: on

Soft Quota (GB) *: 1000

Hard Quota (GB) *: 1600

Protocol: NFS
Date created: 3/13/2019 12:44
Date modified: 3/20/2019 15:07

Primary tier (with snapshots): 0.0 B
Object tier Google: 0.0 B
Total: 0.0 B

Remote Replication

| Remote Site | Rep Status | RPO | Meeting RPO | Last Replicated | Access Rules | Replication Role |
|---|---|---|---|---|---|---|
| snir-dr | Collision | 10 m | ✗ | 3/20/2019 18:14 | ✓ | |

Snapshots 1 item

| Time | Name | Local | Remote | Object tier |
|---|---|---|---|---|
| 3/13/2019 12:46 | DR_Mar.13-104601 | ✓ | --- | --- |

Shares (2)

| Shared directory | Status | Share name | User mapping | Snapshot name |
|---|---|---|---|---|
| / | Enabled | active/root | Everybody ‣ Root | |
| / | Enabled | active/DR_Mar.13-104601_dr_root | Root ‣ Nobody | DR_Mar.13-104601 |

3. The share details appear.

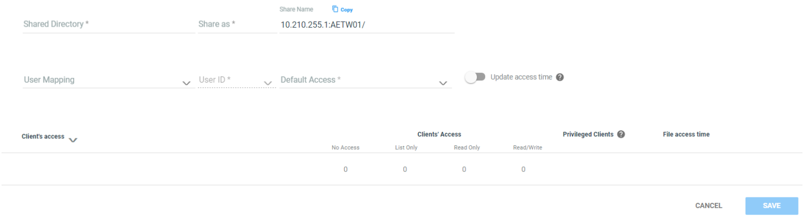

Shared Directory *

Share as *

Share Name: 10.210.255.1:AETW01/

User Mapping

User ID *

Default Access *

Update access time: off

| Client's access | Clients' Access | | | | Privileged Clients | File access time |
|---|---|---|---|---|---|---|
| | No Access | List Only | Read Only | Read/Write | | |
| | 0 | 0 | 0 | 0 | | |

4. Modify the required fields as described (see Section 4.2.4.1 - Adding Data Container Shares).

5. Click **Save**.

## 4.2.5 Deleting Data Containers

**To delete a data container:**

1. In the Assets menu, click (**DATA CONTAINERS**).

2. On the row of the data container you want to delete, click ⋮ (**Options**) .

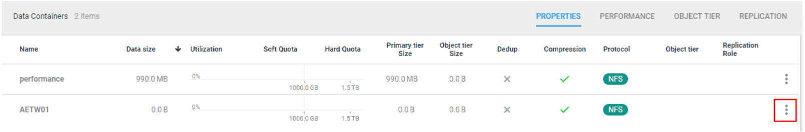

Data Containers 2 items

| Name | Data size | Utilization | Soft Quota | Hard Quota | Primary tier Size | Object tier Size | Dedup | Compression | Protocol | Object tier | Replication Role |
|---|---|---|---|---|---|---|---|---|---|---|---|
| performance | 990.0 MB | 0% | 1000.0 GB | 1.5 TB | 990.0 MB | 0.0 B | ✗ | ✓ | NFS | | |
| AETW01 | 0.0 B | 0% | 1000.0 GB | 1.5 TB | 0.0 B | 0.0 B | ✗ | ✓ | NFS | | |

3. Click **Delete**.

> A data container with shares and/or files cannot be deleted.

## 4.2.6 Creating Data Container Snapshots and Schedules

You can create local snapshots (logical backup points) of a data container and assign shares to the snapshot to allow user access. If you setup an object tier, you can move snapshots that have not been accessed for a specific period to the object tier for cold storage, or move them manually.

You can create snapshots on a regular basis through a schedule, and delete the snapshots after a defined time or move them to the object tier after a specific time.

### 4.2.6.1 Creating Data Container Snapshots

**To create a data container snapshot:**

1. In the Assets menu, click (**DATA CONTAINERS**).

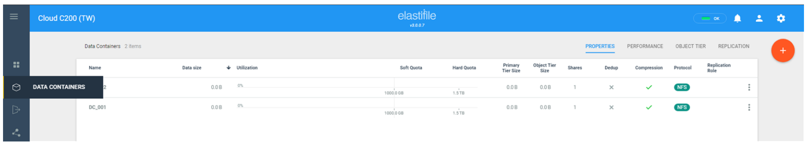

Cloud C200 (TW)

Data Containers 2 items

| Name | Data size | Utilization | Soft Quota | Hard Quota | Primary Tier Size | Object Tier Size | Shares | Dedup | Compression | Protocol | Replication Role |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 2 | 0.0 B | 0% | 1000.0 GB | 1.5 TB | 0.0 B | 0.0 B | 1 | ✕ | ✓ | NFS | |
| DC_001 | 0.0 B | 0% | 1000.0 GB | 1.5 TB | 0.0 B | 0.0 B | 1 | ✕ | ✓ | NFS | |

2. Click the required data container.

3. Click **Create Snapshot**.

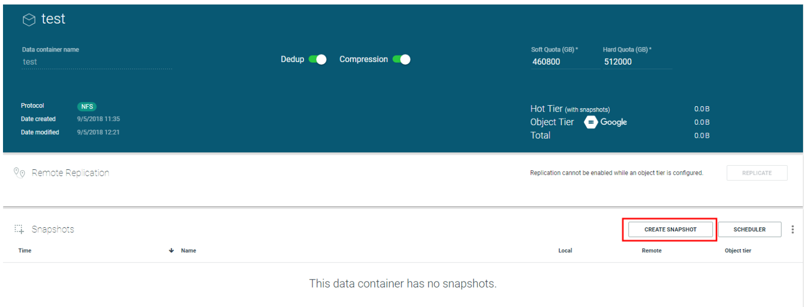

test

Data container name: test

Dedup: on
Compression: on

Soft Quota (GB) *: 460800
Hard Quota (GB) *: 512000

Protocol: NFS
Date created: 9/5/2018 11:35
Date modified: 9/5/2018 12:21

Hot Tier (with snapshots): 0.0 B
Object Tier Google: 0.0 B
Total: 0.0 B

Remote Replication

Replication cannot be enabled while an object tier is configured.

| Time | Name | Local | Remote | Object tier |
|---|---|---|---|---|

This data container has no snapshots.

4. Complete the new snapshot details as follows:

- If the object tier is enabled (see Section 7 - Managing the ClearTier Object Tier):

  a. In **Snapshot name**, type a name for the snapshot.

  b. In the **Move to object tier** list, select one of the following:

  - **Never**
  - **Hours** and add a value.
  - **Days** and add a value.

  c. If you select hours or days in **Move to object tier**, then in the text area, type the number of hours or days.

  d. In the **Delete** area, in the list, select one of the following:

  - **Never**
  - **Hours** and add a value.
  - **Days** and add a value.

  e. If you select hours or days in **Delete**, then in the text area, type the number of hours or days.

  f. Click **CREATE**.

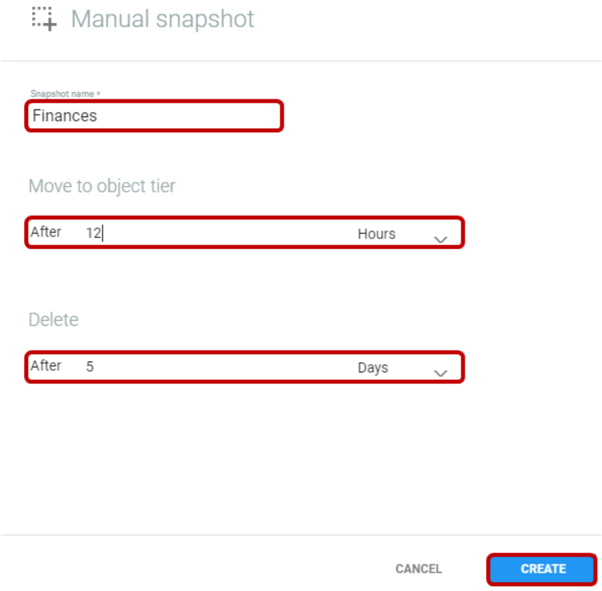

Manual snapshot

Snapshot name *: Finances

Move to object tier: After 12 Hours

Delete: After 5 Days

- If the object tier is not configured (see Section 7 - Managing the ClearTier Object Tier):

  a. In **Snapshot name**, type a name for the snapshot.

  b. In the **Delete** area, in the list, select one of the following:

     - **Never**
     - **Hours** and add a value.
     - **Days** and add a value.

  c. Click **CREATE**.

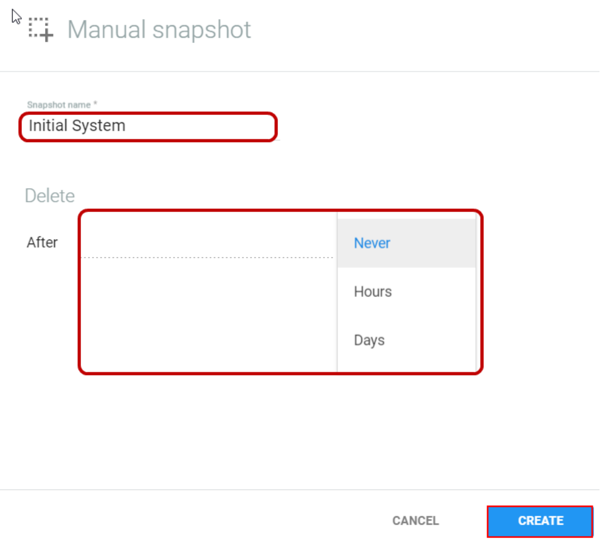

Manual snapshot

Snapshot name *: Initial System

Delete

After: Never / Hours / Days

CANCEL | CREATE

5. The snapshot you created appears in the data container window.

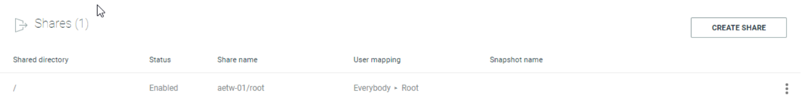

Shares (1) — CREATE SHARE

| Shared directory | Status | Share name | User mapping | Snapshot name |
|---|---|---|---|---|
| / | Enabled | aetw-01/root | Everybody ‣ Root | |

### 4.2.6.2 Moving a Snapshot to the Object Tier

You can manually move snapshots to the object tier for cold storage.

**To move a data container snapshot to the object tier:**

1. In the Assets menu, click (**DATA CONTAINERS**).

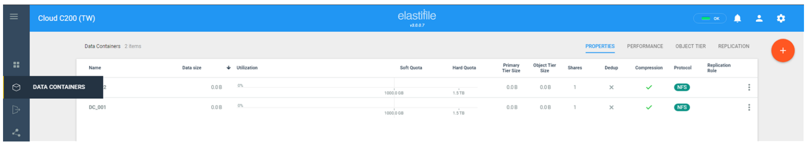

Cloud C200 (TW) — elastifile v3.0.0.7 — OK

DATA CONTAINERS

Data Containers 2 items — PROPERTIES | PERFORMANCE | OBJECT TIER | REPLICATION

| Name | Data size | Utilization | Soft Quota | Hard Quota | Primary Tier Size | Object Tier Size | Shares | Dedup | Compression | Protocol | Replication Role |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 2 | 0.0 B | 0% | 1000.0 GB | 1.5 TB | 0.0 B | 0.0 B | 1 | ✕ | ✓ | NFS | |
| DC_001 | 0.0 B | 0% | 1000.0 GB | 1.5 TB | 0.0 B | 0.0 B | 1 | ✕ | ✓ | NFS | |

2. Click the required data container.

3. On the row of the snapshot you want to move to the object tier, click ⋮ (**Options**).

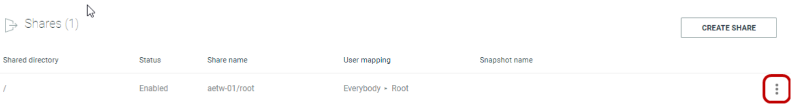

Shares (1) — CREATE SHARE

| Shared directory | Status | Share name | User mapping | Snapshot name |
|---|---|---|---|---|
| / | Enabled | aetw-01/root | Everybody ‣ Root | |

4. Click **Move to object tier**.

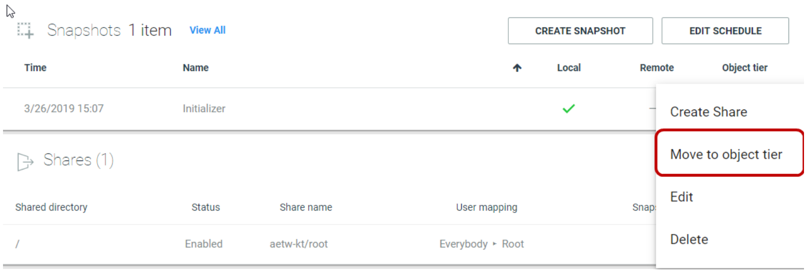

Snapshots 1 item — View All — CREATE SNAPSHOT | EDIT SCHEDULE

| Time | Name | Local | Remote | Object tier |
|---|---|---|---|---|
| 3/26/2019 15:07 | Initializer | ✓ | – | |

Shares (1)

| Shared directory | Status | Share name | User mapping | Snap |
|---|---|---|---|---|
| / | Enabled | aetw-kt/root | Everybody ‣ Root | |

If the object tier is not set up, click set it up and proceed as described in Section 7.1 - Enabling the Object Tier.

5. Click **MOVE TO OBJECT TIER**.

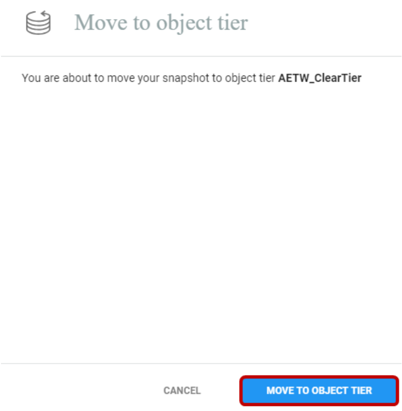

Move to object tier

You are about to move your snapshot to object tier **AETW_ClearTier**

After a while, the snapshot location changes from **Local** to **Object tier**.

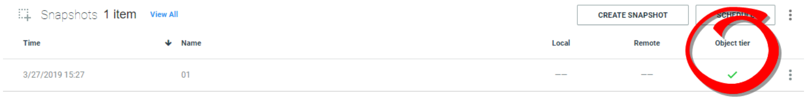

Snapshots 1 item View All

| Time | Name | Local | Remote | Object tier |
|---|---|---|---|---|
| 3/27/2019 15:27 | 01 | — | — | ✓ |

## 4.2.6.3 Creating a Data Container Snapshot Schedule

**To create a data container snapshot schedule:**

1. In the Assets menu, click (**DATA CONTAINERS**).

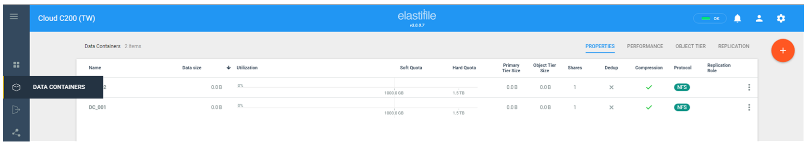

Data Containers 2 items

| Name | Data size | Utilization | Soft Quota | Hard Quota | Primary Tier Size | Object Tier Size | Shares | Dedup | Compression | Protocol | Replication Role |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 2 | 0.0 B | 0% | 1000.0 GB | 1.5 TB | 0.0 B | 0.0 B | 1 | ✕ | ✓ | NFS | |
| DC_001 | 0.0 B | 0% | 1000.0 GB | 1.5 TB | 0.0 B | 0.0 B | 1 | ✕ | ✓ | NFS | |

2. Click the required data container.

3. Click **EDIT SCHEDULE**.

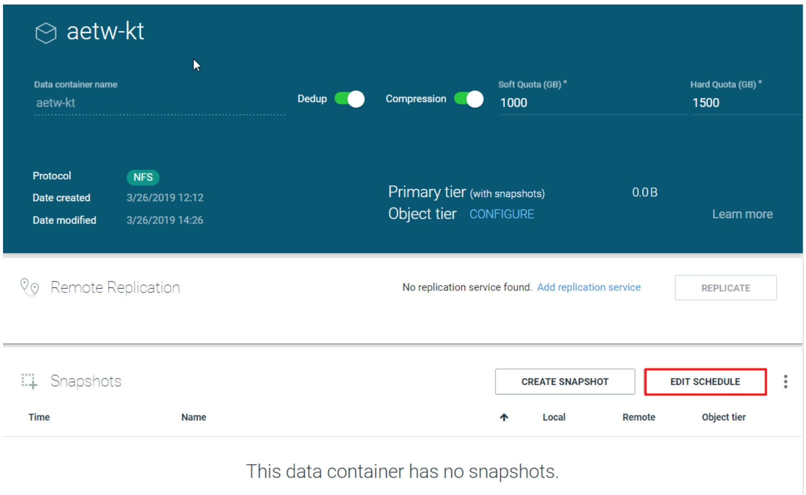

aetw-kt

Data container name: aetw-kt
Dedup
Compression
Soft Quota (GB) *: 1000
Hard Quota (GB) *: 1500

Protocol: NFS
Date created: 3/26/2019 12:12
Date modified: 3/26/2019 14:26

Primary tier (with snapshots) 0.0 B
Object tier CONFIGURE

Remote Replication

No replication service found. Add replication service

Snapshots

| Time | Name | Local | Remote | Object tier |
|---|---|---|---|---|

This data container has no snapshots.

4. Complete the new schedule details as follows:

- If the object tier is configured:

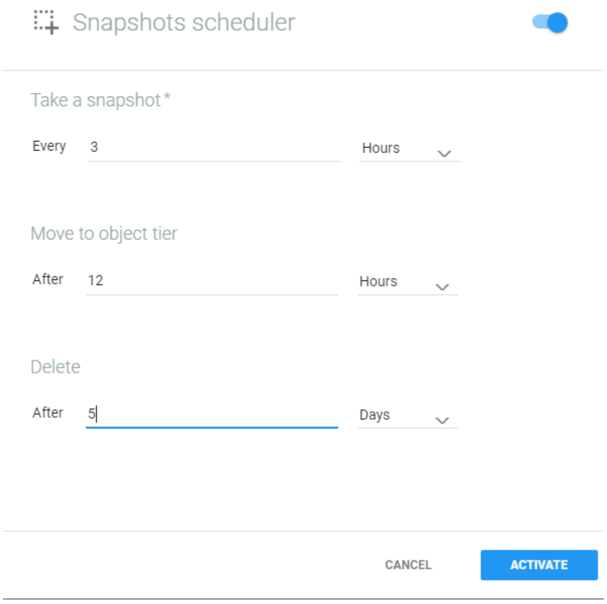

Snapshots scheduler

Take a snapshot*

Every 3 Hours

Move to object tier

After 12 Hours

Delete

After 5 Days

a. In the **Take a snapshot** list, select one of the following:

  - **Never**
  - **Hours**
  - **Days**

b. If you select hours or days in **Take a snapshot**, then in the text area, type the number of hours or days.

c. In **Move to object tier** list, select one of the following:

  - **Never**
  - **Hours**
  - **Days**

d. If you select hours or days in **Move to object tier**, then in the text area, type the number of hours or days.

e. In the **Delete** list, select one of the following:

  - **Never**
  - **Hours**
  - **Days**

f. If you select hours or days in **Delete**, then in the text area, type the number of hours or days.

g. Click **ACTIVATE**.

- If the object tier is not configured:

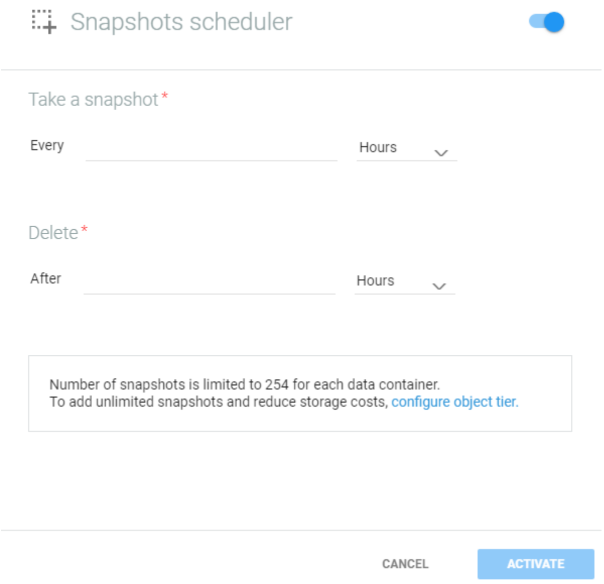

Snapshots scheduler

Take a snapshot*

Every Hours

Delete*

After Hours

Number of snapshots is limited to 254 for each data container.
To add unlimited snapshots and reduce storage costs, configure object tier.

a. In the **Take a snapshot** list, select one of the following:

  - **Never**
  - **Hours**
  - **Days**

b. If you select hours or days in **Take a snapshot**, then in the text area, type the number of hours or days.

c. In the **Delete** list, select one of the following:

  - **Never**
  - **Hours**
  - **Days**

d. If you select hours or days in **Delete**, then in the text area, type the number of hours or days.

e. Click **ACTIVATE**.

> Use ● to enable/disable the schedule.

### 4.2.6.4 Enabling and Disabling the Snapshot Schedule

You can enable and disable the snapshot schedule for a data container.

> Disabling the schedule does not delete it.

**To enable or disable the snapshot schedule:**

1. In the Assets menu, click (**DATA CONTAINERS**).

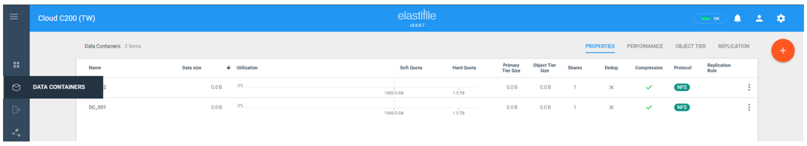

| Name | Data size | Utilization | Soft Quota | Hard Quota | Primary Tier Size | Object Tier Size | Shares | Dedup | Compression | Protocol | Replication Role |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 2 | 0.0 B | 0% | 1000.0 GB | 1.5 TB | 0.0 B | 0.0 B | 1 | ✕ | ✓ | NFS | |
| DC_001 | 0.0 B | 0% | 1000.0 GB | 1.5 TB | 0.0 B | 0.0 B | 1 | ✕ | ✓ | NFS | |

2. Click the required data container.

3. Use either of the following methods:

   - Snapshots pane:

     a. In the Snapshots pane, click ⋮ (**Options**).

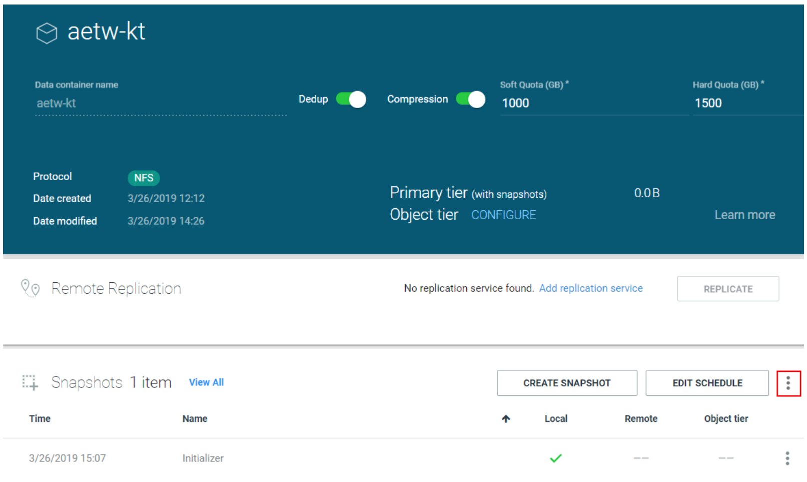

Data container name: aetw-kt
Dedup ● Compression ●
Soft Quota (GB) *: 1000
Hard Quota (GB) *: 1500
Protocol: NFS
Date created: 3/26/2019 12:12
Date modified: 3/26/2019 14:26
Primary tier (with snapshots) 0.0B
Object tier CONFIGURE
No replication service found. Add replication service

| Time | Name | Local | Remote | Object tier |
|---|---|---|---|---|
| 3/26/2019 15:07 | Initializer | ✓ | –– | –– |

     b. Click **Enable scheduler** or **Disable scheduler**.

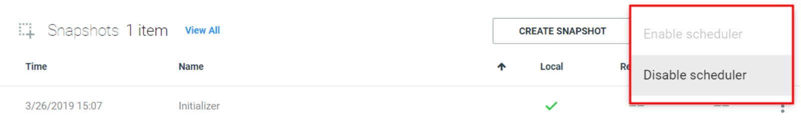

| Time | Name | Local |
|---|---|---|
| 3/26/2019 15:07 | Initializer | ✓ |

- Schedule window

a. Click **Edit Schedule**.

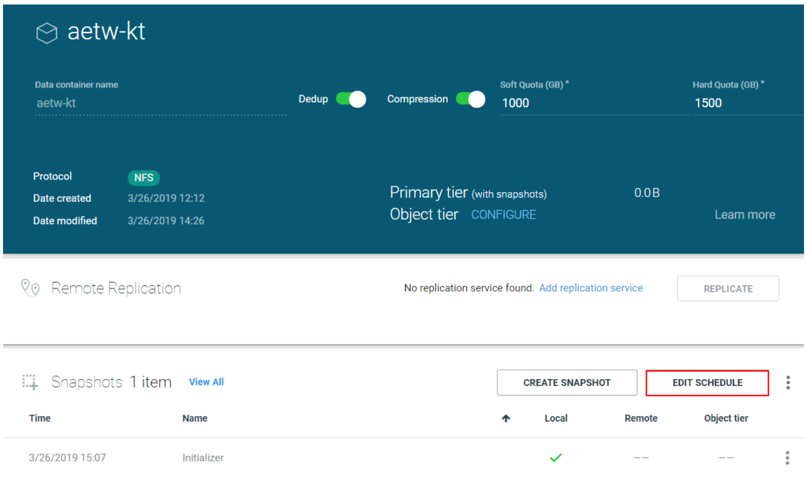

aetw-kt

Data container name: aetw-kt

Dedup ●

Compression ●

Soft Quota (GB) *: 1000

Hard Quota (GB) *: 1500

Protocol: NFS

Date created: 3/26/2019 12:12

Date modified: 3/26/2019 14:26

Primary tier (with snapshots) 0.0 B

Object tier CONFIGURE

Remote Replication

No replication service found. Add replication service

Snapshots 1 item View All

| Time | Name | Local | Remote | Object tier |
|---|---|---|---|---|
| 3/26/2019 15:07 | Initializer | ✓ | -- | -- |

b. Click ● to toggle (enable or disable) the schedule.

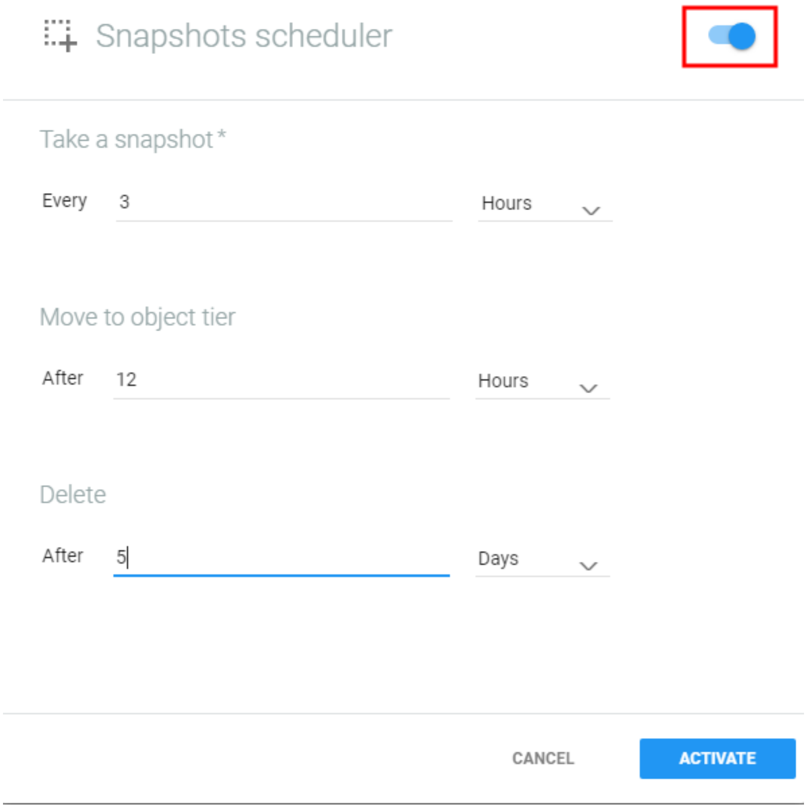

Snapshots scheduler

Take a snapshot *

Every 3 Hours

Move to object tier

After 12 Hours

Delete

After 5 Days

c. Click **Update**.

### 4.2.6.5 Adding Snapshot Shares

**To add a snapshot share:**

1. In the Assets menu, click ⬡ (**DATA CONTAINERS**).

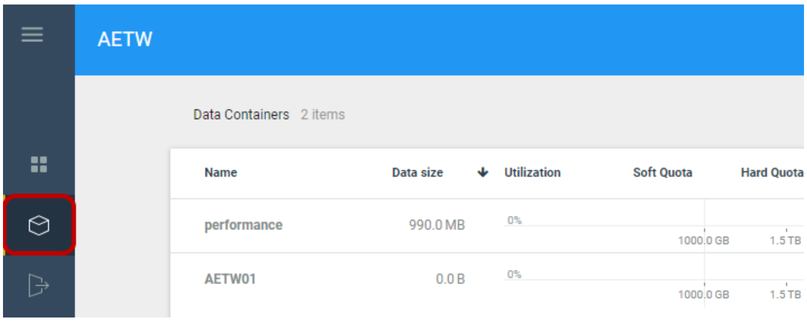

| Name | Data size | Utilization | Soft Quota | Hard Quota |
|---|---|---|---|---|
| performance | 990.0 MB | 0% | 1000.0 GB | 1.5 TB |
| AETW01 | 0.0 B | 0% | 1000.0 GB | 1.5 TB |

2. Click the required data container.

3. On the row of the snapshot for which you want to add a share, click ⋮ (**Options**).

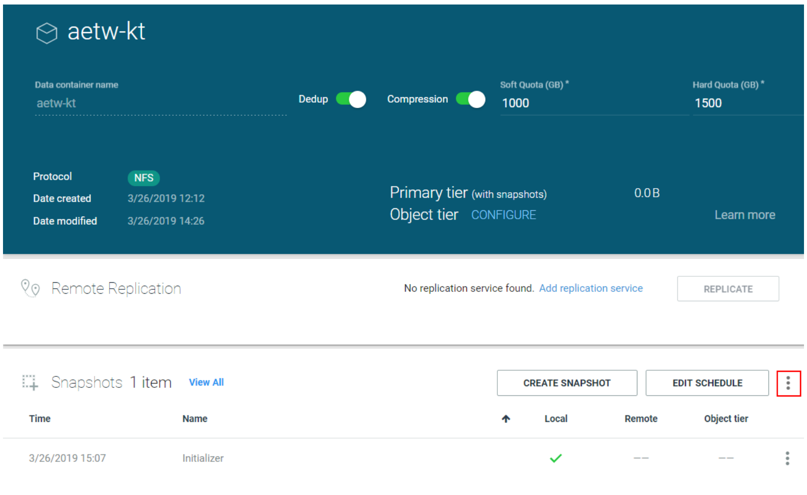

| Time | Name | Local | Remote | Object tier |
|---|---|---|---|---|
| 3/26/2019 15:07 | Initializer | ✓ | -- | -- |

4. Click **Create Share**.

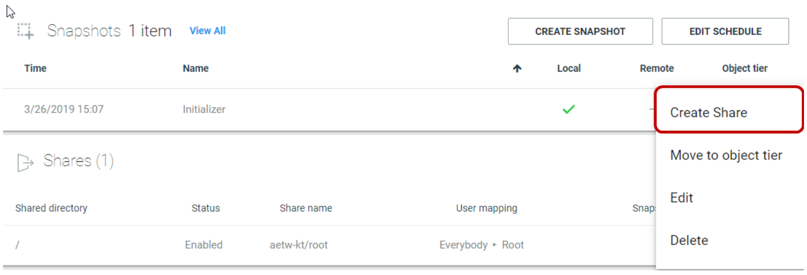

| Time | Name | Local | Remote | Object tier |
|---|---|---|---|---|
| 3/26/2019 15:07 | Initializer | ✓ | | |

| Shared directory | Status | Share name | User mapping |
|---|---|---|---|
| / | Enabled | aetw-kt/root | Everybody ▸ Root |

5. Configure the fields as described in the table below:

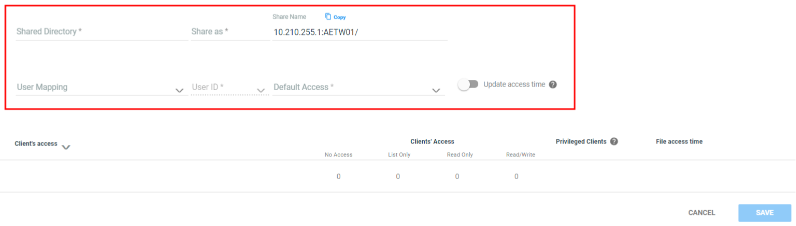

Share Name: 10.210.255.1:AETW01/

| Client's access | No Access | List Only | Read Only | Read/Write | Privileged Clients | File access time |
|---|---|---|---|---|---|---|
| | 0 | 0 | 0 | 0 | | |

| Field | Description |
|---|---|
| **Shared Directory** | The path of the snapshot that will be shared with clients.<br>If the path does not exist, the following message appears: **New directory will be created**. In this case:<br>1. Select the check box **New directory will be created**.<br>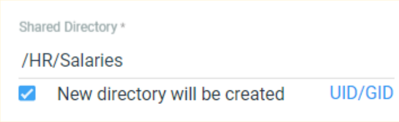<br>Shared Directory *<br>/HR/Salaries<br>☑ New directory will be created<br> <br>2. Click **UID/GID**. The following window appears:<br>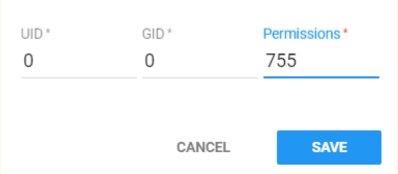<br>UID * 0<br>GID * 0<br>Permissions * 755<br> <br>3. Set the UID, GID and Permissions for the directory that will be created.<br>4. Click **Save**. |
| **Share As** | Optional. A directory alias that will be added to the Share Name. If you don't define an alias, the directory path name is used. |
| **Share Name** | Read only. Comprised from the snapshot name and the directory path/alias.<br>To copy the **Share Name** to the clipboard, click Copy above the field. |
| **User Mapping** | The directory user mapping rule (squash). |
| **User ID** | Target user for **User Mapping**. |
| **Default Access** | Set the default access permissions for all clients as follows:<br>- No access<br>- List only<br>- Read only |
| **Update time access** | For snapshots using the object tier (see Section 7 - Managing the ClearTier Object Tier):<br>- On - clients that mount this snapshot modify the file access timestamp.<br>- Off - prevents clients from modifying the file access timestamp. |

5. Click **Client's Access**.

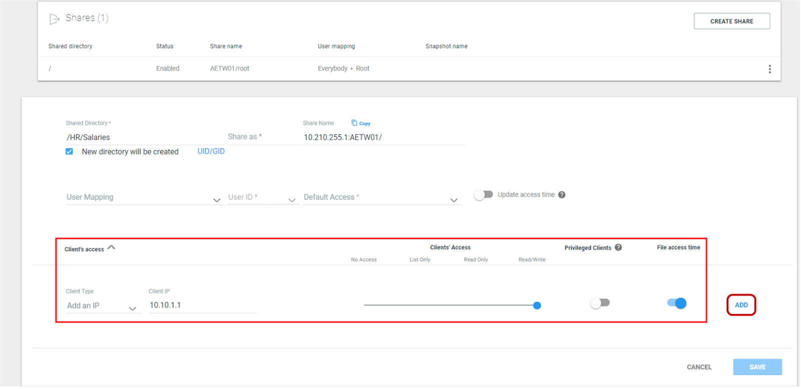

Shares (1)

| Shared directory | Status | Share name | User mapping | Snapshot name |
|---|---|---|---|---|
| / | Enabled | AETW01/root | Everybody ▸ Root | |

Shared Directory *: /HR/Salaries

Share as *

Share Name: 10.210.255.1:AETW01/

☑ New directory will be created

Client Type: Add an IP

Client IP: 10.10.1.1

6. In the Client IP field, type either the IP address or the Subnet details of the client.

7. Use the default access rule (above) or drag the slider to the required access rule option:
   - No Access
   - List Only
   - Read Only

8. If the snapshot user mapping (above) is set to **Map root to nobody**, toggle **Privileged Clients** and add at least one client as a privileged client in order to allow root access to create the file system.

9. For snapshots using the object tier (see Section 7 - Managing the ClearTier Object Tier):
   - On - clients that mount this snapshot modify the file access timestamp.
   - Off - prevents clients from modifying the file access timestamp.

10. Click **Add**. The client access rule appears.

11. Repeat Steps 6 to 10 for each additional client for which you need to define access rules.

12. Click **Save**.

### 4.2.6.6 Editing Snapshot Shares

**To edit a snapshot share:**

1. In the Assets menu, click ⬡

   (**DATA CONTAINERS**) and click the required data container.

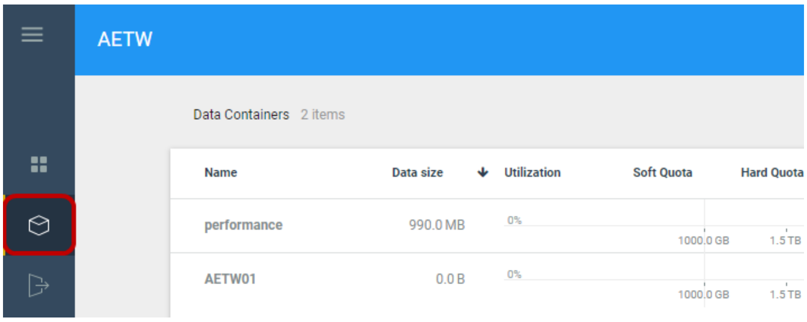

Data Containers 2 items

| Name | Data size | Utilization | Soft Quota | Hard Quota |
|---|---|---|---|---|
| performance | 990.0 MB | 0% | 1000.0 GB | 1.5 TB |
| AETW01 | 0.0 B | 0% | 1000.0 GB | 1.5 TB |

2. Click the required snapshot share.

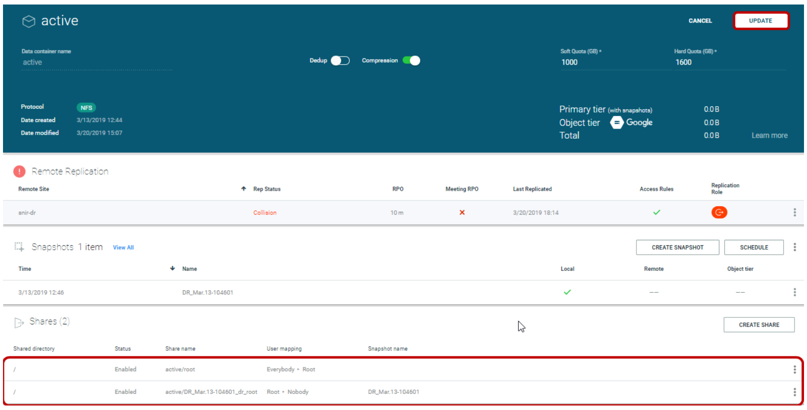

active

Data container name: active

Dedup | Compression

Soft Quota (GB) *: 1000

Hard Quota (GB) *: 1600

Protocol: NFS

Date created: 3/13/2019 12:44

Date modified: 3/20/2019 15:07

Primary tier (with snapshots): 0.0 B

Object tier Google: 0.0 B

Total: 0.0 B

Remote Replication

| Remote Site | Rep Status | RPO | Meeting RPO | Last Replicated | Access Rules | Replication Role |
|---|---|---|---|---|---|---|
| snir-dr | Collision | 10 m | ✗ | 3/20/2019 18:14 | ✓ | |

Snapshots 1 item

| Time | Name | Local | Remote | Object tier |
|---|---|---|---|---|
| 3/13/2019 12:46 | DR_Mar.13-104601 | ✓ | -- | -- |

Shares (2)

| Shared directory | Status | Share name | User mapping | Snapshot name |
|---|---|---|---|---|
| / | Enabled | active/root | Everybody ▸ Root | |
| / | Enabled | active/DR_Mar.13-104601_dr_root | Root ▸ Nobody | DR_Mar.13-104601 |

3. The share details appear.

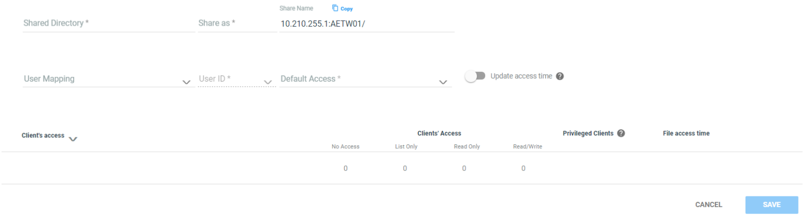

Shared Directory *

Share as *

Share Name: 10.210.255.1:AETW01/

User Mapping

User ID *

Default Access *

Update access time

Client's access

| | Clients' Access | | | | Privileged Clients | File access time |
|---|---|---|---|---|---|---|
| | No Access | List Only | Read Only | Read/Write | | |
| | 0 | 0 | 0 | 0 | | |

4. Modify the required fields as described (see Section 4.2.6.5 - Adding Snapshot Shares).

5. Click **Save**.

### 4.2.6.7 Deleting Snapshot Shares

**To delete a snapshot share:**

1. In the Assets menu, click ⬡ (**DATA CONTAINERS**).

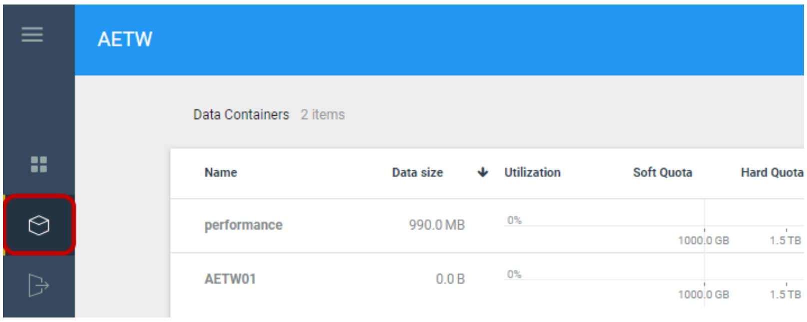

AETW

Data Containers 2 items

| Name | Data size | Utilization | Soft Quota | Hard Quota |
|---|---|---|---|---|
| performance | 990.0 MB | 0% | 1000.0 GB | 1.5 TB |
| AETW01 | 0.0 B | 0% | 1000.0 GB | 1.5 TB |

2. Click the required data container.

3. On the row of the snapshot share you want to delete, click ⋮ (**Options**) and click **Delete**.

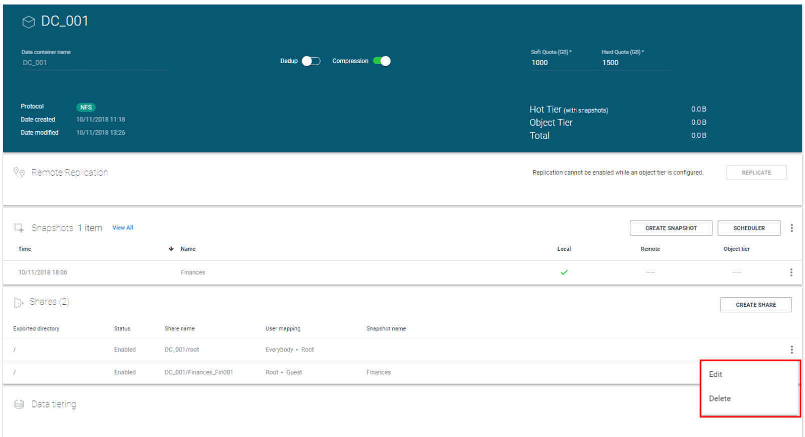

DC_001

Data container name: DC_001

Soft Quota (GB) *: 1000 | Hard Quota (GB) *: 1500

Protocol: NFS
Date created: 10/11/2018 11:18
Date modified: 10/11/2018 13:26

Hot Tier (with snapshots) 0.0 B
Object Tier 0.0 B
Total 0.0 B

Replication cannot be enabled while an object tier is configured.

| Time | Name | Local | Remote | Object tier |
|---|---|---|---|---|
| 10/11/2018 18:06 | Finances | ✓ | — | — |

| Exported directory | Status | Share name | User mapping | Snapshot name |
|---|---|---|---|---|
| / | Enabled | DC_001/root | Everybody ▸ Root | |
| / | Enabled | DC_001/Finances_Fin001 | Root ▸ Guest | Finances |

4. Click **YES, DELETE**.

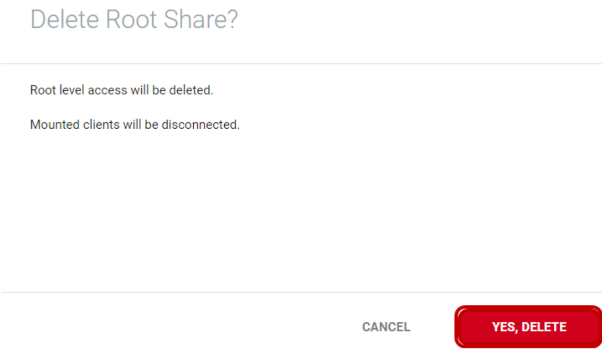

Delete Root Share?

Root level access will be deleted.

Mounted clients will be disconnected.

## 4.2.6.8 Deleting Snapshots

**To delete a snapshot:**

1. In the Assets menu, click (**DATA CONTAINERS**).

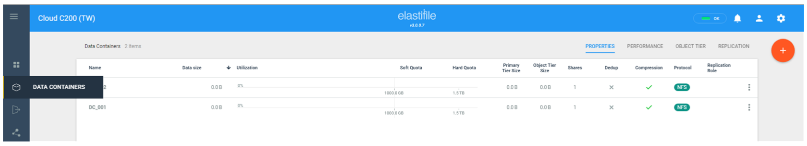

Cloud C200 (TW)

Data Containers 2 items

| Name | Data size | Utilization | Soft Quota | Hard Quota | Primary Tier Size | Object Tier Size | Shares | Dedup | Compression | Protocol | Replication Role |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 2 | 0.0 B | 0% | 1000.0 GB | 1.5 TB | 0.0 B | 0.0 B | 1 | ✕ | ✓ | NFS | |
| DC_001 | 0.0 B | 0% | 1000.0 GB | 1.5 TB | 0.0 B | 0.0 B | 1 | ✕ | ✓ | NFS | |

2. Click the required data container.

3. On the row of the snapshot you want to delete, click ⋮ (**Options**) and click **Delete**.

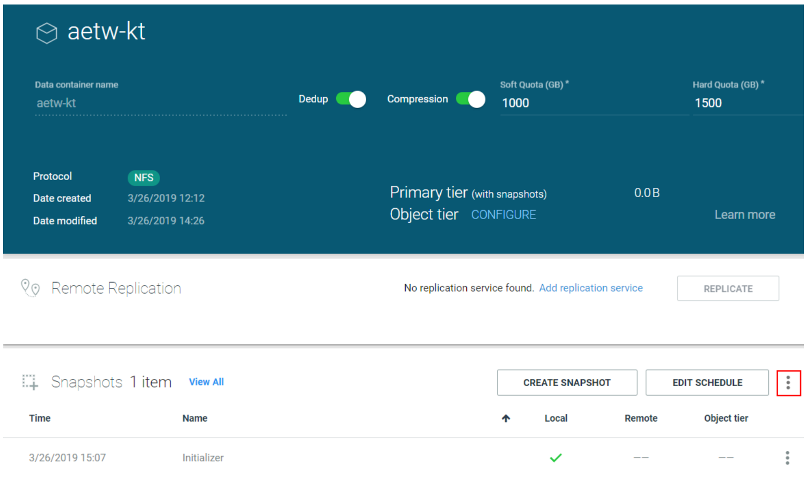

aetw-kt

Data container name: aetw-kt

Soft Quota (GB) *: 1000 | Hard Quota (GB) *: 1500

Protocol: NFS
Date created: 3/26/2019 12:12
Date modified: 3/26/2019 14:26

Primary tier (with snapshots) 0.0 B
Object tier CONFIGURE

No replication service found. Add replication service

| Time | Name | Local | Remote | Object tier |
|---|---|---|---|---|
| 3/26/2019 15:07 | Initializer | ✓ | — | — |

4. Click **YES, DELETE**.

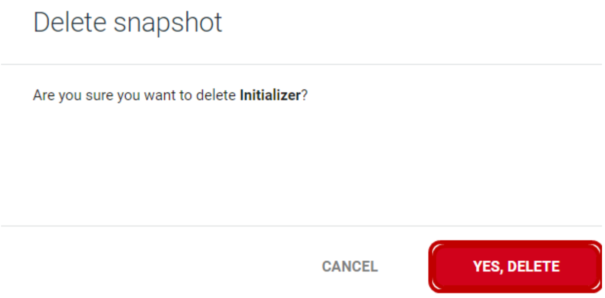

Delete snapshot

Are you sure you want to delete **Initializer**?

## 4.2.7 Associating a Data Container with an Object Tier

You can configure an object tier for data containers, and then set object tier policies for Data Container Snapshots and Snapshot Schedules (see Section 4.2.6 - Creating Data Container Snapshots and Schedules)

**To configure an object tier from a Data Container:**

1. In the Assets menu, click ⬡ (**DATA CONTAINERS**).

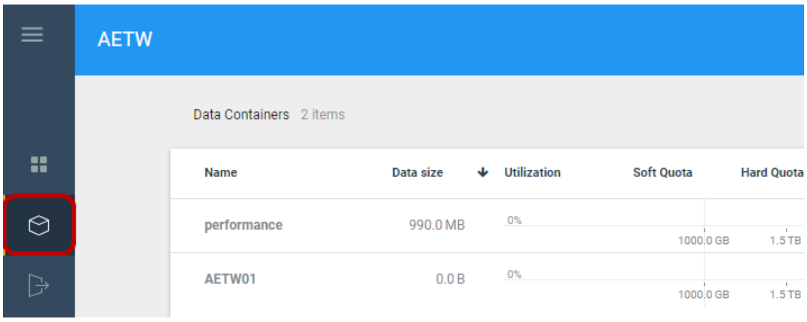

AETW

Data Containers 2 items

| Name | Data size | Utilization | Soft Quota | Hard Quota |
|---|---|---|---|---|
| performance | 990.0 MB | 0% | 1000.0 GB | 1.5 TB |
| AETW01 | 0.0 B | 0% | 1000.0 GB | 1.5 TB |

2. Click the required data container.

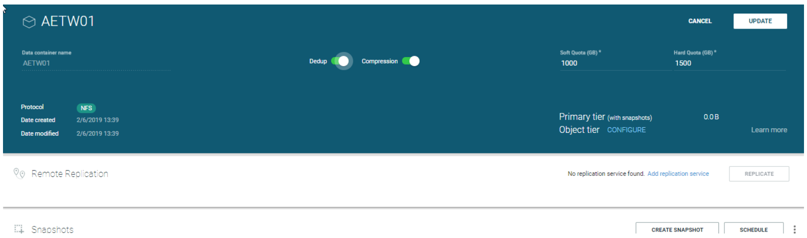

AETW01

Data container name: AETW01

Dedup: on

Compression: on

Soft Quota (GB) *: 1000

Hard Quota (GB) *: 1500

Protocol: NFS

Date created: 2/6/2019 13:39

Date modified: 2/6/2019 13:39

Primary tier (with snapshots): 0.0 B

Object tier: CONFIGURE

Remote Replication

No replication service found. Add replication service

Snapshots

3. Click **CONFIGURE**.

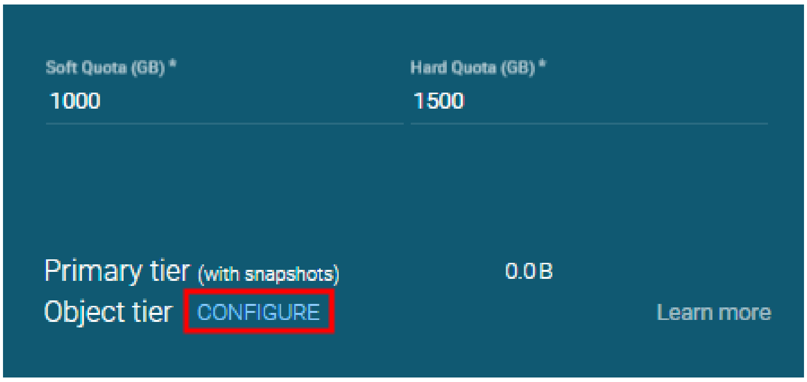

Soft Quota (GB) *: 1000

Hard Quota (GB) *: 1500

Primary tier (with snapshots): 0.0 B

Object tier: CONFIGURE

If the object tier is not set up, click set it up and proceed as described in Section 7.1 - Enabling the Object Tier.

4. In the **Pair object tier** window, **Object tier** list, select your required object tier.

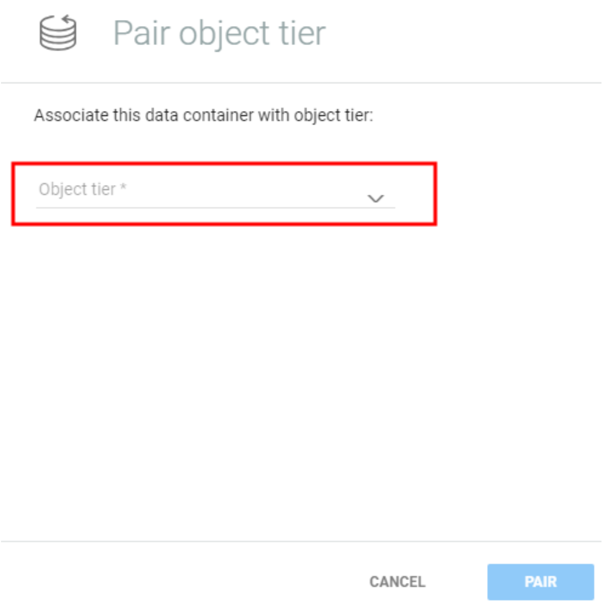
Pair object tier

Associate this data container with object tier:

Object tier *

If the object tier is not setup, then in the **Pair object tier** window, click **Create object tier** and follow the procedure in Section 7.1 - Enabling the Object Tier

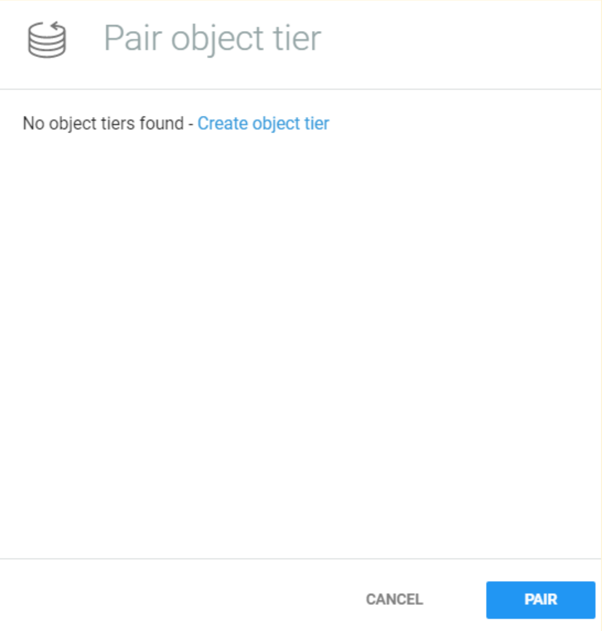
Pair object tier

No object tiers found - Create object tier

5. Click **PAIR**. The object tier is paired.

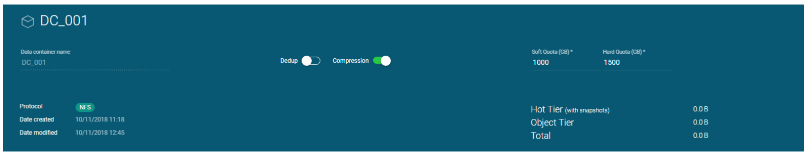
DC_001

Data container name: DC_001

Dedup ○ Compression ●

Soft Quota (GB) *: 1000
Hard Quota (GB) *: 1500

Protocol: NFS
Date created: 10/11/2018 11:18
Date modified: 10/11/2018 12:45

| | |
|---|---|
| Hot Tier (with snapshots) | 0.0 B |
| Object Tier | 0.0 B |
| Total | 0.0 B |

## 4.3 Disaster Recovery for Data Containers

Data containers can be asynchronously replicated to remote sites for disaster recovery. For more information on enabling asynchronous replication, see Section 6 - Managing Asynchronous Disaster Recovery.

### 4.3.1 Replicating Data Containers for Disaster Recovery

**To replicate a data container for disaster recovery:**

1. On the local site, in the Assets menu, click (**DATA CONTAINERS**) and click the required data container.

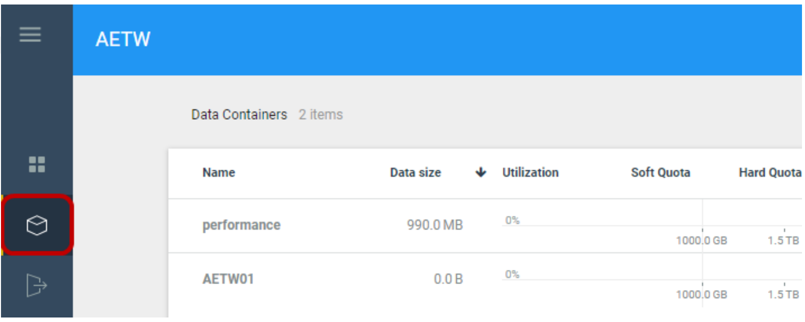

Data Containers 2 items

| Name | Data size | Utilization | Soft Quota | Hard Quota |
|---|---|---|---|---|
| performance | 990.0 MB | 0% | 1000.0 GB | 1.5 TB |
| AETW01 | 0.0 B | 0% | 1000.0 GB | 1.5 TB |

2. Click **REPLICATE**.

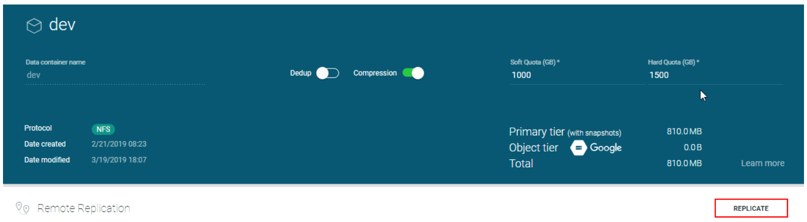

dev

Data container name: dev

Dedup (off) Compression (on)

Soft Quota (GB) *: 1000

Hard Quota (GB) *: 1500

Protocol: NFS
Date created: 2/21/2019 08:23
Date modified: 3/19/2019 18:07

Primary tier (with snapshots) 810.0 MB
Object tier Google 0.0 B
Total 810.0 MB

If a replication service does not yet exist, click **Add replication service** and follow the instructions as detailed in Section 6.1 - Enabling an Initial Replication Service for Asynchronous Replication from Step 2 till the end of the Section 6.2 - Pairing Local and Remote Sites.

3. Configure the following fields as described in the table below:

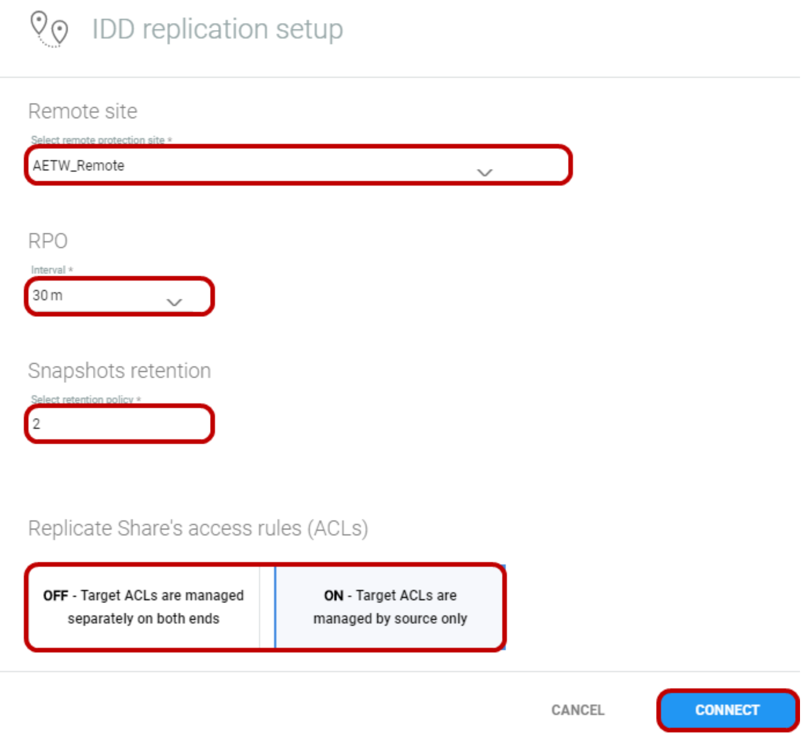

IDD replication setup

Remote site

Select remote protection site *
AETW_Remote

RPO

Interval *
30 m

Snapshots retention

Select retention policy *
2

Replicate Share's access rules (ACLs)

**OFF** - Target ACLs are managed separately on both ends

**ON** - Target ACLs are managed by source only

| Field | Description |
|---|---|
| **Remote Site** | The required remote site. |
| **RPO** | The required Recovery Point Objective. |
| **Snapshots retention** | Number of snapshots retained on the remote site. |
| **Replicate Share's access rules (ACLs)** | If during recovery you will migrate the local clients to the remote site, select **ON**.<br>If you select ON, you will not be able to modify the ACLs until the remote DC is active. |

4. Click **CONNECT**. When the data container replication configuration has been setup successfully, a notification appears at the bottom of the window. Click **CLOSE**.

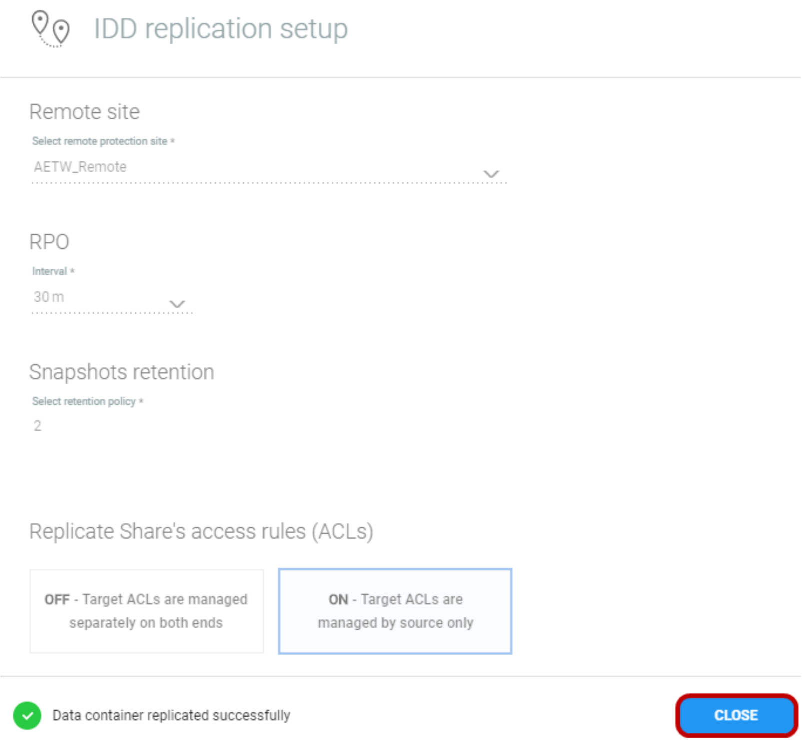

IDD replication setup

Remote site

Select remote protection site *: AETW_Remote

RPO

Interval *: 30 m

Snapshots retention

Select retention policy *: 2

Replicate Share's access rules (ACLs)

**OFF** - Target ACLs are managed separately on both ends

**ON** - Target ACLs are managed by source only

✓ Data container replicated successfully

5. The **Remote Replication** details for the data container appears.

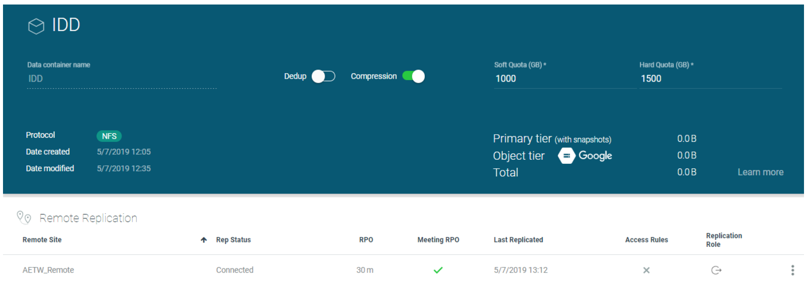

IDD

Data container name: IDD

Dedup ○ Compression ●

Soft Quota (GB) *: 1000

Hard Quota (GB) *: 1500

Protocol: NFS
Date created: 5/7/2019 12:05
Date modified: 5/7/2019 12:35

Primary tier (with snapshots): 0.0 B
Object tier Google: 0.0 B
Total: 0.0 B

Learn more

Remote Replication

| Remote Site | Rep Status | RPO | Meeting RPO | Last Replicated | Access Rules | Replication Role |
|---|---|---|---|---|---|---|
| AETW_Remote | Connected | 30 m | ✓ | 5/7/2019 13:12 | × | |

The data container starts replicating itself as soon as the connection is created. The time required to complete replication depends on the amount of data in the data container and the connection bandwidth.

The RPO is enforced only after the initial replication has successfully completed.

**To view the data containers that are replicated to particular remote site:**

1. In the Assets menu, click **REMOTE SITES**.

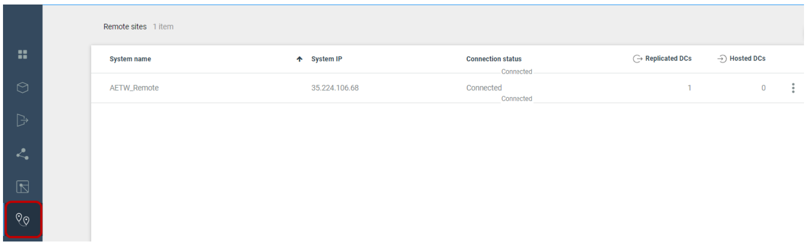

Remote sites 1 item

| System name | System IP | Connection status | Replicated DCs | Hosted DCs |
|---|---|---|---|---|
| AETW_Remote | 35.224.106.68 | Connected | 1 | 0 |

2. Click the remote site name. The replicated data containers are listed.

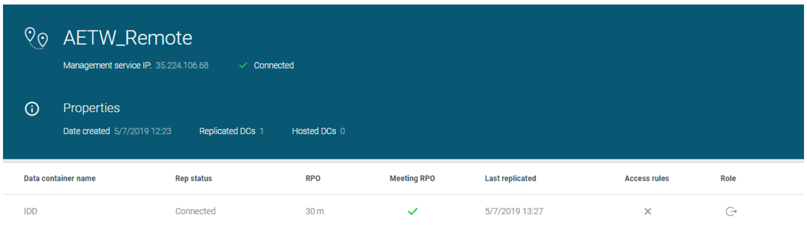

AETW_Remote

Management service IP: 35.224.106.68 ✓ Connected

Properties

Date created 5/7/2019 12:23 Replicated DCs 1 Hosted DCs 0

| Data container name | Rep status | RPO | Meeting RPO | Last replicated | Access rules | Role |
|---|---|---|---|---|---|---|
| IDD | Connected | 30 m | ✓ | 5/7/2019 13:27 | × | |

## 4.3.2 Swapping Data Container Status between Active and Passive

During normal operations, the local data container status is active and the replication on the remote site is passive.

When a disaster recovery situation occurs, you may need to change the status of the remote data container to active so all read/write operations are directed to it. In addition, if the local site is accessible, you need to change the status of its local pair to passive to prevent a "collision", which halts the replication.

### 4.3.2.1 Swapping Data Container Status when the Local Site is Accessible

**To make a remote data container active if the local site is accessible:**

1. Connect to the remote site Management Console.
2. On the remote site, in the Assets menu, click (**DATA CONTAINERS**) and click the required data container.
3. Click **MAKE ACTIVE**.

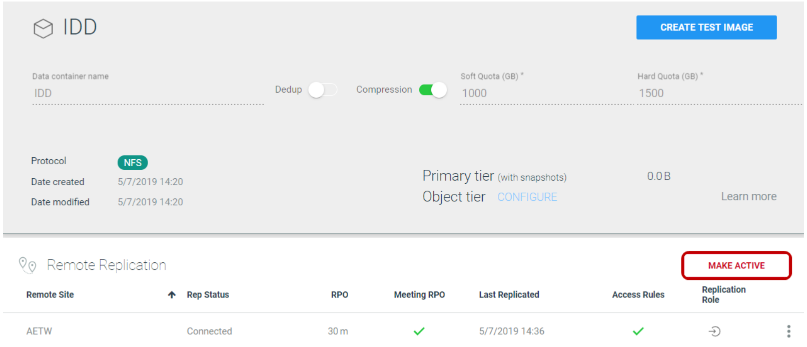

IDD

Data container name: IDD

Dedup ○ Compression ●

Soft Quota (GB) *: 1000

Hard Quota (GB) *: 1500

Protocol NFS

Date created 5/7/2019 14:20

Date modified 5/7/2019 14:20

Primary tier (with snapshots) 0.0 B

Object tier CONFIGURE

Remote Replication

| Remote Site | Rep Status | RPO | Meeting RPO | Last Replicated | Access Rules | Replication Role |
|---|---|---|---|---|---|---|
| AETW | Connected | 30 m | ✓ | 5/7/2019 14:36 | ✓ | |

4. Select the **Switch roles** check box and click **MAKE ACTIVE**.

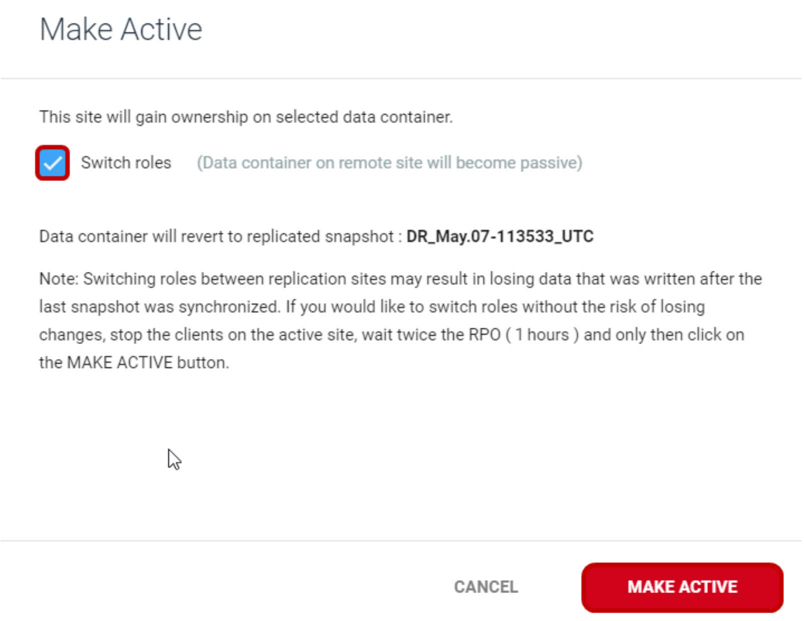

Make Active

This site will gain ownership on selected data container.

☑ Switch roles (Data container on remote site will become passive)

Data container will revert to replicated snapshot : **DR_May.07-113533_UTC**

Note: Switching roles between replication sites may result in losing data that was written after the last snapshot was synchronized. If you would like to switch roles without the risk of losing changes, stop the clients on the active site, wait twice the RPO ( 1 hours ) and only then click on the MAKE ACTIVE button.

5. The snapshot data is copied to the current data container, the status of its local pair is changed to passive and when completed, the status **done** appears.
6. Click **CLOSE**.

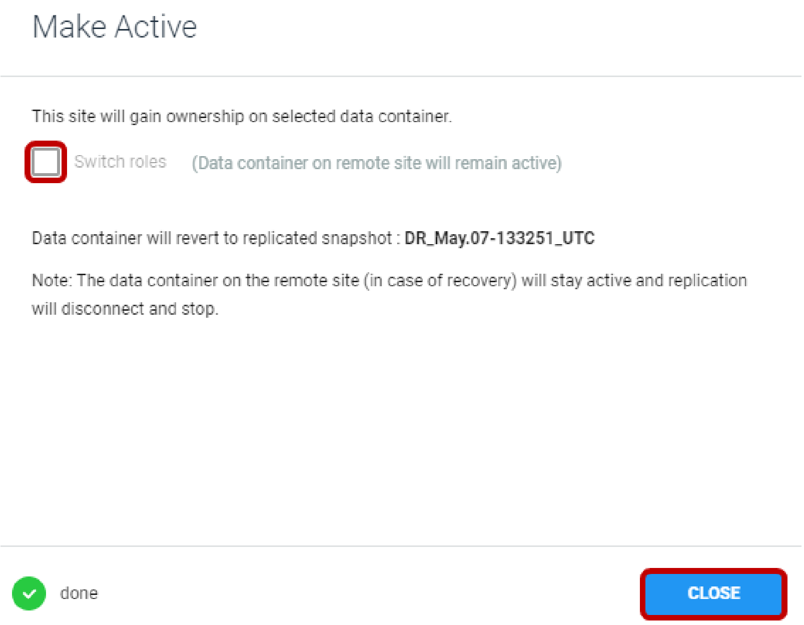

Make Active

This site will gain ownership on selected data container.

□ Switch roles (Data container on remote site will remain active)

Data container will revert to replicated snapshot : **DR_May.07-133251_UTC**

Note: The data container on the remote site (in case of recovery) will stay active and replication will disconnect and stop.

✓ done

The time required to copy the snapshot to the current data container depends on the size of the snapshot.

#### 4.3.2.2 Swapping Data Container Status when the Local Site is Inaccessible

**To make a remote data container active if the local site is inaccessible:**

1. Connect to the remote site Management Console.
2. On the remote site, in the Assets menu, click ⬡ (**DATA CONTAINERS**) and click the required data container.

3. Click **MAKE ACTIVE**.

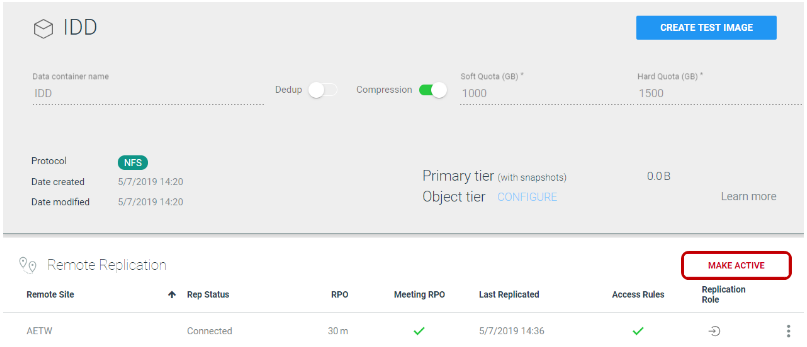

IDD

Data container name: IDD

Dedup

Compression

Soft Quota (GB) *: 1000

Hard Quota (GB) *: 1500

Protocol: NFS

Date created: 5/7/2019 14:20

Date modified: 5/7/2019 14:20

Primary tier (with snapshots) 0.0 B

Object tier CONFIGURE

Remote Replication

| Remote Site | Rep Status | RPO | Meeting RPO | Last Replicated | Access Rules | Replication Role |
|---|---|---|---|---|---|---|
| AETW | Connected | 30 m | ✓ | 5/7/2019 14:36 | ✓ | |

4. Clear the **Switch roles** check box and click **MAKE ACTIVE**.

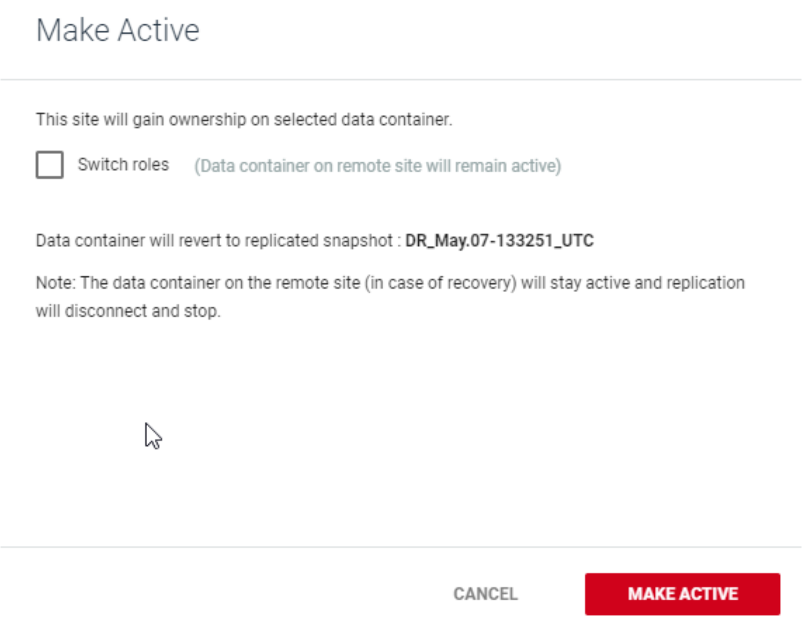

Make Active

This site will gain ownership on selected data container.

□ Switch roles (Data container on remote site will remain active)

Data container will revert to replicated snapshot : **DR_May.07-133251_UTC**

Note: The data container on the remote site (in case of recovery) will stay active and replication will disconnect and stop.

5. The snapshot data is copied to the current data container and when completed, the status **done** appears.
6. Click **CLOSE**.

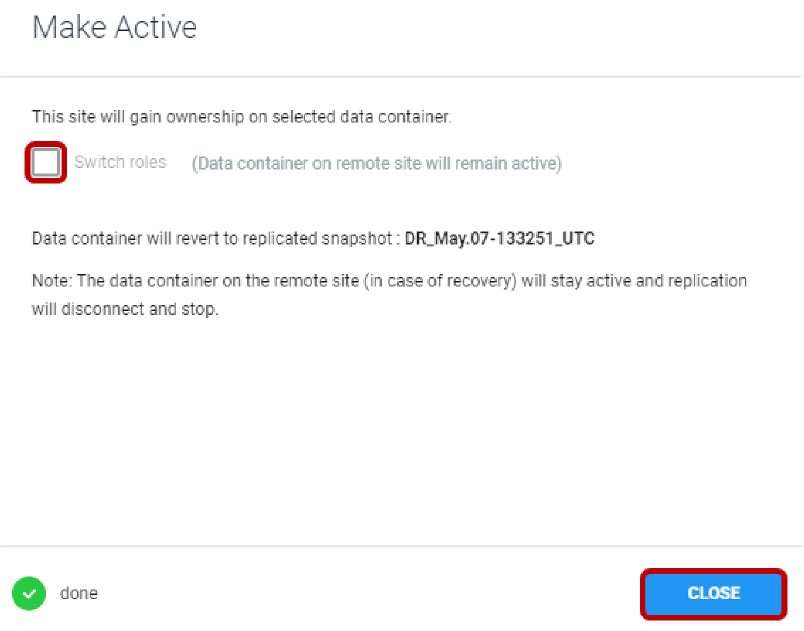

Make Active

This site will gain ownership on selected data container.

□ Switch roles (Data container on remote site will remain active)

Data container will revert to replicated snapshot : **DR_May.07-133251_UTC**

Note: The data container on the remote site (in case of recovery) will stay active and replication will disconnect and stop.

✓ done

The time required to copy the snapshot to the current data container depends on the size of the snapshot.

**To make the local site passive after it becomes accessible again:**

1. On the local site, in the Assets menu, click (**DATA CONTAINERS**) and click the required data container.

2. Under Remote Replication, click ⋮ (**Options**) .

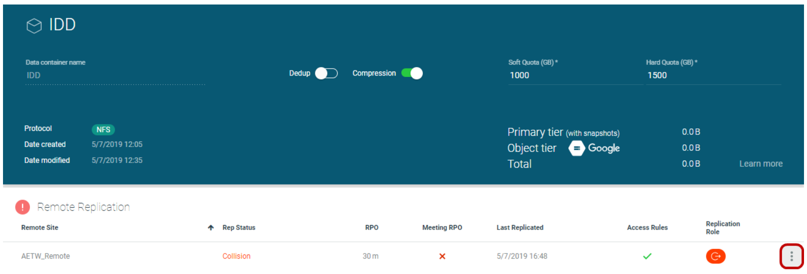

IDD

Data container name: IDD

Dedup ○ Compression ●

Soft Quota (GB) *: 1000

Hard Quota (GB) *: 1500

Protocol: NFS

Date created: 5/7/2019 12:05

Date modified: 5/7/2019 12:35

Primary tier (with snapshots): 0.0 B

Object tier Google: 0.0 B

Total: 0.0 B

Learn more

Remote Replication

| Remote Site | Rep Status | RPO | Meeting RPO | Last Replicated | Access Rules | Replication Role | |
|---|---|---|---|---|---|---|---|
| AETW_Remote | Collision | 30 m | ✕ | 5/7/2019 16:48 | ✓ | | ⋮ |

3. Click **Make local passive**.

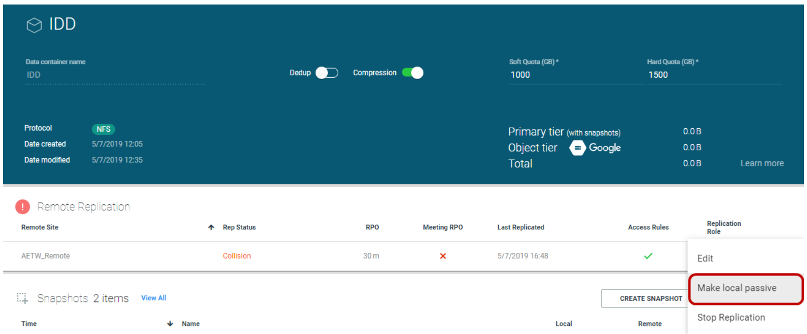

IDD

Data container name: IDD

Dedup ○ Compression ●

Soft Quota (GB) *: 1000

Hard Quota (GB) *: 1500

Protocol: NFS

Date created: 5/7/2019 12:05

Date modified: 5/7/2019 12:35

Primary tier (with snapshots): 0.0 B

Object tier Google: 0.0 B

Total: 0.0 B

Learn more

Remote Replication

| Remote Site | Rep Status | RPO | Meeting RPO | Last Replicated | Access Rules | Replication Role |
|---|---|---|---|---|---|---|
| AETW_Remote | Collision | 30 m | ✕ | 5/7/2019 16:48 | ✓ | |

Snapshots 2 items View All

CREATE SNAPSHOT

| Time | Name | Local | Remote |
|---|---|---|---|

4. Click **YES, MAKE PASSIVE**.

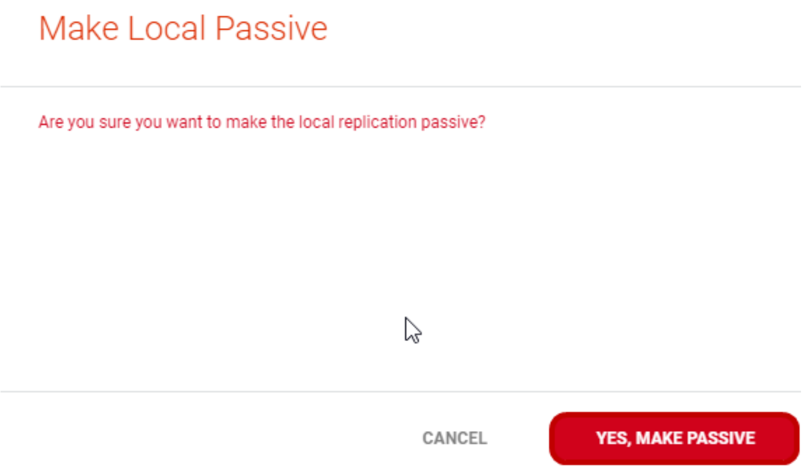

Make Local Passive

Are you sure you want to make the local replication passive?

5. When you want to make the active site active again, perform the procedure detailed in Section 4.3.2.1 - Swapping Data Container Status when the Local Site is Accessible.

### 4.3.3 Creating Data Container Test Images

You can create a test image of a passive data container that is hosted on your ECFS system The test image is created from the latest replicated snapshot to verify read/write operations and data consistency.

**To create a data container test image:**

1. On the passive site, in the Assets menu, click (**DATA CONTAINERS**) and click the required data container.

2. Click **CREATE TEST IMAGE**.

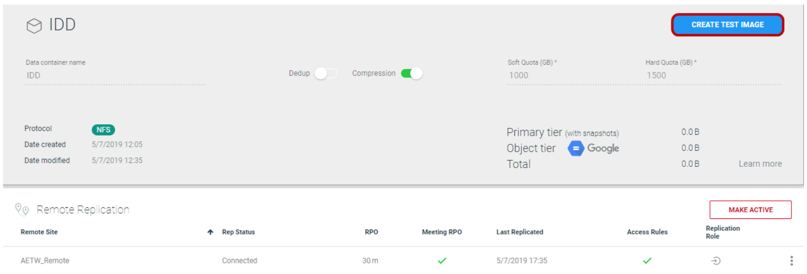

IDD

CREATE TEST IMAGE

Data container name: IDD
Dedup: off
Compression: on
Soft Quota (GB) *: 1000
Hard Quota (GB) *: 1500

Protocol: NFS
Date created: 5/7/2019 12:05
Date modified: 5/7/2019 12:35

| | | |
|---|---|---|
| Primary tier (with snapshots) | | 0.0 B |
| Object tier | Google | 0.0 B |
| Total | | 0.0 B |

Learn more

Remote Replication

MAKE ACTIVE

| Remote Site | Rep Status | RPO | Meeting RPO | Last Replicated | Access Rules | Replication Role |
|---|---|---|---|---|---|---|
| AETW_Remote | Connected | 30 m | ✓ | 5/7/2019 17:35 | ✓ | |

3. In **Data container name**, type a new data container name and click **CREATE**.

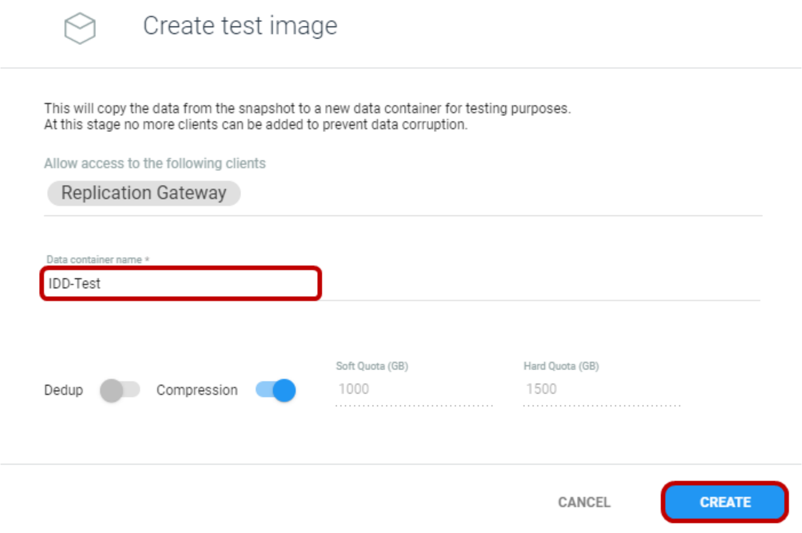

Create test image

This will copy the data from the snapshot to a new data container for testing purposes.
At this stage no more clients can be added to prevent data corruption.

Allow access to the following clients: Replication Gateway

Data container name *: IDD-Test

Dedup: off
Compression: on
Soft Quota (GB): 1000
Hard Quota (GB): 1500

4. When the test image is created, click **CLOSE**.

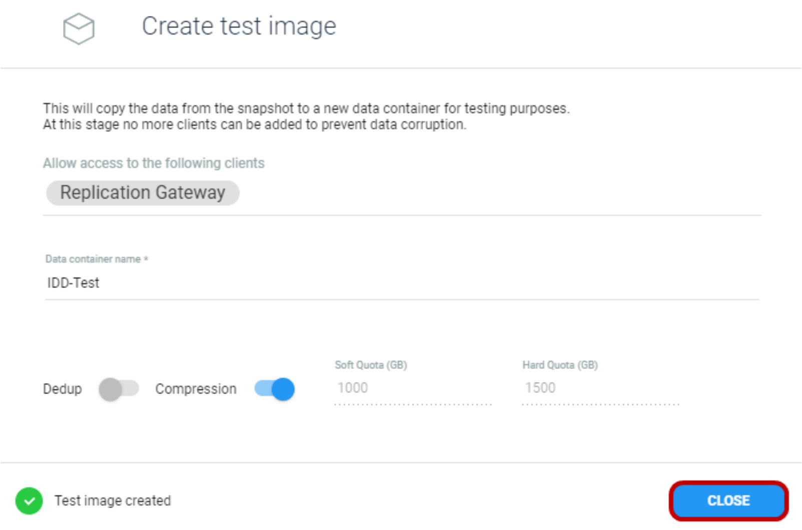

Create test image

This will copy the data from the snapshot to a new data container for testing purposes.
At this stage no more clients can be added to prevent data corruption.

Allow access to the following clients: Replication Gateway

Data container name *: IDD-Test

Dedup: off
Compression: on
Soft Quota (GB): 1000
Hard Quota (GB): 1500

Test image created

5. The snapshot data is copied to the new data container.

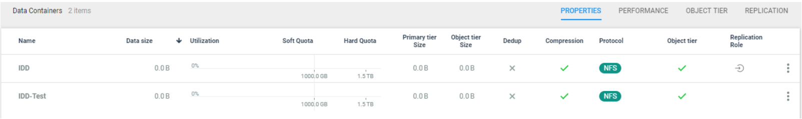

Data Containers 2 items

PROPERTIES PERFORMANCE OBJECT TIER REPLICATION

| Name | Data size | Utilization | Soft Quota | Hard Quota | Primary tier Size | Object tier Size | Dedup | Compression | Protocol | Object tier | Replication Role |
|---|---|---|---|---|---|---|---|---|---|---|---|
| IDD | 0.0 B | 0% | 1000.0 GB | 1.5 TB | 0.0 B | 0.0 B | ✕ | ✓ | NFS | ✓ | ⇥ |
| IDD-Test | 0.0 B | 0% | 1000.0 GB | 1.5 TB | 0.0 B | 0.0 B | ✕ | ✓ | NFS | ✓ | |

The time required to copy depends on the size of the snapshot.

Once the data is copied, you can modify the new data container share to add clients that will verify data integrity.

## 4.3.4 Stopping Replication for a Data Container

**To stop replication for a data container:**

1. In the Assets menu, click (**DATA CONTAINERS**) and click the required data container.

2. On the required **Remote site** row, click ⋮ (**Options**). and click **Stop Replication**.

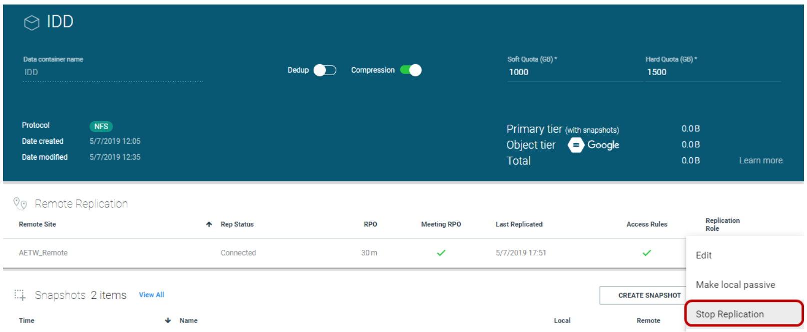

IDD

Data container name: IDD

Dedup ○ Compression ●

Soft Quota (GB) *: 1000

Hard Quota (GB) *: 1500

Protocol: NFS

Date created: 5/7/2019 12:05

Date modified: 5/7/2019 12:35

Primary tier (with snapshots) 0.0 B

Object tier Google 0.0 B

Total 0.0 B

Learn more

Remote Replication

| Remote Site | Rep Status | RPO | Meeting RPO | Last Replicated | Access Rules | Replication Role |
|---|---|---|---|---|---|---|
| AETW_Remote | Connected | 30 m | ✓ | 5/7/2019 17:51 | ✓ | |

Snapshots 2 items View All

| Time | Name | Local | Remote |
|---|---|---|---|

3. In the **Stop Replication** window, select the **Allow the replication to finish and only then stop.** check box if you want replication to complete, then click **YES, STOP**.

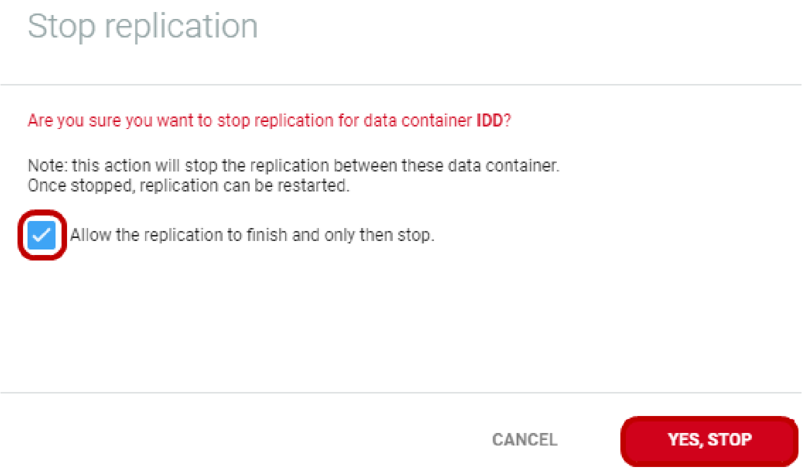

Stop replication

Are you sure you want to stop replication for data container **IDD**?

Note: this action will stop the replication between these data container. Once stopped, replication can be restarted.

☑ Allow the replication to finish and only then stop.

## 4.4 Upscaling Your System

You can upscale your system by adding additional hosts as described in the ECFS Installation Guide for your environment.

# 5. Monitoring System Performance

You can view system performance metrics analytics and alerts regarding ECFS's current performance and recent activity through the following:

- Dashboard - displays an aggregated view of status and analytics for key system parameters.
- Assets - displays the inventory for the asset you selected.
- System View - displays essential information for each host.

## 5.1 Monitoring through the Dashboard

The dashboard view displays essential information such as system capacity, performance, maintenance information and an infrastructure review.

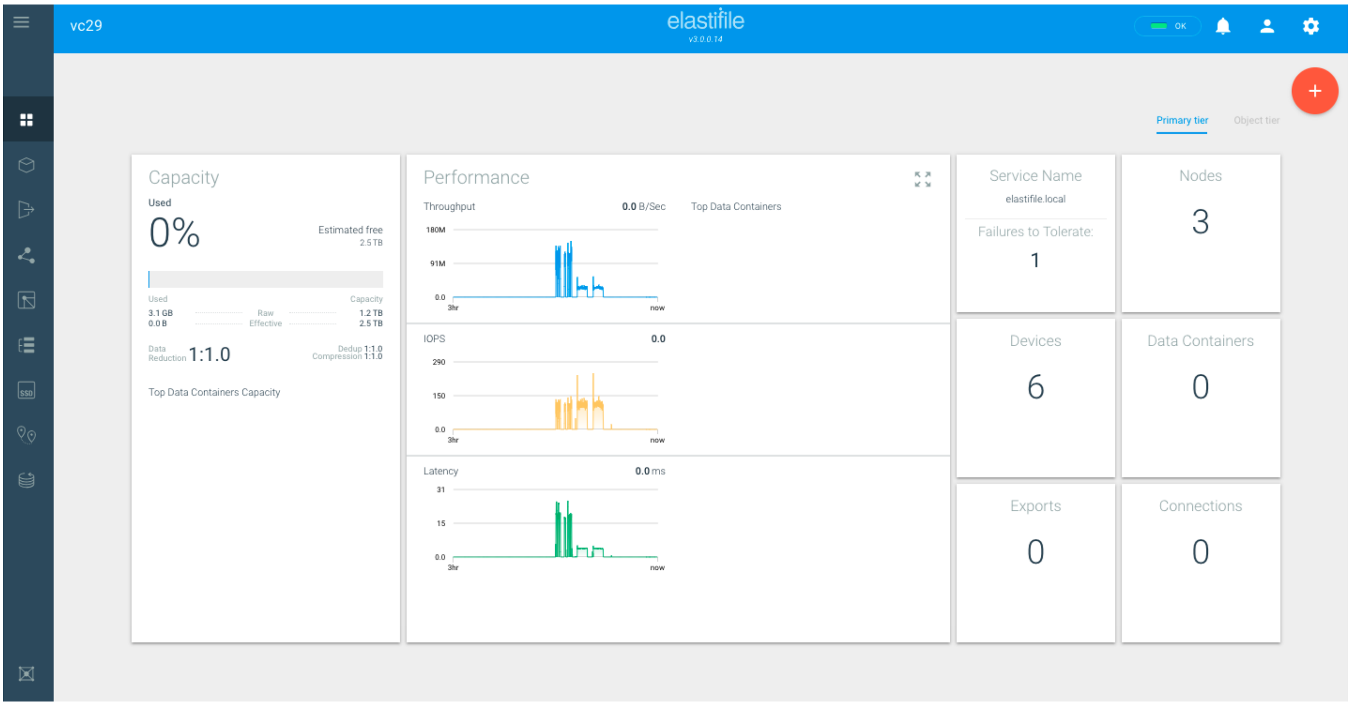

| Item | Description |
|---|---|
| **Capacity** | See Section 5.2 - Capacity. |
| **Performance** | See Section 5.3 - Performance. |
| **Service Name** | Information displayed here depends on the load balancer option you selected during deployment:<br>• If using **Round robin DNS**: the NFS service name, is displayed.<br>• If using the **Cloud load balancer**: the cloud load balancer NFS IP address is displayed. |
| **Nodes** | The number of storage nodes in the cluster |
| **Failures to Tolerate** | Displays the number of servers that can fail without putting the ECFS cluster in degraded mode, as configured during deployment. |

| Item | Description |
|---|---|
| **Devices** | Displays the number of devices in the system that are:<br>• Active (large numerals)<br>• Down<br>• Suspended<br>• Read only<br>• Approaching limit |
| **Data Containers** | Displays the number of data containers in the system that are:<br>• Active (large numerals)<br>• Suspended |
| **Exports** | Displays the number of shares (exports) in the system that are:<br>• Active (large numerals)<br>• Suspended |
| **Connections** | Displays the number of mounts that have been initiated by the clients.<br>If the Virtual Controller Mode is configured to Storage backend and Client frontend, **Connections** is not displayed. |

## 5.2 Capacity

The dashboard capacity pane displays the information on system capacity.

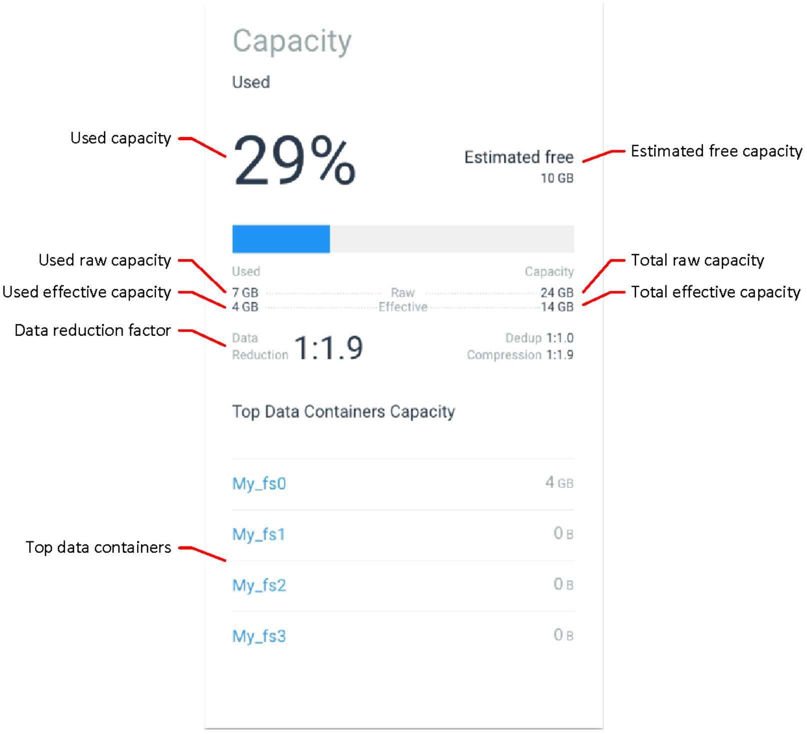

| Item | Description |
|---|---|
| **Used capacity** | Used capacity as a percentage of the total capacity (whether raw or effective). |
| **Used raw capacity** | Bytes used on system devices. |
| **Used effective capacity** | Actual data size before taking into account data reduction and replication level. |
| **Data reduction factor** | Level of data reduction based on deduplication and compression. |

| Item | Description |
|---|---|
| **Total effective capacity** | Predicted data size based on current reduction factor. |
| **Total raw capacity** | Total bytes available on system devices. |
| **Estimated free** | Estimated free capacity on system devices ("total effective capacity" minus "used effective capacity"). |
| **Top data containers** | List of data containers contributing the most to the total effective capacity. |

## 5.3 Performance

### 5.3.1 Performance Summary Pane

The dashboard performance pane displays a summary of the system performance metrics.

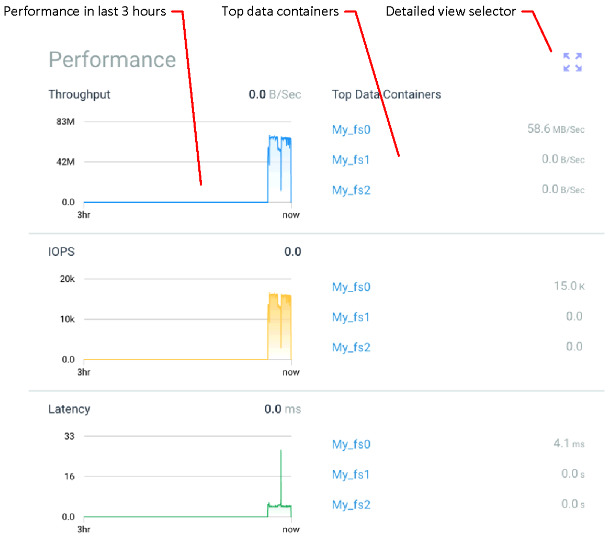

Top Data Containers

| | |
|---|---|
| My_fs0 | 58.6 MB/Sec |
| My_fs1 | 0.0 B/Sec |
| My_fs2 | 0.0 B/Sec |
| My_fs0 | 15.0 K |
| My_fs1 | 0.0 |
| My_fs2 | 0.0 |
| My_fs0 | 4.1 ms |
| My_fs1 | 0.0 s |
| My_fs2 | 0.0 s |

**To display a summary of performance for Throughput, IOPS and Latency:**

- Click the required (**Throughput**, **IOPS** or **Latency**) parameter.

## 5.3.2 Detailed Performance Window

**To display the detailed performance window:**

- Click ⛶ (**DETAILED VIEW SELECTOR**).

The detailed performance window is divided into the following main elements:

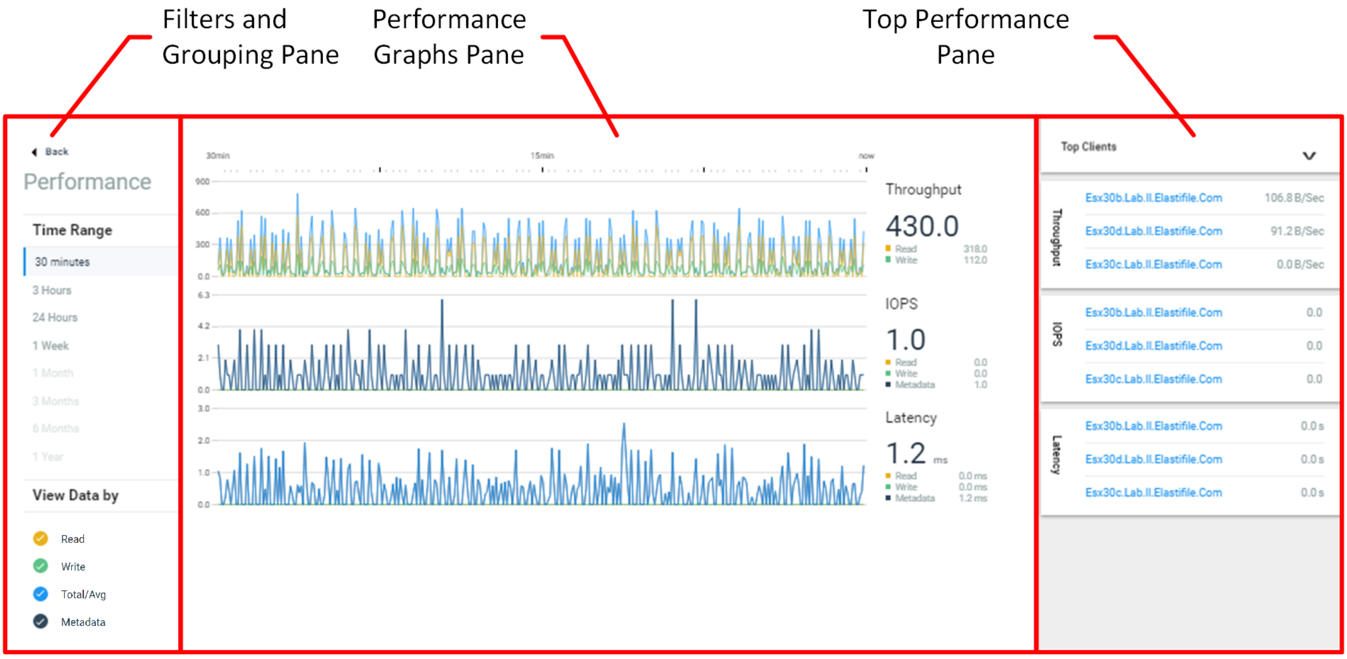

| | | |
|---|---|---|
| Throughput | Esx30b.Lab.Il.Elastifile.Com | 106.8 B/Sec |
| | Esx30d.Lab.Il.Elastifile.Com | 91.2 B/Sec |
| | Esx30c.Lab.Il.Elastifile.Com | 0.0 B/Sec |
| IOPS | Esx30b.Lab.Il.Elastifile.Com | 0.0 |
| | Esx30d.Lab.Il.Elastifile.Com | 0.0 |
| | Esx30c.Lab.Il.Elastifile.Com | 0.0 |
| Latency | Esx30b.Lab.Il.Elastifile.Com | 0.0 s |
| | Esx30d.Lab.Il.Elastifile.Com | 0.0 s |
| | Esx30c.Lab.Il.Elastifile.Com | 0.0 s |

| Item | Description |
|---|---|
| Filters and Grouping Pane | Displays the time range filter options and view data filter options. See Filters and Grouping Pane. |
| Performance Graphs Pane | Displays filtered metric data in graphical formats. See Performance Graph Pane. |
| Top Performance Pane | Displays top data containers and clients. See Top Performance Pane. |

### 5.3.3 Filters and Grouping Pane

This pane provides options for filtering and grouping the data to be displayed.

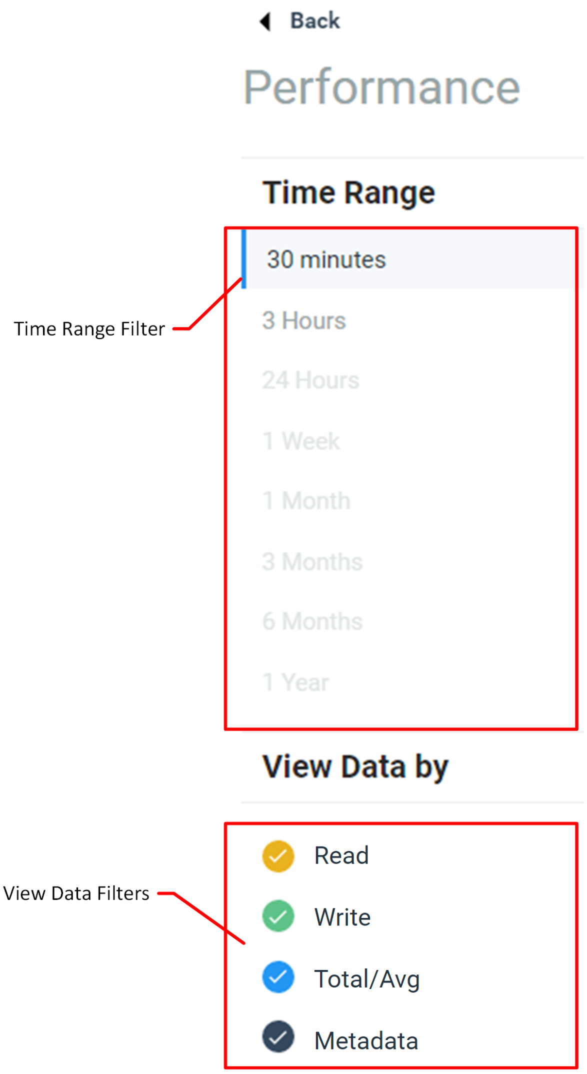

The Filters and Grouping items are described in the following table:

| Item | Description |
|---|---|
| **Time Range** | Defines the performance period you want to review. |
| **View Data** | Select the check boxes for toggle view (on / off). This will show/ hide the associated information from:<br>• All the graphs |

### 5.3.4 Performance Graph Pane

This pane displays the filtered and grouped metric data as graphs.

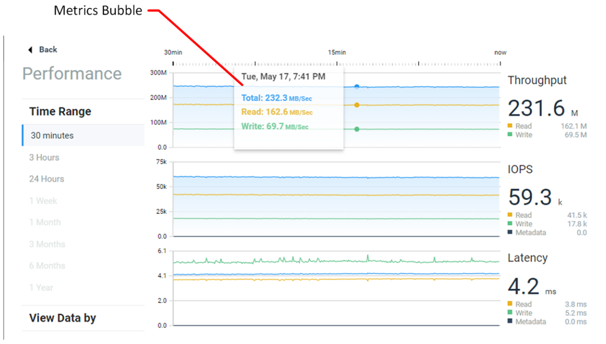
Metrics Bubble

The performance graph are described in the following table:

| Item | Description |
|---|---|
| Metrics bubble | Hover over the graph to display the metric bubble displaying metrics for that particular point on the graph. |

## 5.4 Monitoring Assets

For each type of asset, you view the inventory and see each asset's properties and performance details. These views can be used to pinpoint issues when troubleshooting the system.

**To display the properties view for assets:**

1. In the Assets Menu, click the required asset (in this example: (**SERVICES**):

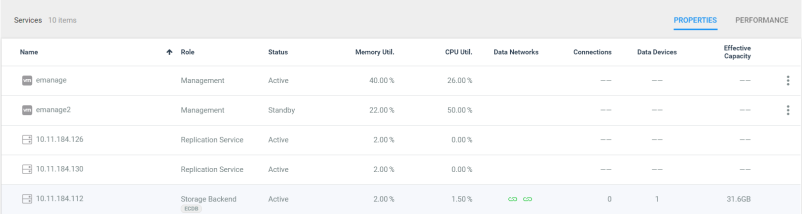

| Name | Role | Status | Memory Util. | CPU Util. | Data Networks | Connections | Data Devices | Effective Capacity | |
|---|---|---|---|---|---|---|---|---|---|
| emanage | Management | Active | 40.00 % | 26.00 % | | –– | –– | –– | ⋮ |
| emanage2 | Management | Standby | 22.00 % | 50.00 % | | –– | –– | –– | ⋮ |
| 10.11.184.126 | Replication Service | Active | 2.00 % | 0.00 % | | –– | –– | –– | |
| 10.11.184.130 | Replication Service | Active | 2.00 % | 0.00 % | | –– | –– | –– | |
| 10.11.184.112 | Storage Backend ECDB | Active | 2.00 % | 1.50 % | | 0 | 1 | 31.6GB | |

2. When viewing data containers, clients and services, click **PERFORMANCE** to display Throughput, IOPS and Latency performance:

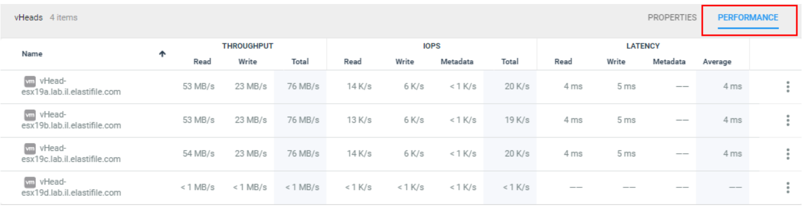

| Name | THROUGHPUT | | | IOPS | | | | LATENCY | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Read | Write | Total | Read | Write | Metadata | Total | Read | Write | Metadata | Average | |
| vHead-esx19a.lab.il.elastifile.com | 53 MB/s | 23 MB/s | 76 MB/s | 14 K/s | 6 K/s | < 1 K/s | 20 K/s | 4 ms | 5 ms | –– | 4 ms | ⋮ |
| vHead-esx19b.lab.il.elastifile.com | 53 MB/s | 23 MB/s | 76 MB/s | 13 K/s | 6 K/s | < 1 K/s | 19 K/s | 4 ms | 5 ms | –– | 4 ms | ⋮ |
| vHead-esx19c.lab.il.elastifile.com | 54 MB/s | 23 MB/s | 76 MB/s | 14 K/s | 6 K/s | < 1 K/s | 20 K/s | 4 ms | 5 ms | –– | 4 ms | ⋮ |
| vHead-esx19d.lab.il.elastifile.com | < 1 MB/s | < 1 MB/s | < 1 MB/s | < 1 K/s | < 1 K/s | < 1 K/s | < 1 K/s | –– | –– | –– | –– | ⋮ |

| Item | Description |
|---|---|
| **Throughput** | Current total throughput of data accessing the asset, with a breakdown into read and write. |
| **IOPS** | Rate of IO operations, with a breakdown into read and write. |
| **Latency** | Average latency, with a breakdown into read and write. |
| ⋮ (**Options**) | For future use |

### 5.4.1 Top Performance Pane

Switch between the data containers and client views to see top performance displays for the following metrics:

- **Throughput** - Top assets that have the highest throughput within the selected time frame.
- **IOPS** - Top assets that have the highest IOPS within the selected time frame.
- **Latency** - Top assets that have the highest latency within the selected time frame.

## 5.5 Monitoring Data Containers

**To display the properties view for data containers:**

1. In the Assets Menu, click (**DATA CONTAINERS**):

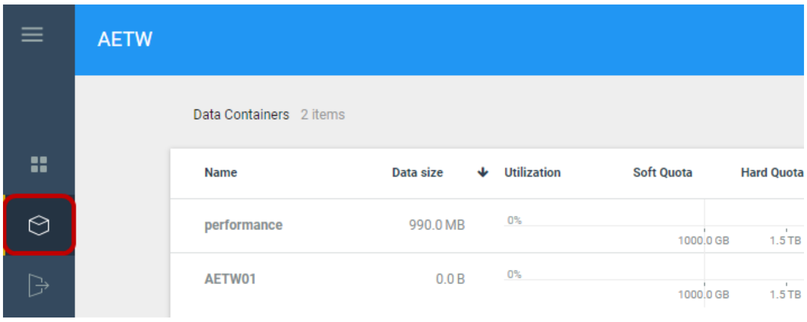

AETW

Data Containers 2 items

| Name | Data size | Utilization | Soft Quota | Hard Quota |
|---|---|---|---|---|
| performance | 990.0 MB | 0% | 1000.0 GB | 1.5 TB |
| AETW01 | 0.0 B | 0% | 1000.0 GB | 1.5 TB |

> You can sort certain columns by clicking the column header. Change the sort order by reclicking the same column.

| Item | Description |
|---|---|
| **Name** | Name you used for your data container. |
| **Data size** | Size of the data stored in this data container, including primary and object tiers. |
| **Utilization** | Percentage of the allocated quota used by the DC on the primary tier. |
| **Soft Quota** | Storage (in GB) which, when passed, raises a notification. |
| **Hard Quota** | Storage (in GB) limit which cannot be surpassed. |
| **Primary Tier Size** | Size of the data from this data container stored in the primary tier. |
| **Object Tier Size** | Size of the data from this data container stored in the object tier. |
| **Dedup** | Indicates if deduplication is enabled or disabled for this data container. |
| **Compression** | Indicates if compression is enabled or disabled for this data container. |
| **Protocol** | File system protocol used by this data container. Currently, only NFSv3 is supported. |
| **Object Tier** | N/A |
| **Replication Role** | Role of the data controller in an Async-DR pair (either source or target). |
| ⋮ (**Options**) | Click to edit or delete the data container on that row. |

**To display detailed information for a specific data container:**

1. In the **Data Container** column, click the required data container.

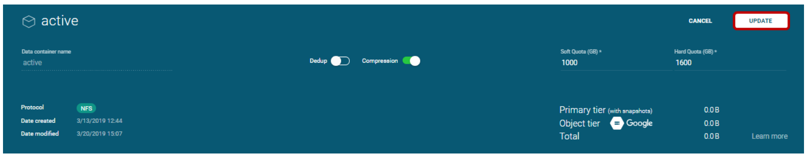

Data container name: active

Dedup: off

Compression: on

Soft Quota (GB) *: 1000

Hard Quota (GB) *: 1600

Protocol: NFS

Date created: 3/13/2019 12:44

Date modified: 3/20/2019 15:07

Primary tier (with snapshots): 0.0 B

Object tier Google: 0.0 B

Total: 0.0 B

Learn more

| Item | Description |
|---|---|
| **Data container name** | The name you want to use for your data container. Field messages appear for invalid names (such as names already used or names that are too short). |
| **Dedup** | Toggle to enable/disable data deduplication for the data container. |
| **Compression** | Toggle to enable/disable compression for the data container. |
| **Soft Quota** | Storage (in GB) which, when passed, raises a notification. |
| **Hard Quota** | Storage (in GB) limit which cannot be surpassed (optional) |
| **Protocol** | Click to select file system protocol (currently only NFSv3 is supported). |
| **Date Created** | Date the data container was created |
| **Date modified** | Date the data container was last modified. |
| **Primary Tier (with snapshots)** | Size of the data from this data container, including snapshots, that is stored on the primary tier. |
| **Object tier** | Size of the data from this data container, including snapshots, that is stored on the object tier. |
| **Total** | Click **Learn more** to view more details on distribution of data container data among tiers. |

In addition, there are four panes for managing the following features in relation to this data container:

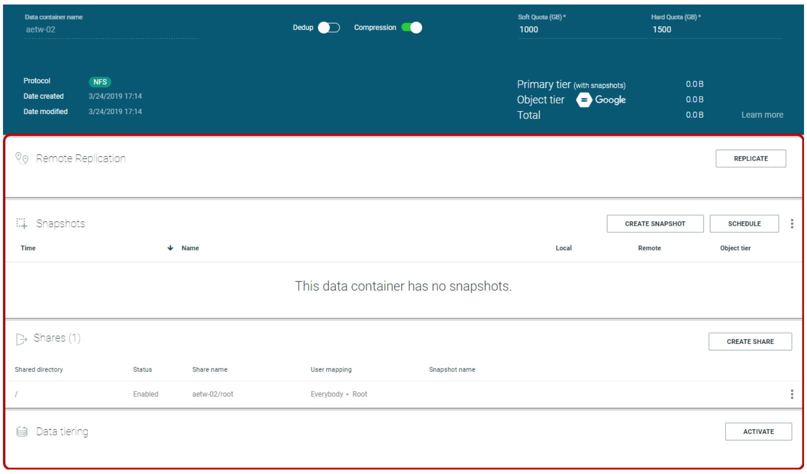

Data container name: aetw-02
Dedup: off
Compression: on
Soft Quota (GB) *: 1000
Hard Quota (GB) *: 1500

Protocol: NFS
Date created: 3/24/2019 17:14
Date modified: 3/24/2019 17:14

| | |
|---|---|
| Primary tier (with snapshots) | 0.0 B |
| Object tier Google | 0.0 B |
| Total | 0.0 B |

Remote Replication

Snapshots

| Time | Name | Local | Remote | Object tier |
|---|---|---|---|---|

This data container has no snapshots.

Shares (1)

| Shared directory | Status | Share name | User mapping | Snapshot name |
|---|---|---|---|---|
| / | Enabled | aetw-02/root | Everybody ▸ Root | |

Data tiering

| | |
|---|---|
| **Remote Replication** | Configure and manage Async DR for this data container. |
| **Snapshots** | Create and manage snapshots for this data container. |
| **Shares** | Configure and manage shares for this data container. |
| **Data tiering** | Configure and manage data tiering for this data container. |

2. Click **PERFORMANCE** to display Throughput, IOPS and Latency performance:

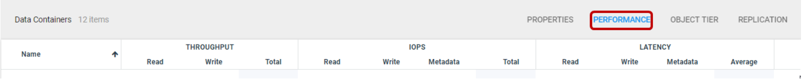

Data Containers 12 items

| Name | THROUGHPUT | | | IOPS | | | | LATENCY | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | Read | Write | Total | Read | Write | Metadata | Total | Read | Write | Metadata | Average |

| Item | Description |
|---|---|
| **Throughput** | Current total throughput of data accessing the asset, with a breakdown into read and write. |
| **IOPS** | Rate of IO operations, with a breakdown into read and write. |
| **Latency** | Average latency, with a breakdown into read and write. |
| ⋮ **(Options)** | Click to edit or delete the data container on that row. |

3. Click **OBJECT TIER** to display object tier statistics for the data containers:

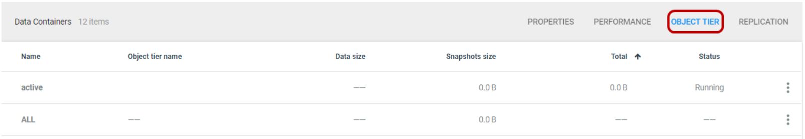

Data Containers 12 items

| Name | Object tier name | Data size | Snapshots size | Total | Status | |
|---|---|---|---|---|---|---|
| active | | –– | 0.0 B | 0.0 B | Running | ⋮ |
| ALL | –– | –– | 0.0 B | –– | –– | ⋮ |

| Item | Description |
|---|---|
| **Name** | Name you used for your data container. |
| **Object tier name** | N/A |
| **Data size** | Size of the cold data stored on the object tier. |
| **Snapshots size** | Size of the snapshots stored on the object tier. |
| **Total** | Total size of data stored on the object tier (cold data and snapshots). |
| **Status** | Object tier status. |
| ⋮ **(Options)** | Click to edit or delete the data container on that row. |

4. Click **REPLICATION** to display object tier statistics for the data containers:

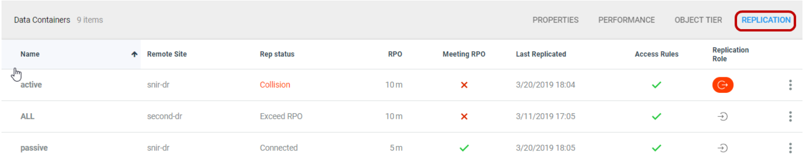

Data Containers 9 items

| Name | Remote Site | Rep status | RPO | Meeting RPO | Last Replicated | Access Rules | Replication Role | |
|---|---|---|---|---|---|---|---|---|
| active | snir-dr | Collision | 10 m | ✗ | 3/20/2019 18:04 | ✓ | | ⋮ |
| ALL | second-dr | Exceed RPO | 10 m | ✗ | 3/11/2019 17:05 | ✓ | | ⋮ |
| passive | snir-dr | Connected | 5 m | ✓ | 3/20/2019 18:05 | ✓ | | ⋮ |

| Item | Description |
|---|---|
| **Name** | Name you used for your data container. |
| **Remote Site** | Name of the remote site to which the data container is paired. |
| **Rep status** | Status of the replication connection. |
| **RPO** | Recovery point objective (RPO) time set for this replication connection. |
| **Meeting RPO** | Indicates if the last sync occured within the required RPO. |
| **Last Replicated** | Last date/time a full sync was completed. |
| **Access Rules** | Indicates if access rules are replicated to the remote site. |
| **Replication Role** | The role of this data container in the replication pair. |
| ⋮ **(Options)** | Click to edit or delete the data container on that row. |

## 5.6 Monitoring Events

You can monitor all the events that the system logged.

**To view the events log:**

1. Click 🔔 (**Alerts**).

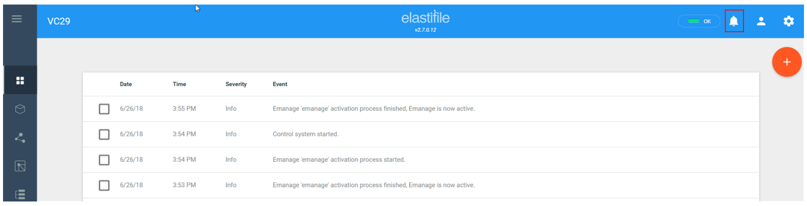

| | Date | Time | Severity | Event |
|---|---|---|---|---|
| □ | 6/26/18 | 3:55 PM | Info | Emanage 'emanage' activation process finished, Emanage is now active. |
| □ | 6/26/18 | 3:54 PM | Info | Control system started. |
| □ | 6/26/18 | 3:54 PM | Info | Emanage 'emanage' activation process started. |
| □ | 6/26/18 | 3:53 PM | Info | Emanage 'emanage' activation process finished, Emanage is now active. |

The **Events** view displays events logged by the system:

| Element | Description |
|---|---|
| **Date** and **Time** | The events view table is organized chronologically, with the latest events listed first. You cannot sort the table according to any other category. |
| **Severity** | Displays the event severity. Possible values are:<br>• Critical<br>• Info |
| **Event** | Event details. |

## 5.7 Monitoring through System View

The system view displays essential information for each host such as EMS network status, power state, devices connected, type of storage, and capacity. In addition, several system actions can also be performed from this window.

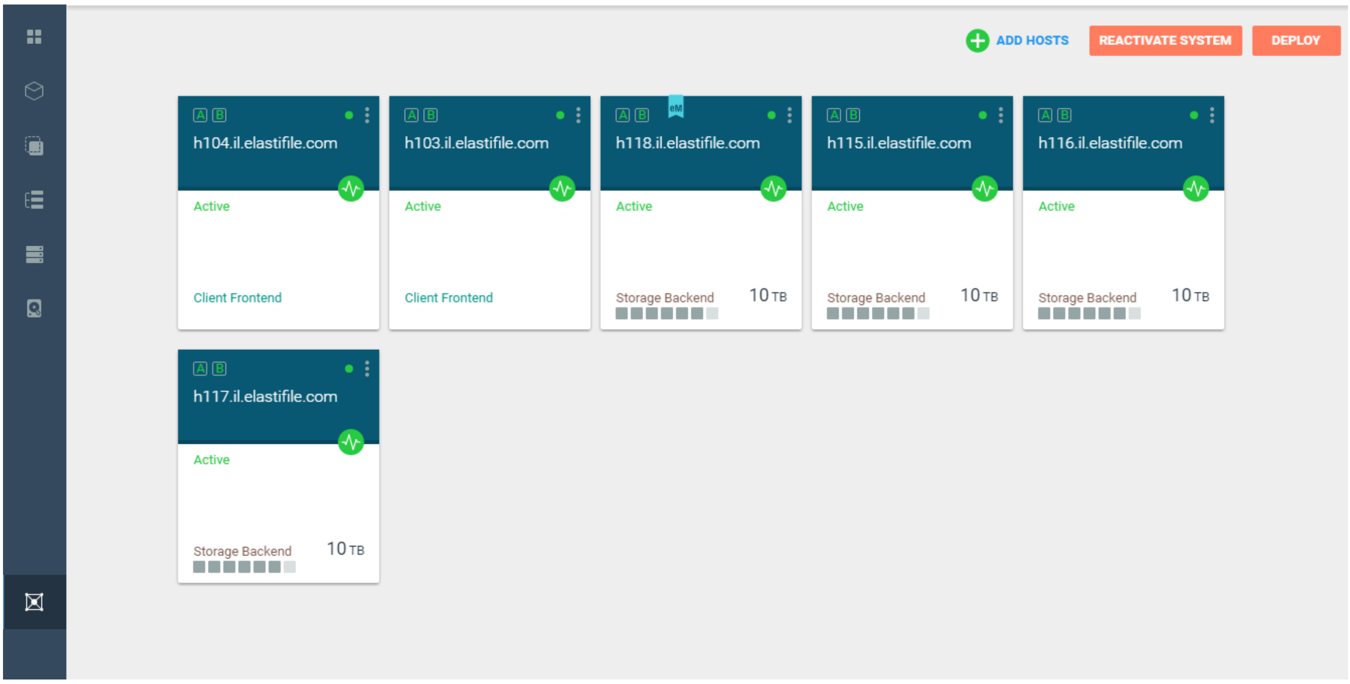

### 5.7.1 Host Information

Each host displays the following information:

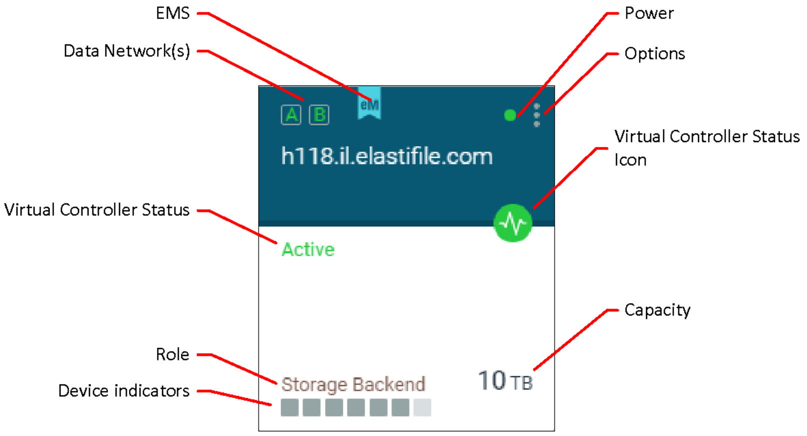

| Item | Description |
|---|---|
| EMS | ECFS Management Service (ECFS's VM which also hosts the Management Console). |
| Network | Displays the status of the data network(s) used by the controller:<br>• Green - network is connected.<br>• Red - network is disconnected.<br>• Grey - network status is unknown. |
| Virtual Controller Status | Display the status of the controller. The color indicates the status as follows:<br>• Green - controller is active.<br>• Red - any other errors. |
| Virtual Controller Status Icon | Icon representing the controller status. |
| Role | Virtual controller role:<br>• Client Frontend/Storage Backend<br>• Hyperconverged |
| Device indicators | Displays devices connected to the server:<br>• Dark gray - devices used by ECFS.<br>• Light gray - devices used by other applications. |
| Capacity | Raw capacity available to the system. |
| ⋮ (**Options**) | • Configure - (disabled)<br>• Install - (disabled)<br>• Remove from system - remove the host from the ECFS system. |
| Power | If green, indicates that the host is powered. |

## 5.8 Alerts

**To see critical events:**

- On the Header, click 28 (**Alerts**).

## 5.9 Notifiers

The notifiers display feedback for actions and errors. There are two types of notifiers:

- Confirmation (black toast).
  - Textual - For inserts. Disappears after 3 seconds.
  - Textual + Undo link - For delete and updates. Disappears after 5 seconds.

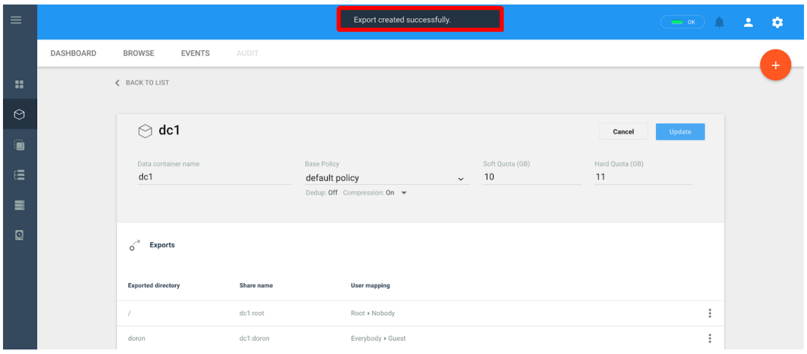

Export created successfully.

dc1

Data container name: dc1
Base Policy: default policy
Dedup: Off Compression: On
Soft Quota (GB): 10
Hard Quota (GB): 11

Exports

| Exported directory | Share name | User mapping |
|---|---|---|
| / | dc1:root | Root ▸ Nobody |
| doron | dc1:doron | Everybody ▸ Guest |

- Errors (orange toast) - Disappears when error is corrected / error status cleared.

## 5.10 System Health Indicator

The System Health indicator, located on the header, shows the current ability of the system to maintain a good service level. Health issues occur when there is a degradation in the service level, and/or an increase in the risk of service degradation or loss.

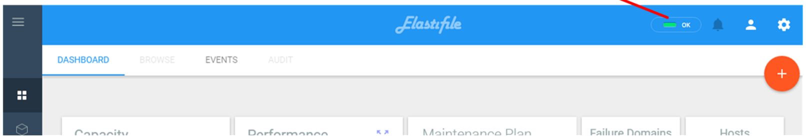

System Health

The System Health Indicator displays the following levels (in order of severity):

- OK - green.
- Attention - yellow.
- Alert - orange.
- Risk - red.
- Critical - flashing red.
- Down - black.

# 6. Managing Asynchronous Disaster Recovery

To asynchronously replicate data containers to remote sites for disaster recovery (as described in Section 4.3 - Disaster Recovery for Data Containers), you must first configure a replication service. Elastifile recommends enabling at least two replication services for load balancing and redundancy.

The maximum number of replication servers should take into consideration the amount of data to replicate. In general, Elastifile recommends one replication service per controller, up to a maximum of 10 replication services per system.

> For asynchronous replication, the infrastructure must be able to support the RPO (Recovery Point Objective) and the maximum volume of data in a replication backup for that RPO.

Setting up asynchronous disaster recovery requires three main stages:

1. Adding replication services to local and remote sites, see Enabling an Initial Replication Service for Asynchronous Replication.
2. Pairing local and remote sites, see Pairing Local and Remote Sites.
3. Connecting local site data containers to remote sites, see Disaster Recovery for Data Containers.

In addition, you can manage your remote site disaster recovery, including the following tasks:

- viewing remote sites, site details, and errors, see Viewing Remote Protection Sites.
- editing local and remote site credentials for pairings, see Editing Local and Remote Site Pairings.
- removing remote sites to manage disaster recovery, see Removing Remote Site Pairing.

## 6.1 Enabling an Initial Replication Service for Asynchronous Replication

**To add an initial replication service:**

1. In the Assets menu, click **REMOTE SITES**, then under 1 **Add replication service**, click **System view**.

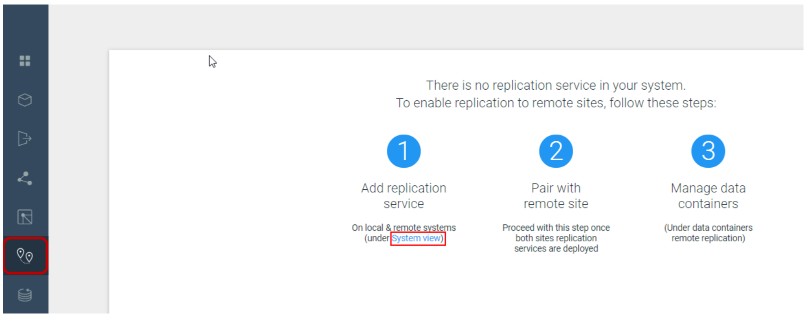

There is no replication service in your system.
To enable replication to remote sites, follow these steps:

1 Add replication service
On local & remote systems (under System view)

2 Pair with remote site
Proceed with this step once both sites replication services are deployed

3 Manage data containers
(Under data containers remote replication)

2. Click **SETUP REPLICATION SERVICE**.

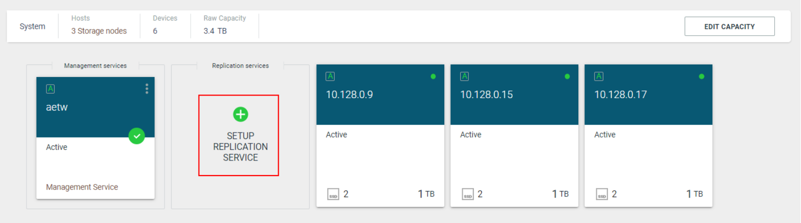

| System | Hosts | Devices | Raw Capacity |
|---|---|---|---|
| | 3 Storage nodes | 6 | 3.4 TB |

3. Select the check boxes as follows:
   - **Enable High availability** - recommended.
   - **Use external network for replication traffic** - select if the remote site is located in a different GCP project or in a different cloud provider.
4. Click **ADD**.

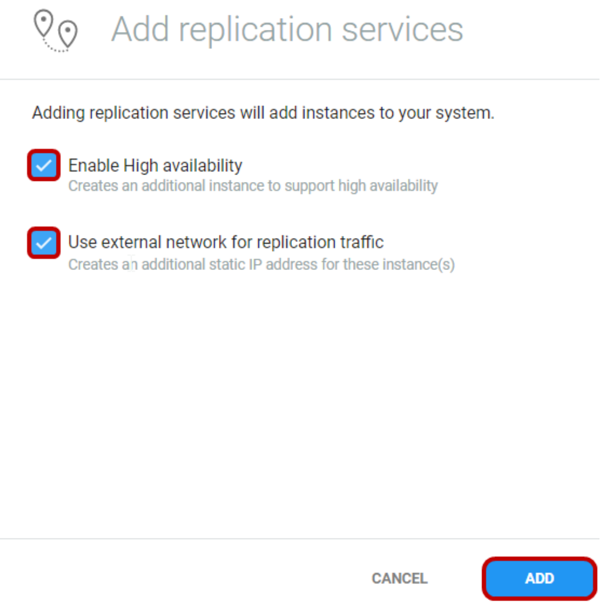

Add replication services

Adding replication services will add instances to your system.

☑ Enable High availability
Creates an additional instance to support high availability

☑ Use external network for replication traffic
Creates an additional static IP address for these instance(s)

5. After the replication services are added, click **CLOSE**.

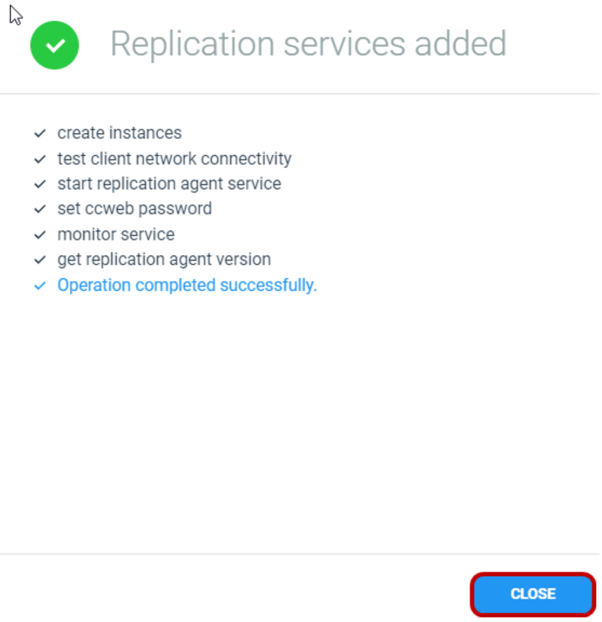

Replication services added

✓ create instances
✓ test client network connectivity
✓ start replication agent service
✓ set ccweb password
✓ monitor service
✓ get replication agent version
✓ Operation completed successfully.

6. Two replication instances with external networks are created.

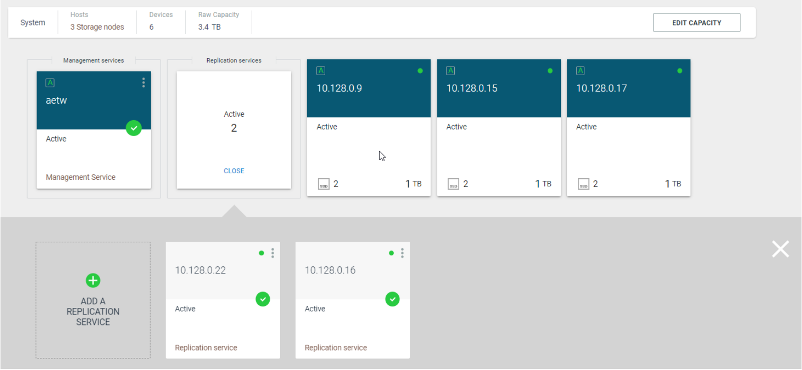

7. Repeat the above procedure for the remote ECFS.

> To add additional replication services for load balancing and redundancy, see Section 6.4 - Adding Additional Replication Services

## 6.2 Pairing Local and Remote Sites

**To pair a local and remote site:**

1. In the Assets menu, click **REMOTE SITES**, then on the FAB click **(Pair with remote site**).

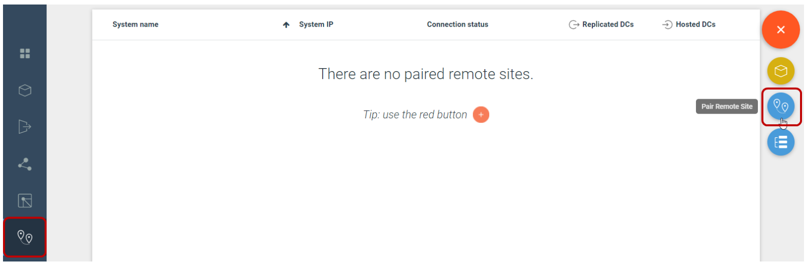

2. The **Pair with remote site** window appears.
3. Under **Remote site details**, in **Management service IP**, type the IP address of the remote site.
4. In **User name**, type the required user name for the remote site.
5. In **Password**, type the remote site password.
6. Under **Local site details**, in **Management service IP**, type the required IP address of the local site.
7. In **User name**, type the required local site user name.
8. In **Password**, type the local site password.

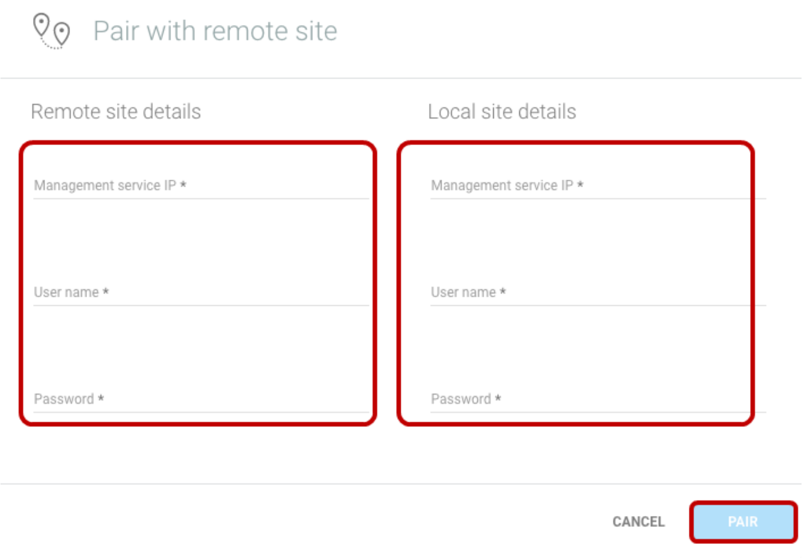

9. Click **PAIR**.

10. The remote site IP address, credentials, and version compatibility are validated. If valid and the connection is successful, the local and remote sites are paired.

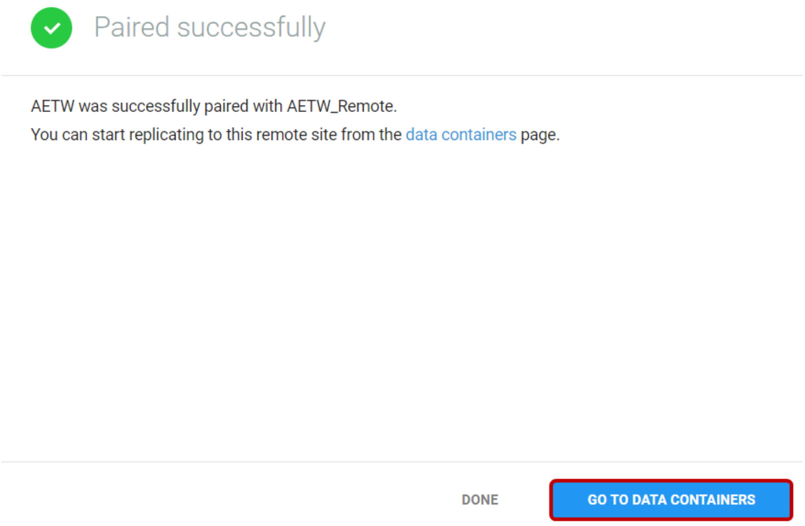

Paired successfully

AETW was successfully paired with AETW_Remote.
You can start replicating to this remote site from the data containers page.

11. Do one of the following:

    - To replicate a data container, click **GO TO DATA CONTAINERS** and refer to Section 4.3.1 - Replicating Data Containers for Disaster Recovery.

    - Otherwise, click **DONE**. A list of **Remote sites** appears.

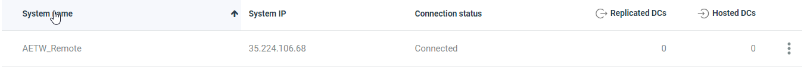

| System name | System IP | Connection status | Replicated DCs | Hosted DCs |
|---|---|---|---|---|
| AETW_Remote | 35.224.106.68 | Connected | 0 | 0 |

## 6.3 Managing Remote Protection Sites

### 6.3.1 Viewing Remote Protection Sites

**To view remote protection sites:**

1. Click **REMOTE SITES**.

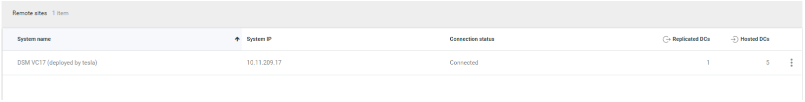

Remote sites 1 item

| System name ↑ | System IP | Connection status | Replicated DCs | Hosted DCs |
|---|---|---|---|---|
| DSM VC17 (deployed by tesla) | 10.11.209.17 | Connected | 1 | 5 |

| Item | Description |
|---|---|
| System Name | The name of the remote protection site. |
| System IP | The IP address of the remote protection site. |
| Connection Status | The connection status of the local site and the remote protection site pairing, including the following:<br>• **Connected**<br>• **Failed**<br>• **Disconnecting** |
| Replicated DCs | The number of replicated data containers on the remote protection site. |
| Hosted DCs | The number of hosted data containers on the local site. |

**To view the details for a remote protection site:**

1. Double click the required remote site name.

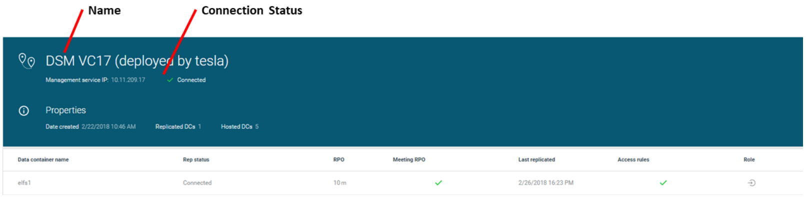

Name

Connection Status

DSM VC17 (deployed by tesla)

Management service IP: 10.11.209.17 ✓ Connected

Properties

Date created 2/22/2018 10:46 AM Replicated DCs 1 Hosted DCs 5

| Data container name | Rep status | RPO | Meeting RPO | Last replicated | Access rules | Role |
|---|---|---|---|---|---|---|
| elfs1 | Connected | 10 m | ✓ | 2/26/2018 16:23 PM | ✓ | ⇥ |

| Item | Description |
|---|---|
| Name | The name of the remote protection site. |
| Management service IP | The IP address of the remote protection site. |
| Connection Status | The status of the local site-remote protection site connection as follows:<br>• **Connected**<br>• **Disconnected**<br>• **Failed**<br>• **Disconnecting** |
| Date Created | The date and time of the remote protection site pairing. |
| Replicated DCs | The number of replicated data containers and their total size on the remote protection site. |
| Hosted DCs | The number of hosted data containers and their total size on the local site. |

| Item | Description |
|---|---|
| Data container name | The names of the data containers on the remote protection site. |
| Rep Status | The status of the local and remote site connection:<br>• **Connected**<br>• **Stopped**<br>• **Failed**<br>• **Disconnecting** |
| RPO | RPO (Recovery Point Objective) or time interval threshold for a connection disruption after which the replicated data container backup exceeds the maximum tolerance for recovery. The RPO is in minutes. |
| Meeting RPO | The RPO status indicating whether the RPO is being met. |
| Last Replication | The date and time of the last replication. |
| Access Rules | Status indicating if ACLs are replicated on the remote site. |
| Role | Replication role of the data container, either:<br>• ↪ - the data container is a replicated container.<br>• ⇥ - the data container is a hosted container. |

## 6.3.2 Editing Local and Remote Site Pairings

Edit local and remote site pairings to manage site credentials.

**To edit a local and remote site pairing:**

1. Click **REMOTE SITES**.

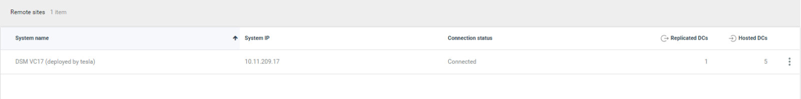

Remote sites 1 item

| System name | System IP | Connection status | Replicated DCs | Hosted DCs |
|---|---|---|---|---|
| DSM VC17 (deployed by tesla) | 10.11.209.17 | Connected | 1 | 5 |

2. On the required remote site row, click ⋮ (**Options**). The options menu appears.

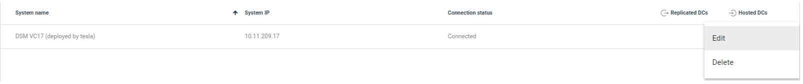

| System name | System IP | Connection status | Replicated DCs | Hosted DCs |
|---|---|---|---|---|
| DSM VC17 (deployed by tesla) | 10.11.209.17 | Connected | | |

3. Click **Edit** and edit the remote site and local site details as required:
   - In the **Remote site details** section, in **Management service IP**, type the IP address of the remote site.
   - In **User name**, type the required user name for the remote site.
   - In **Password**, type the remote site password.

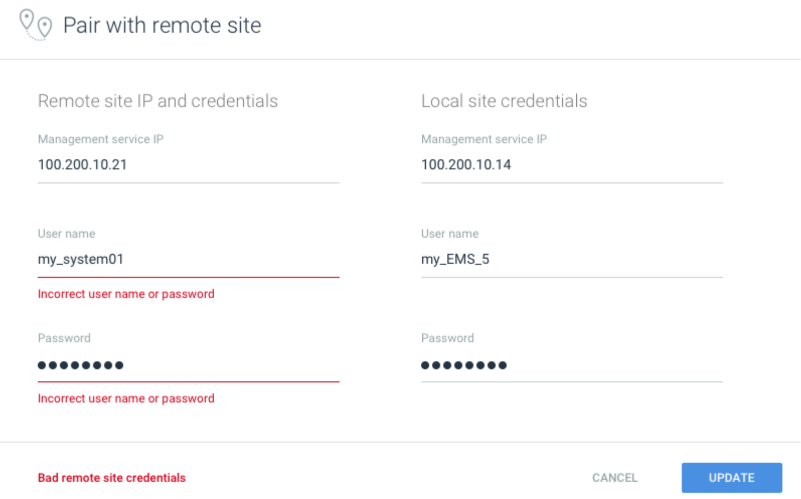

Pair with remote site

Remote site IP and credentials

Management service IP: 100.200.10.21

User name: my_system01

Incorrect user name or password

Password: ●●●●●●●●

Incorrect user name or password

Local site credentials

Management service IP: 100.200.10.14

User name: my_EMS_5

Password: ●●●●●●●●

Bad remote site credentials

- In the **Local site details**, in **Management service IP**, type the required IP address of the local site.
- In **User name**, type the required local site user name.
- In **Password**, type the local site password.

4. Click **UPDATE**. The remote site IP address, credentials, and version compatibility are validated. If valid and the connection is successful, the local and remote sites are paired.

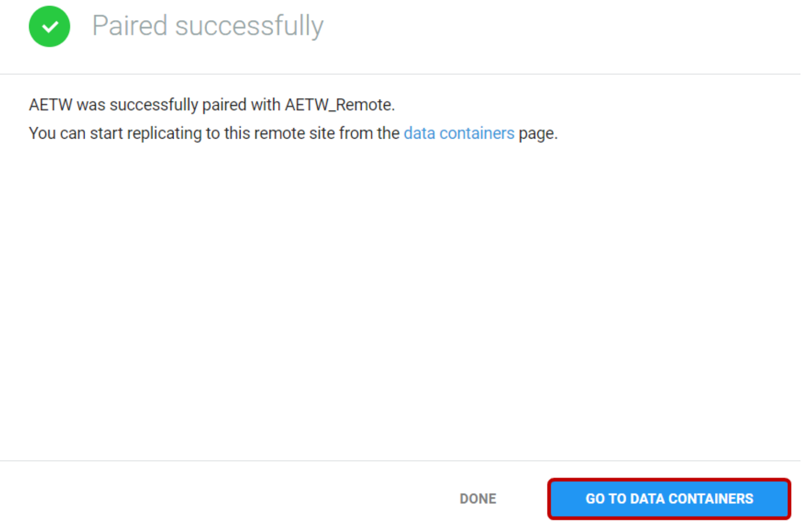

Paired successfully

AETW was successfully paired with AETW_Remote.
You can start replicating to this remote site from the data containers page.

## 6.3.3 Removing Remote Site Pairing

Remove remote protection sites to manage disaster recovery sites.

All hosted and replicated data containers connected to a remote protection site must be disconnected before removing a remote protection site, see Stopping Replication for a Data Container.

**To remove a remote protection site pairing:**

1. Click **REMOTE SITES**.

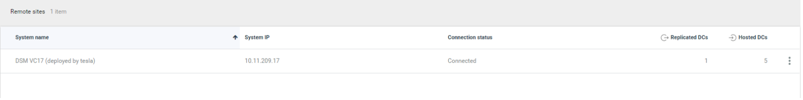

Remote sites 1 item

| System name | System IP | Connection status | Replicated DCs | Hosted DCs |  |
|---|---|---|---|---|---|
| DSM VC17 (deployed by tesla) | 10.11.209.17 | Connected | 1 | 5 | ⋮ |

2. On the required remote site row, click ⋮ (**Options**).

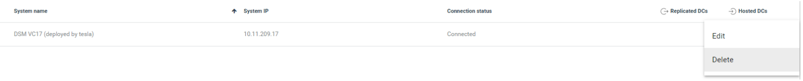

| System name | System IP | Connection status | Replicated DCs | Hosted DCs |
|---|---|---|---|---|
| DSM VC17 (deployed by tesla) | 10.11.209.17 | Connected |  |  |

3. Click **Delete** then click **YES, DELETE**. The local and remote site pairing is removed.

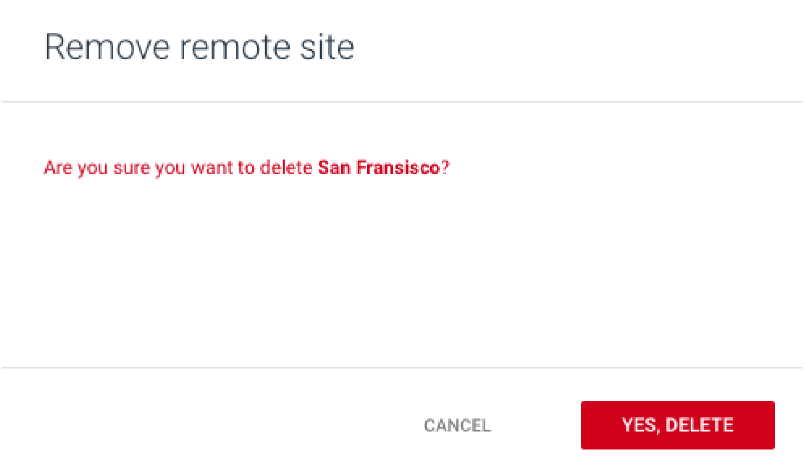

Remove remote site

Are you sure you want to delete **San Fransisco**?

## 6.4 Adding Additional Replication Services

**To add another replication service:**

1. Click ☒ **SYSTEM VIEW**.

2. Click **VIEW ALL**.

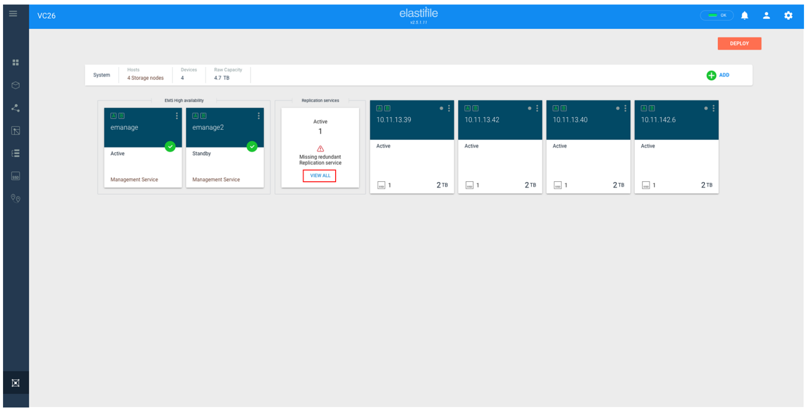

3. Click **ADD A REPLICATION SERVICE**.

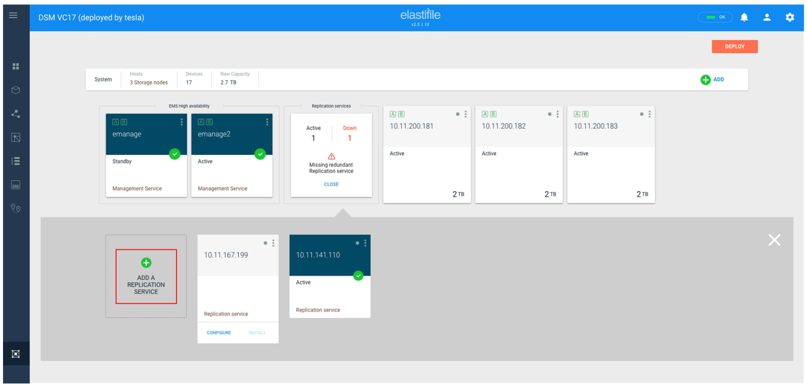

4. Perform Steps 1 to 6 in Section 0.1 - Enabling an Initial Replication Service for Asynchronous Replication.

5. In **System View**, click **INSTALL**.

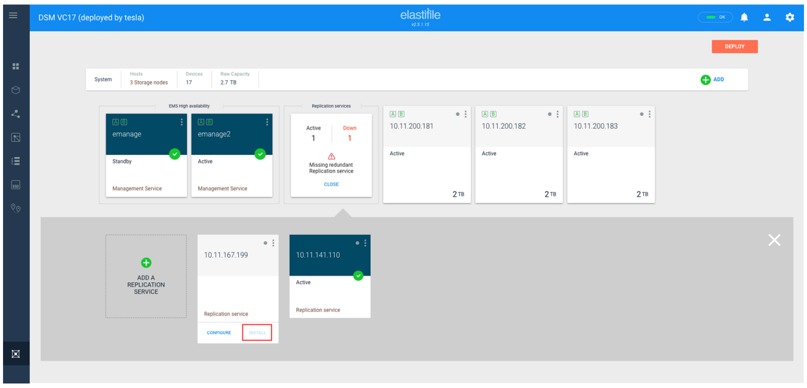

6. When the replication service status changes to **Ready to Deploy**, click **DEPLOY**.

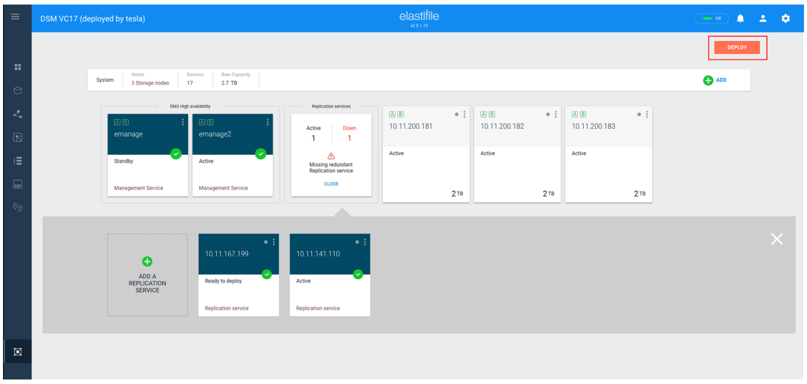

7. In the Deployment Summary window, click **DEPLOY**.

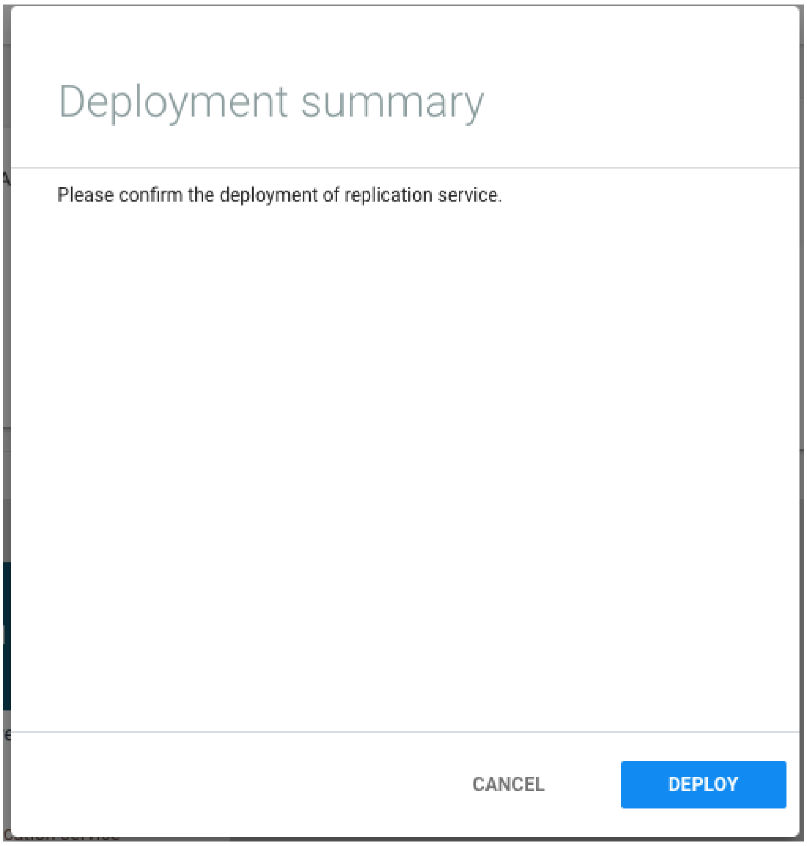
Deployment summary

Please confirm the deployment of replication service.

8. When the replication service has been deployed, do one of the following:

   - Click **GO TO DASHBOARD**.

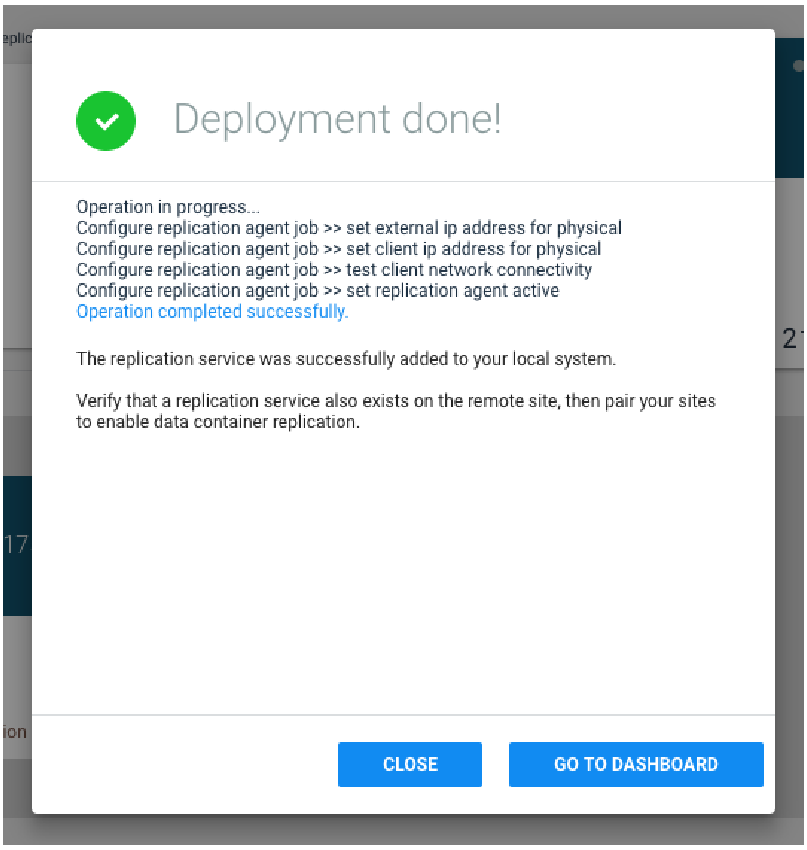

Deployment done!

Operation in progress...
Configure replication agent job >> set external ip address for physical
Configure replication agent job >> set client ip address for physical
Configure replication agent job >> test client network connectivity
Configure replication agent job >> set replication agent active
Operation completed successfully.

The replication service was successfully added to your local system.

Verify that a replication service also exists on the remote site, then pair your sites to enable data container replication.

   - Click **CLOSE** to return to **System view**. The Replication service status initially displays **Init**, then **Active**.

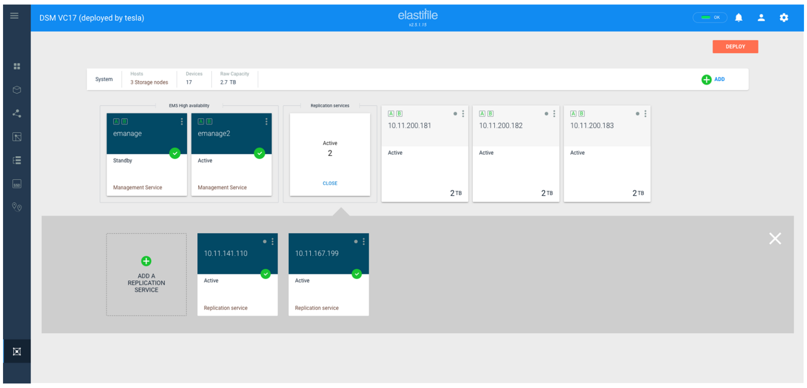

# 7. Managing the ClearTier Object Tier

In the age of data growth, primary storage cost reduction is a major concern. A large portion of the data created by business applications is hardly accessed after a certain time but is needed for analytics purposes or historical fact-checking. Physically moving the data to another cheaper system is cumbersome both for the system admin and more so for the application users who don't know which system to check. Running analytics on multiple systems is nearly impossible.

Elastifile ClearTier moves data automatically to a cheaper object store (either cloud services or on-prem third-party solutions) while keeping the files available via the same NFS mount point.

Frequently used data (hot files) will enjoy the high performance of a primary storage solution, while older files (cold files) will be moved in the background to the object tier. All the files metadata will be stored on the primary storage so the whole namespace will appear as a single file system to the application regardless where the files reside. File updates will be done on the primary storage tier in sub-file chunks and will not require to warm a full file just to update a portion of it.

ClearTier also gives you the option to move a full snapshot from the primary tier to the object tier, in order to free precious space from the primary tier while keeping as many backups as you like. Since each snapshot maintains the full directory structure and files attributes, it can serve as a backup and archive solution, answering business archive requirements. Snapshots on the object tier can be accessed using Elastifile CloudConnect if you need to recover a specific snapshot.

Based on Elastifile CloudConnect technology, ClearTier optimizes object storage by compressing data before moving it to the object tier and by utilizing deduplication to ensure that only unique data is written and existing data is just referenced by a new write command.

## 7.1 Enabling the Object Tier

Use one of the following procedures to enable the object tier as follows:

- Section 7.1.1 - Enabling the Object Tier Automatically - use this option to create an object tier in the same GCP project as the ECFS.
- Section 7.1.2 - Enabling the Object Tier Manually - use this option to create an object tier in a different GCP project to the ECFS, or with a different cloud provider.

### 7.1.1 Enabling the Object Tier Automatically

You can create an object tier automatically if the object tier will be located in the same GCP project as the ECFS.

When you create an object tier automatically, two replication services with only internal networks are created.

If you want to create an object tier automatically but with external networks in addition to the internal ones, first create the replication servers manually as described in Section 7.1.3 - Creating Replication Services for an Object Tier with External Networks, then perform the following procedure.

If you will create an object tier without external networks, you must enable private Google access as a prerequisite for enabling the object tier automatically.

For more information on how to check and configure private Google access, see: https://elastifile.freshdesk.com/support/solutions/articles/42000046543-private-google-access

**To enable the object tier automatically:**

1. In the Assets menu, click **OBJECT TIERS**.

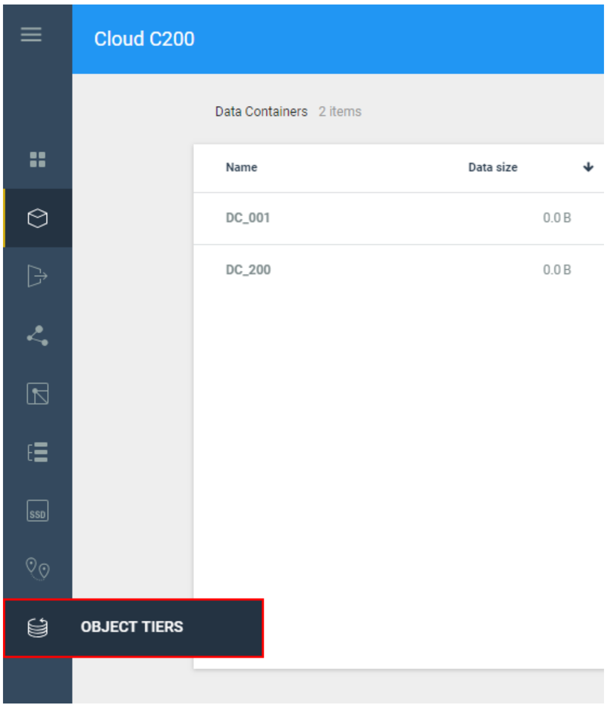

2. Click **Activate Clear Tier on the current project**.

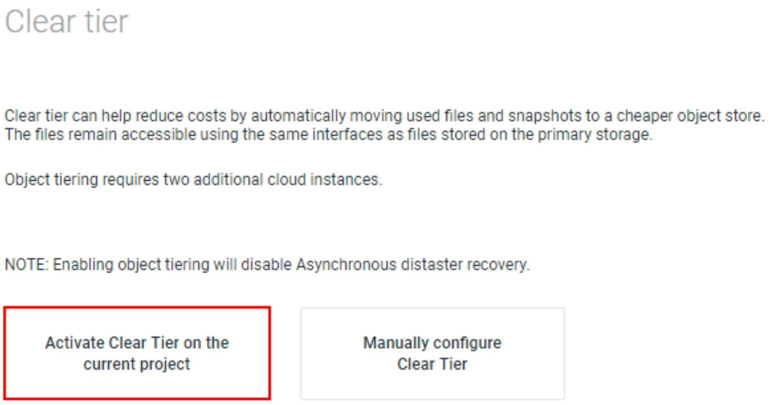
Clear tier

Clear tier can help reduce costs by automatically moving used files and snapshots to a cheaper object store. The files remain accessible using the same interfaces as files stored on the primary storage.

Object tiering requires two additional cloud instances.

NOTE: Enabling object tiering will disable Asynchronous distaster recovery.

Activate Clear Tier on the current project
Manually configure Clear Tier

3. After several minutes, your object tier is created. Click **CLOSE**.

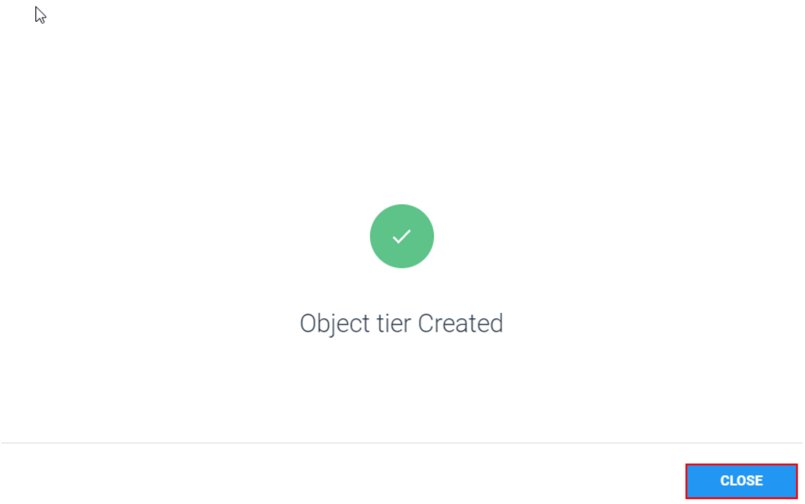

Object tier Created

### 7.1.2 Enabling the Object Tier Manually

You can create an object tier manually if, for example, the object tier will be located in a different GCP project to the ECFS, or with a different cloud provider.

> When you create an object tier manually, replication services with external networks are created.

**To enable the object tier manually:**

1. In the Assets menu, click **OBJECT TIERS**.

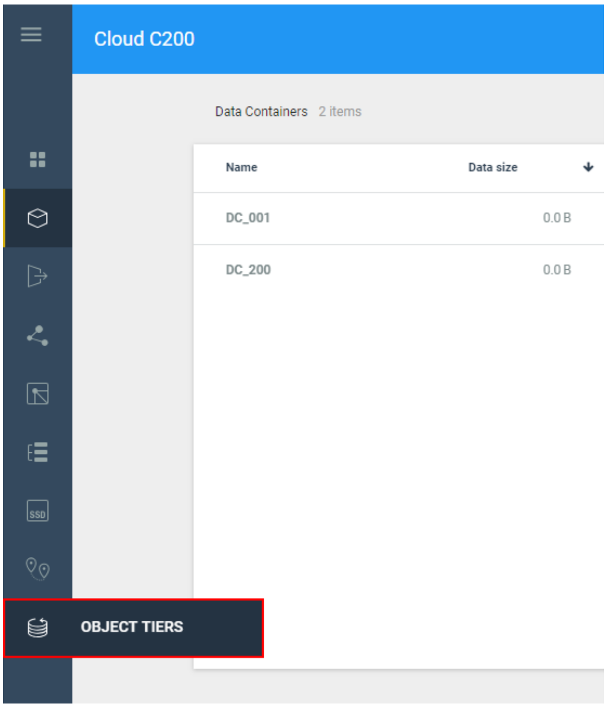

Cloud C200

Data Containers 2 items

| Name | Data size |
|---|---|
| DC_001 | 0.0 B |
| DC_200 | 0.0 B |

2. In the Clear tier window, click **Manually configure Clear Tier**.

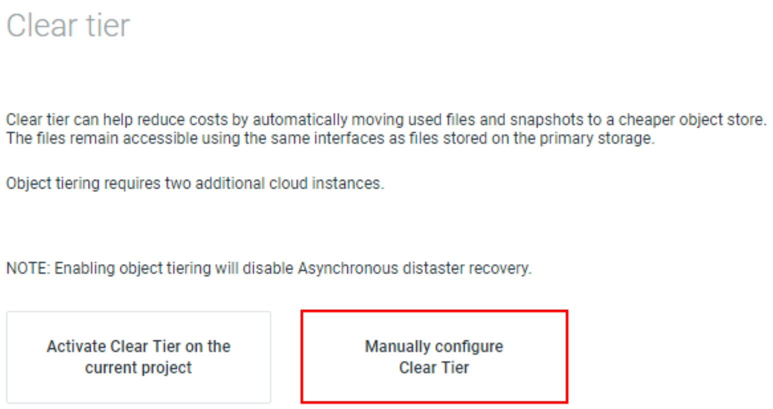

Clear tier

Clear tier can help reduce costs by automatically moving used files and snapshots to a cheaper object store. The files remain accessible using the same interfaces as files stored on the primary storage.

Object tiering requires two additional cloud instances.

NOTE: Enabling object tiering will disable Asynchronous distaster recovery.

Activate Clear Tier on the current project
Manually configure Clear Tier

3. Select one of the following providers:
   - Google
   - Amazon Web Service (S3)
   - Cloudian HyperStore

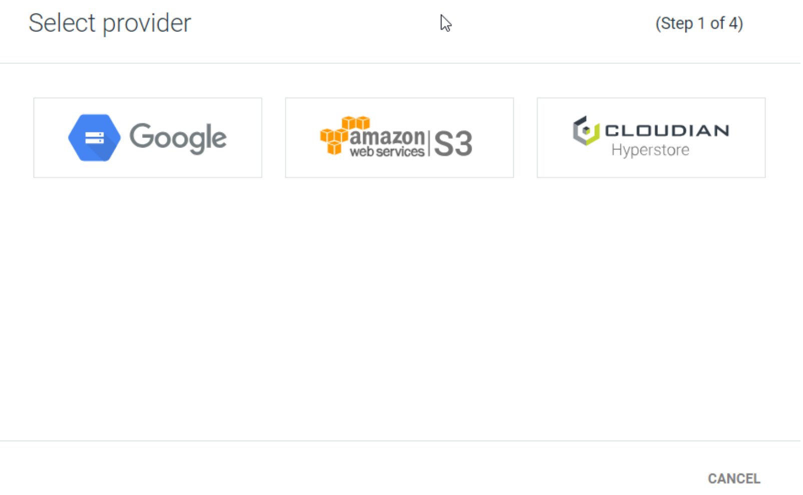

Select provider (Step 1 of 4)

4. In the **Configure project** window, in **Name**, type the name of your cloud project, the region, endpoint and alert level.

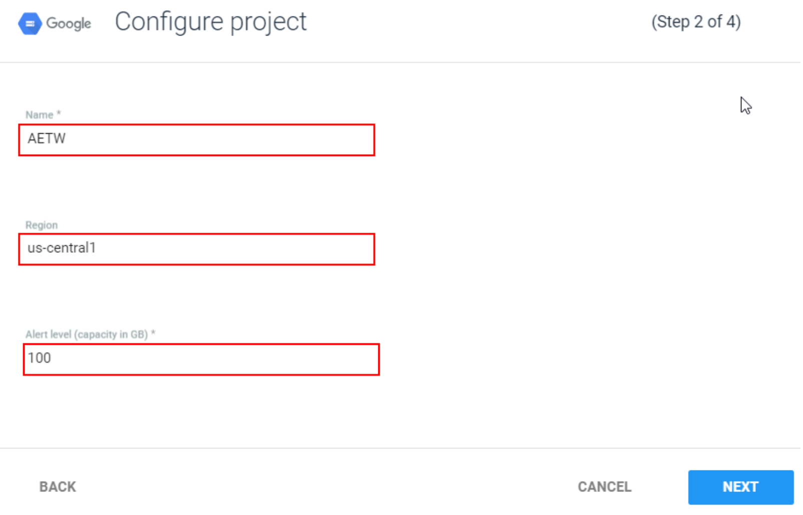

Google Configure project (Step 2 of 4)

Name * AETW

Region us-central1

Alert level (capacity in GB) * 100

- Google and Amazon Web Service (S3) do not require an end point.

5. Click **NEXT**.

6. If your selected provider is:
   - Google:
     - In the Upload JSON file window, drag and drop your credentials JSON file to this window.
   - Amazon Web Service (S3) or Cloudian HyperStore
     a. In **Access Key**, type your access key.
     b. In **Service Key**, type your service key.
     c. Click **SAVE**.

7. After several minutes, your object tier is created. Click **CLOSE**.

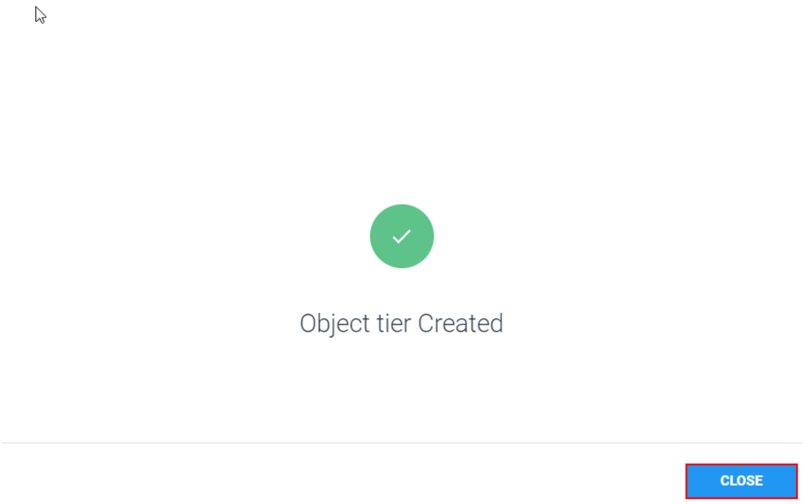
Object tier Created

### 7.1.3 Creating Replication Services for an Object Tier with External Networks

**To create replication services with an external network for an object tier:**

1. In the Assets menu, click ☒ **SYSTEM VIEW**.

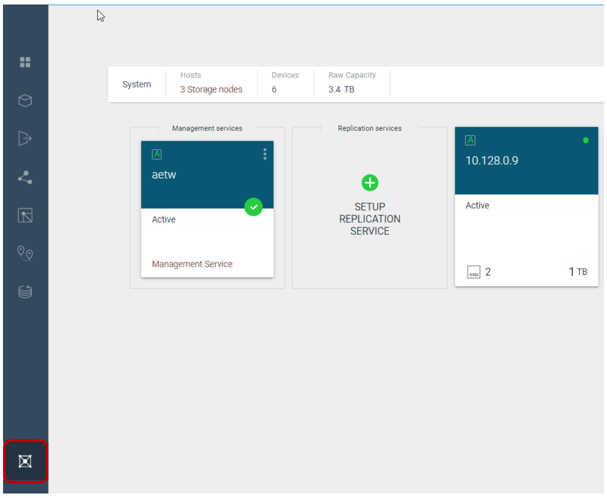
System | Hosts 3 Storage nodes | Devices 6 | Raw Capacity 3.4 TB

Management services: aetw, Active, Management Service

Replication services: SETUP REPLICATION SERVICE

10.128.0.9, Active, 2, 1 TB

2. Click **SETUP REPLICATION SERVICE** and select both check boxes (**Enable High availability** and **Use external network for replication traffic**).
3. Click **ADD**.

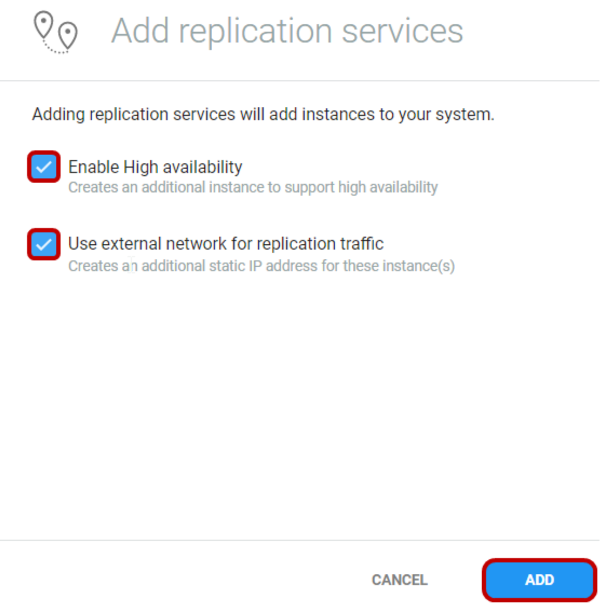

Add replication services

Adding replication services will add instances to your system.

☑ Enable High availability
Creates an additional instance to support high availability

☑ Use external network for replication traffic
Creates an additional static IP address for these instance(s)

4. After the replication services are added, click **CLOSE**.

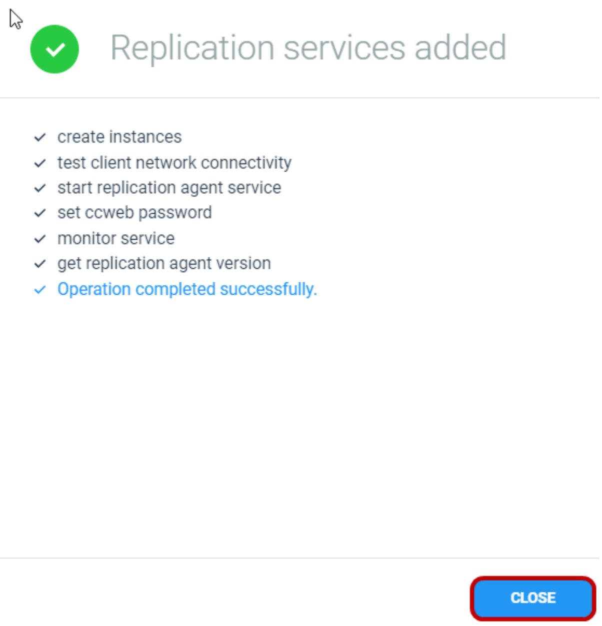

Replication services added

✓ create instances
✓ test client network connectivity
✓ start replication agent service
✓ set ccweb password
✓ monitor service
✓ get replication agent version
✓ Operation completed successfully.

5. Two replication instances with external networks are created. You can now create the object tier as described in Section 7.1.1 - Enabling the Object Tier Automatically.

## 7.2 Managing Your Data Tiering Policy

The **Data tiering** window shows the target data tiering policy and the actual use of primary and object tier storage.

The data tiering policy consists of two conditions:

- the percentage of total storage used by primary storage and the object tier.
- the time period an object must not be accessed for it to be moved to an object tier.

**To modify your data tiering policy:**

1. In the Assets menu, click **OBJECT TIERS**.

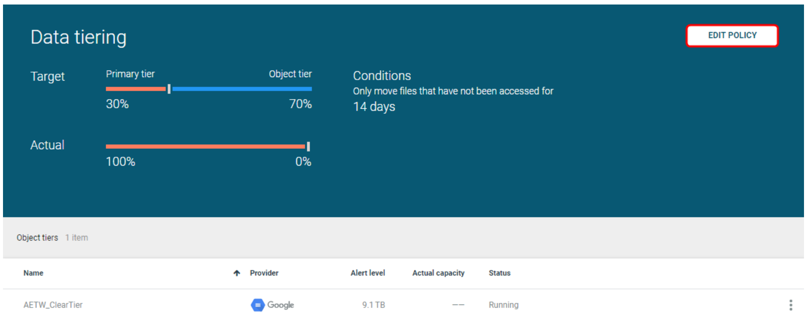

Data tiering

Target: Primary tier 30% / Object tier 70%

Conditions
Only move files that have not been accessed for
14 days

Actual: 100% / 0%

Object tiers 1 item

| Name | Provider | Alert level | Actual capacity | Status |
|---|---|---|---|---|
| AETW_ClearTier | Google | 9.1 TB | –– | Running |

2. Click **Edit Policy**.

3. Modify the following as required:
   - **Set Primary/Object tier ratio** - click any location on the bar or drag the bar slider to the required ratio.
   - **Restriction** - type the number of days that an object was not accessed for it to be moved to an object tier.
   - Click **UPDATE**.

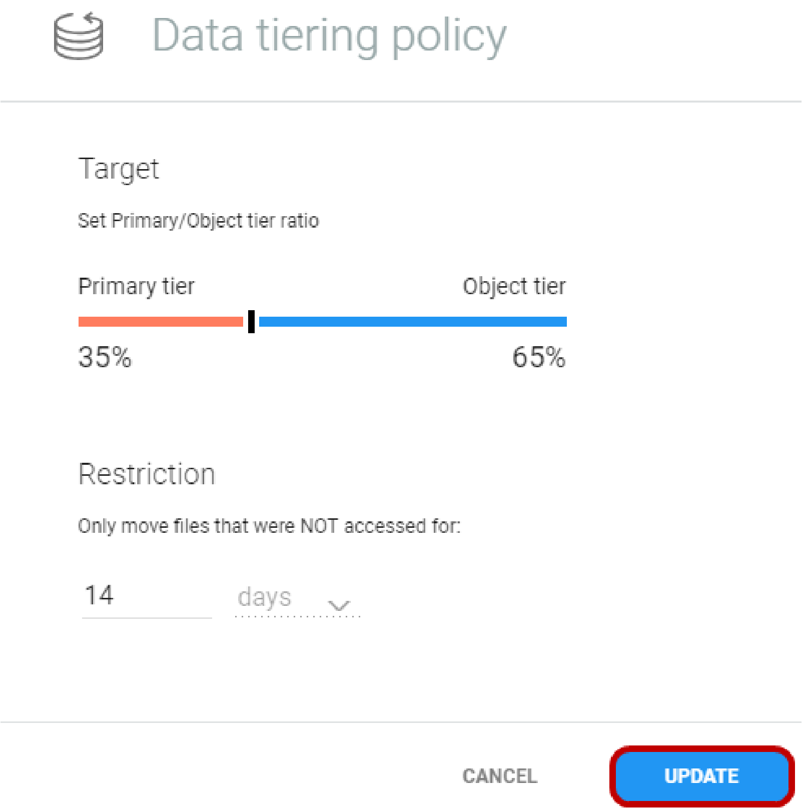

Data tiering policy

Target

Set Primary/Object tier ratio

Primary tier 35% / Object tier 65%

Restriction

Only move files that were NOT accessed for:

14 days

## 7.3 Managing Object Tiers

### 7.3.1 Deleting Object Tiers

You cannot delete an object tier that is used by data containers.

**To delete an object tier:**

1. In the Assets menu, click 🛢 **OBJECT TIERS**.

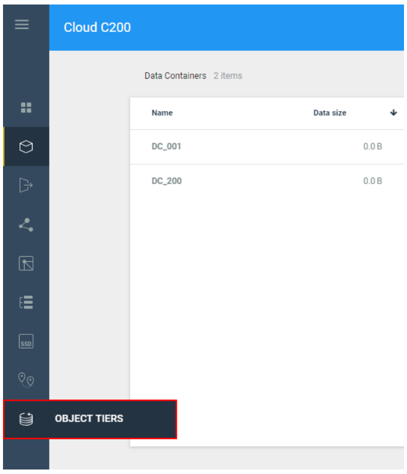

Cloud C200

Data Containers 2 items

| Name | Data size |
|---|---|
| DC_001 | 0.0 B |
| DC_200 | 0.0 B |

2. On the required object tier row, click ⋮ (**Options**) and click **Delete**.

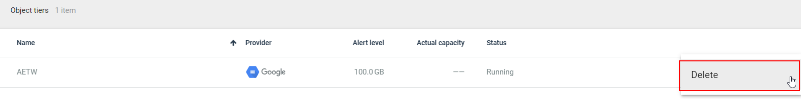

Object tiers 1 item

| Name | Provider | Alert level | Actual capacity | Status |
|---|---|---|---|---|
| AETW | Google | 100.0 GB | — | Running |

# 8. Administering Your System

Click (**ADMINISTRATION**) to open the following system menu items:

- System Settings - reconfigure the system using the installation wizard.
- System Operations - contains various system operation options such as shutting down and restarting a system.
- Log Collection - lists log files and enables you to generate logs on demand.
- REST API Docs - describes the REST API for accessing the ECFS Management System.
- About - displays information about licensing and system properties.

## 8.1 Changing System Settings

You can change system settings using a similar interface to the configuration wizard that you used during system installation.

**To access the System Settings:**

1. In the header, click (**ADMINISTRATION**).
2. Click **System Settings**.
3. Click any of the subcategories and modify the settings as required (see the *ECFS Installation Guide* for your environment).

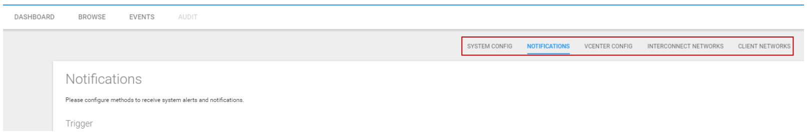

Notifications

Please configure methods to receive system alerts and notifications.

Trigger

## 8.2 Performing System Operations

In the **System Operations** window, you can perform various administration tasks as described below.

**To access System Operations:**

1. In the header, click (**ADMINISTRATION**).
2. Click **System Operations** and select the required operation as described in the following table.

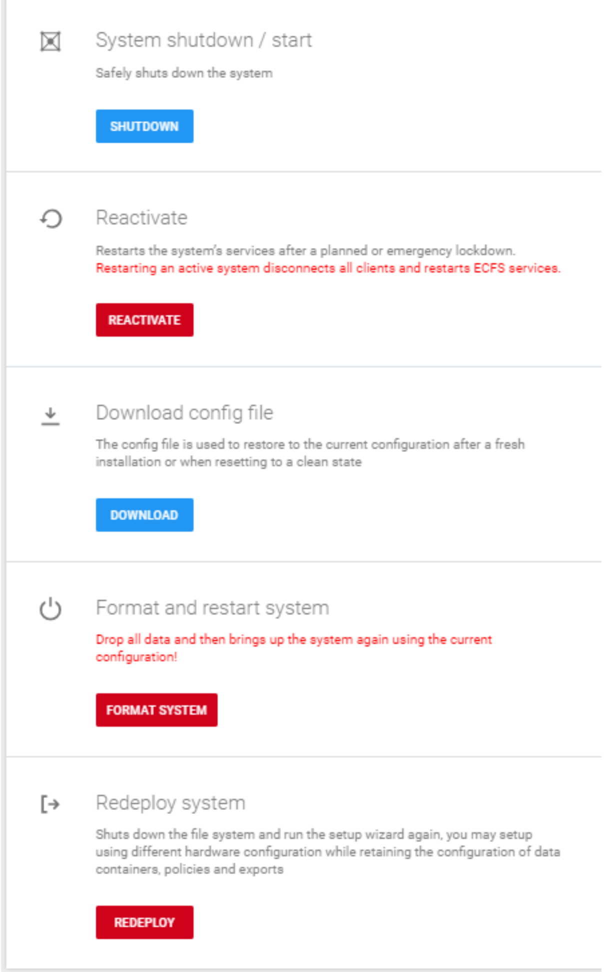

System shutdown / start

Safely shuts down the system

Reactivate

Restarts the system's services after a planned or emergency lockdown.
Restarting an active system disconnects all clients and restarts ECFS services.

Download config file

The config file is used to restore to the current configuration after a fresh installation or when resetting to a clean state

Format and restart system

Drop all data and then brings up the system again using the current configuration!

Redeploy system

Shuts down the file system and run the setup wizard again, you may setup using different hardware configuration while retaining the configuration of data containers, policies and exports

| Item | Description |
|---|---|
| **Shutdown/Start System** | Safely shuts down the system or starts a system that has been shutdown. |
| **Reactivate** | Restarts the system's services after a planned or emergency lockdown. |
| **Download config file** | Downloads a json configuration file you can use to restore to the current configuration after a fresh installation or when resetting to a clean state. |
| **Format and restart system** | Clears all data and then brings up the system again using the current configuration. |
| **Redeploy System** | Shuts down the file system and run the setup wizard again. You can then setup using different hardware configuration. |

## 8.3 Generating and Viewing Logs

**To access logs:**

1. In the header, click ⚙ (**ADMINISTRATION**).
2. Click **Log Collection**. A list of logs generated by you or the system appears.
3. You can download or delete each log by clicking the appropriate icon at the end of the row.

**To generate a logs:**

1. Click **GENERATE LOG**.
2. Select the required log: **System log** or **Single host log**.
3. If you selected **Single host log**:
    a. Click **Select host**. The list of hosts appears.
    b. Click your required host.
4. If you also require trace files, select the **Add trace files (larger file size)** check box.
5. Click **GENERATE**.

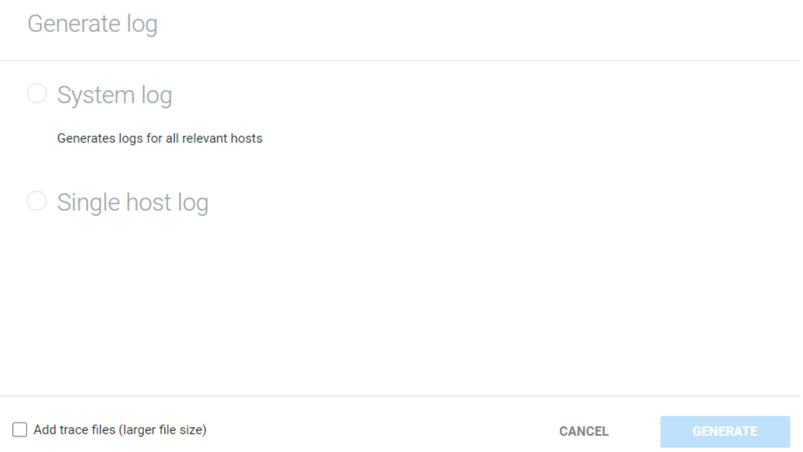

Generate log

○ System log

Generates logs for all relevant hosts

○ Single host log

□ Add trace files (larger file size)

6. The **Generating report** message appears and your log will appear in the log list.

## 8.4 Viewing REST API Documentation

ECFS includes a REST API for accessing the ECFS Management System.

**To access the REST API online documentation:**

1. In the header, click ⚙ (**ADMINISTRATION**).
2. Click **REST API docs**.

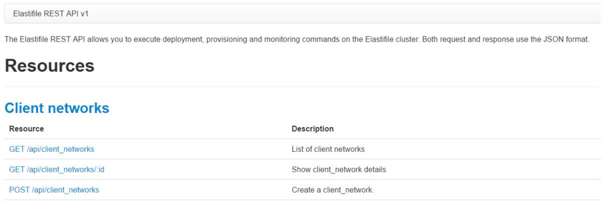

Elastifile REST API v1

The Elastifile REST API allows you to execute deployment, provisioning and monitoring commands on the Elastifile cluster. Both request and response use the JSON format.

Resources

Client networks

| Resource | Description |
|---|---|
| GET /api/client_networks | List of client networks |
| GET /api/client_networks/:id | Show client_network details |
| POST /api/client_networks | Create a client_network. |

To return to ECFS Management Console, click the back arrow in your browser.

## 8.5 Licensing Information and System Properties

**To display licensing information and system properties:**

1. In the header, click (**ADMINISTRATION**).

2. Click **About**.

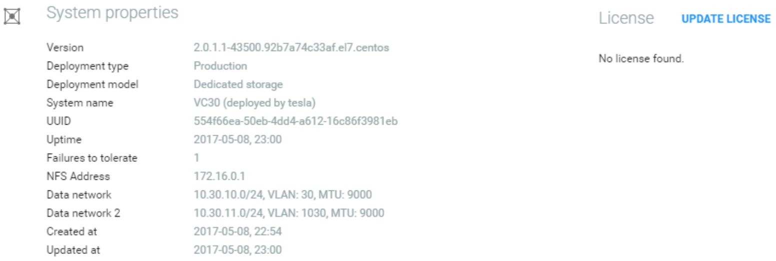

System properties

| | |
|---|---|
| Version | 2.0.1.1-43500.92b7a74c33af.el7.centos |
| Deployment type | Production |
| Deployment model | Dedicated storage |
| System name | VC30 (deployed by tesla) |
| UUID | 554f66ea-50eb-4dd4-a612-16c86f3981eb |
| Uptime | 2017-05-08, 23:00 |
| Failures to tolerate | 1 |
| NFS Address | 172.16.0.1 |
| Data network | 10.30.10.0/24, VLAN: 30, MTU: 9000 |
| Data network 2 | 10.30.11.0/24, VLAN: 1030, MTU: 9000 |
| Created at | 2017-05-08, 22:54 |
| Updated at | 2017-05-08, 23:00 |

License

No license found.

### 8.5.1 Updating the License

**To update the license:**

1. In the header, click (**ADMINISTRATION**).

2. Click **About**, then click **UPDATE LICENSE**.

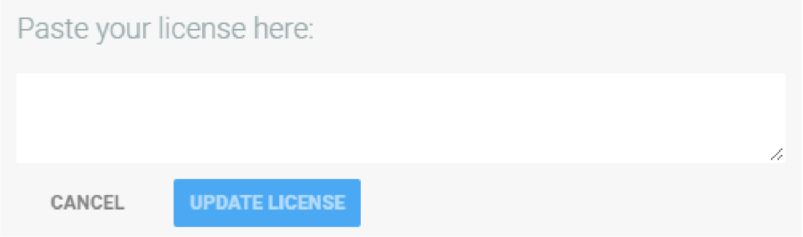

Paste your license here:

3. Copy the contents of the license.txt file you received from Elastifile and paste in the text box.

4. Click **UPDATE LICENSE**.

## 8.6 Upgrading the ECFS System

You can upgrade the ECFS system without disrupting system operations.

To obtain the .tar upgrade file, contact Elastifile.

**To upgrade ECFS:**

1. In the header, click (**ADMINISTRATION**).

2. Click **About**.
3. Drag and drop the .tar file into the designated area, or browse to and select the file.
4. Select the upgrade method as explained in the window.
5. Click **UPGRADE SYSTEM**.

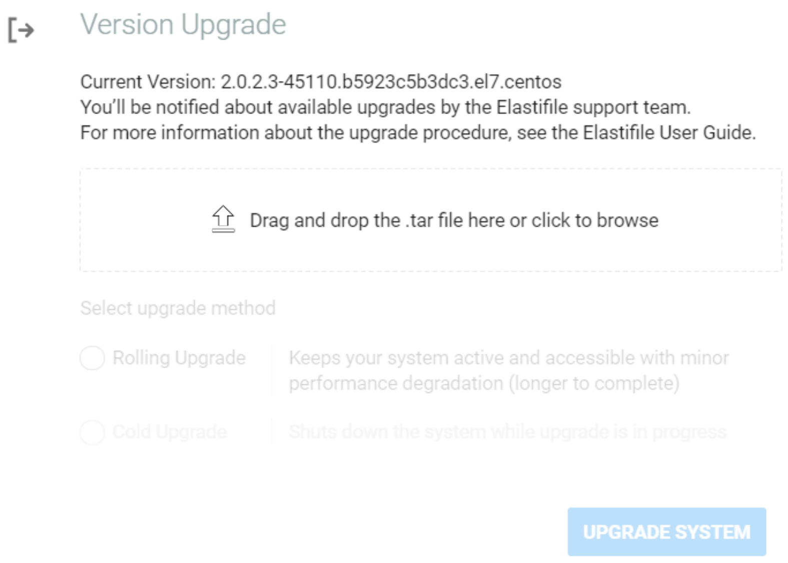

Version Upgrade

Current Version: 2.0.2.3-45110.b5923c5b3dc3.el7.centos
You'll be notified about available upgrades by the Elastifile support team.
For more information about the upgrade procedure, see the Elastifile User Guide.

Drag and drop the .tar file here or click to browse

Select upgrade method

○ Rolling Upgrade | Keeps your system active and accessible with minor performance degradation (longer to complete)

○ Cold Upgrade | Shuts down the system while upgrade is in progress

## 8.7 User Profile

Click 👤 (**USER**) to open the following user menu items:

- Change Password
- Logout

### 8.7.1 Changing Your Password

When you log in to ECFS for the first time, Elastifile recommends that you change your password.

**To update your password:**

1. Click 👤 (**USER**), then click **Change Password**.
2. In the appropriate fields, enter your current password and your new password.
3. Click **SAVE**.

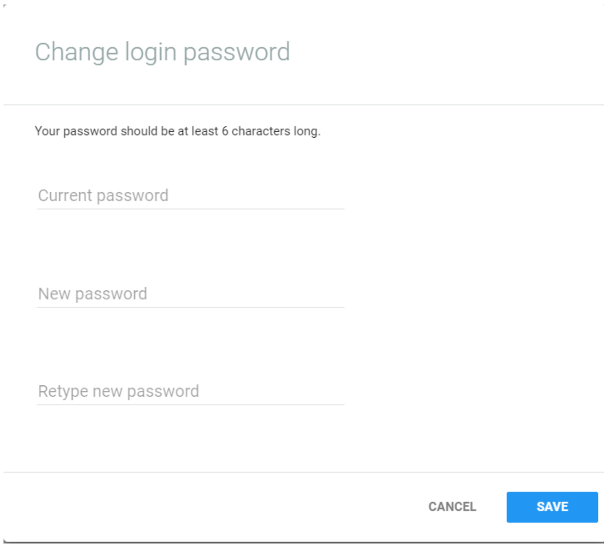

Change login password

Your password should be at least 6 characters long.

Current password

New password

Retype new password

### 8.7.2 Logging Out from ECFS

**To logout from ECFS:**

1. Click 👤 (**USER**), then click **Logout**.

# 9. Troubleshooting

## 9.1 System Lockdown

A system lockdown can occur when one or more ECFS subsystems malfunctions.

**To resolve a system lockdown:**

1. For each faulty host, do the following:
    a. Confirm the faulty host is up and running.
    b. Confirm the faulty host is connected.
2. In System View perform the following:
    a. Confirm faulty hosts' power indicator is green.
    b. Confirm the faulty hosts' data networks indicators are green.
    c. Confirm the Virtual Controller Status displays **Awaiting Reactivation**.
3. Click ⚙ (**ADMINISTRATION**), select System Operations and click **REACTIVATE**.

If the system does not reactivate, contact Elastifile Support.

## 9.2 Resolving Failed Devices

If ECFS detects a device with degraded performance, it disables the device by removing it from the data path and marks it as failed.

ECFS indicates failed devices in several views:

- System view - the number of devices and raw capacity are updated to reflect the failed disk. In addition, the failed disk in the host is red.
  You can hover over the failed disk to see more information on the failure.

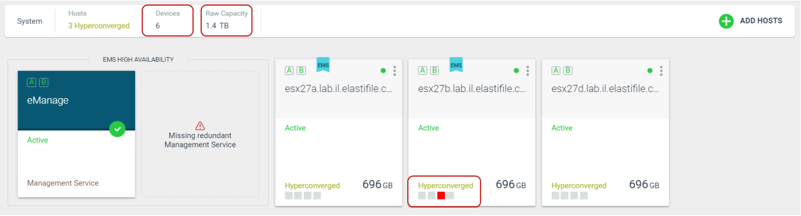

- Host view - the Status column displays **Failed** for any row with a failed device.

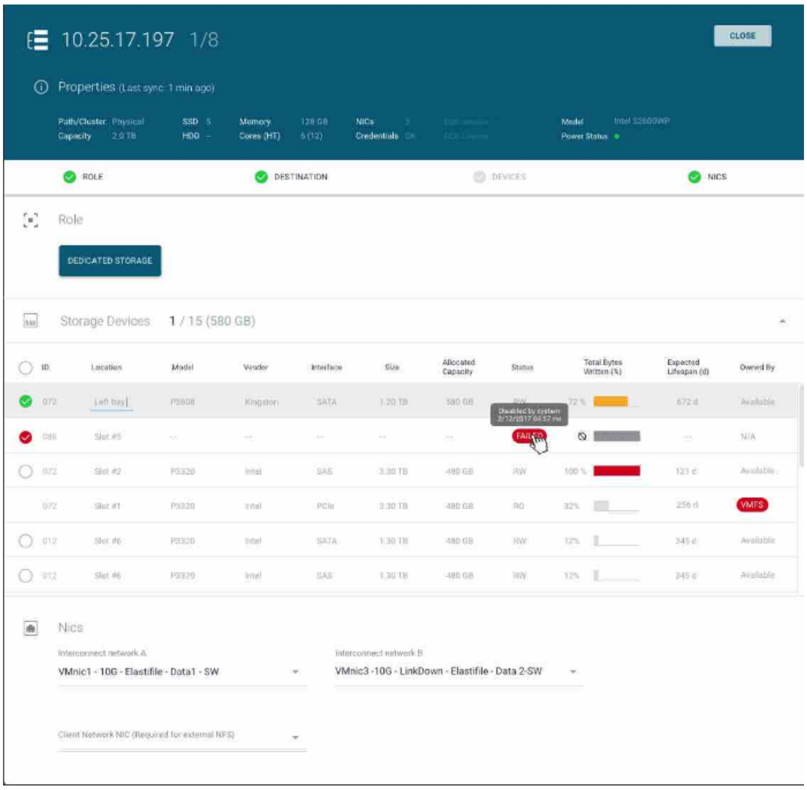

10.25.17.197 1/8

Properties (Last sync: 1 min ago)

Path/Cluster: Physical
Capacity: 2.0 TB
SSD: 5
HDD: –
Memory: 128 GB
Cores (HT): 5 (12)
NICs: 3
Credentials: OK
Model: Intel S2600WP
Power Status: ●

Role

Storage Devices 1 / 15 (580 GB)

| | ID | Location | Model | Vendor | Interface | Size | Allocated Capacity | Status | Total Bytes Written (%) | Expected Lifespan (d) | Owned By |
|---|---|---|---|---|---|---|---|---|---|---|---|
| ✓ | 072 | Left bay | P3608 | Kingston | SATA | 1.20 TB | 580 GB | RW | 72 % | 672 d | Available |
| ✓ | 088 | Slot #5 | -- | -- | -- | -- | -- | FAILED | | -- | N/A |
| ○ | 072 | Slot #2 | P3320 | Intel | SAS | 3.30 TB | 480 GB | RW | 100 % | 121 d | Available |
| | 072 | Slot #1 | P3320 | Intel | PCIe | 3.30 TB | 480 GB | RO | 32% | 256 d | VMFS |
| ○ | 012 | Slot #6 | P3320 | Intel | SATA | 1.30 TB | 480 GB | RW | 12% | 345 d | Available |
| ○ | 012 | Slot #6 | P3320 | Intel | SAS | 1.30 TB | 480 GB | RW | 12% | 345 d | Available |

Disabled by system 2/12/2017 04:57 PM

Nics

Interconnect network A: VMnic1 - 10G - Elastifile - Data1 - SW

Interconnect network B: VMnic3 -10G - LinkDown - Elastifile - Data 2-SW

Client Network NIC (Required for external NFS)

- Device view - the Status column displays **Failed** for any row with a failed device.

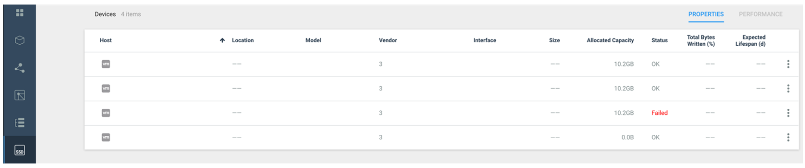

Devices 4 items

| Host | Location | Model | Vendor | Interface | Size | Allocated Capacity | Status | Total Bytes Written (%) | Expected Lifespan (d) |
|---|---|---|---|---|---|---|---|---|---|
| vm | —— | | 3 | | —— | 10.2GB | OK | —— | —— |
| vm | —— | | 3 | | —— | 10.2GB | OK | —— | —— |
| vm | —— | | 3 | | —— | 10.2GB | Failed | —— | —— |
| vm | —— | | 3 | | —— | 0.0B | OK | —— | —— |

**To resolve a failed device:**

1. Click ☒ **SYSTEM VIEW** and locate the host with the failed device.
2. On the host with the faulty device, click ⋮ (**Options**), and click **Remove from system** and follow the instructions.
3. Power off the host and replace the failed device.
4. Power the host up.
5. Click ☒ **SYSTEM VIEW** and click **HOSTS**.
6. Select the required host and follow the on-screen instructions.

## 9.3 Events List

The following table lists possible events displayed in the Events View (see Section 5.6 - Monitoring Events).

| Event ID | Event Name | Class | Severity |
|---|---|---|---|
| 1 | DiskIOError | Disk | Minor |
| 2 | TransientDiskError | Disk | Warning |
| 3 | ClientConnectionDown | Protocols | Info |
| 4 | ClientConnectionEstablished | Protocols | Info |
| 5 | MPoolResourceLimitExceeded | Platform | Major |
| 6 | EThreadsResourceLimitExceeded | Platform | Major |
| 7 | RPCHandlerResourceLimitExceeded | Platform | Major |
| 8 | PStoreConcurrencyResourceLimitExceeded | Disk | Minor |
| 9 | PStoreDirectBucketResourceLimitExceeded | Disk | Minor |
| 11 | ControlSystemStarted | Platform | Info |
| 12 | MountFailed | Protocols | Warning |
| 13 | PANIC | Platform | Major |
| 15 | DiskFormatted | Disk | Info |
| 16 | SlowDevice | Disk | Major |
| 18 | CRCError | Disk | Major |
| 19 | PStoreRotatingLogWaitForIndirectMap | Disk | Minor |
| 1003 | NewMapDeployed | File system | Info |
| 1004 | TransferToMapCompleted | File system | Info |
| 1005 | NodeAdded | Platform | Info |
| 1006 | NodeRemoved | Platform | Info |
| 1007 | DiskAdded | Disk | Info |
| 1008 | DiskRemoved | Disk | Info |
| 1011 | NICAdded | Platform | Info |
| 1012 | NICRemoved | Platform | Info |
| 1013 | ClusterProtocolServiceStopped | Platform | Info |
| 1014 | ClusterMetadataServiceStopped | Platform | Info |
| 1015 | ClusterStorageServiceStopped | Platform | Info |
| 1016 | ClusterProtocolServiceStarted | Platform | Info |
| 1017 | ClusterMetadataServiceStarted | Platform | Info |
| 1018 | ClusterStorageServiceStarted | Platform | Info |
| 1020 | EnodeFencing | Control | Critical |
| 1021 | ClusterLockDown | Control | Critical |
| 1023 | NodeInError | Control | Critical |
| 1025 | RocMapDistributionStarted | File system | Info |
| 1026 | RocMapDistributionFinished | File system | Info |
| 1027 | RocMapDistributionFailed | File system | Info |
| 1028 | OrcMapDistributionStarted | File system | Info |

| Event ID | Event Name | Class | Severity |
|---|---|---|---|
| 1029 | OrcMapDistributionFinished | File system | Info |
| 1030 | OrcMapDistributionFailed | File system | Info |
| 1031 | NodeInPanic | Control | Info |
| 1032 | NodeRestartedAsFEOnly | Control | Info |
| 1033 | EnterROCDegradedMode | File system | Info |
| 1034 | ExitROCDegradedMode | File system | Info |
| 1035 | ProtocolServicesWriteDisabled | File system | Critical |
| 1036 | ProtocolServicesWriteEnabled | File system | Info |
| 1038 | OwnershipRecoveryStarted | File system | Info |
| 1039 | OwnershipRecoveryFinished | File system | Info |
| 1040 | EcdbMapDistributionStarted | Control | Info |
| 1041 | EcdbMapDistributionFinished | Control | Info |
| 1042 | EcdbMapDistributionFailed | Control | Info |
| 1043 | NodeRestarted | Control | Info |
| 1044 | EcoresUnReachable | Platform | Critical |
| 1045 | DiskFailed | Disk | Info |
| 1048 | PMAPoolDepleted | Platform | Warning |
| 1049 | EnterORCDegradedMode | Control | Info |
| 1050 | ExitORCDegradedMode | Control | Info |
| 1051 | EnterECDBDegradedMode | Control | Info |
| 1052 | ExitECDBDegradedMode | Control | Info |
| 1053 | RocMapDistributionNoCapacity | Control | Critical |
| 5005 | DcHardQuotaExceeded | File system | Major |
| 5006 | DcSoftQuotaExceeded | File system | Major |
| 5007 | EnodeAdded | Platform | Info |
| 5008 | EnodeRemoved (REMOVE_ENODE_EVENT) | Platform | Info |
| 5009 | RemoveEnodeStarted (REMOVE_ENODE_STARTED) | Platform | Info |
| 5010 | DcAdded (CREATE_DATA_CONTAINER_EVENT) | File system | Info |
| 5011 | DcRemoved (REMOVE_DATA_CONTAINER_EVENT) | File system | Info |
| 5012 | ExportAdded (CREATE_EXPORT_EVENT) | File system | Info |
| 5013 | ExportRemoved (REMOVE_EXPORT_EVENT) | File system | Info |
| 5014 | ExportAclAdded (CREATE_EXPORT_CLIENT_EVENT) | File system | Info |
| 5015 | ExportAclRemoved (REMOVE_EXPORT_CLIENT_EVENT) | File system | Info |
| 5016 | InterfaceDown (ENODE_NIC_DOWN) | Platform | Info |
| 5018 | SystemLowCapacity (SYSTEM_LOW_CAPACITY) | File system | Major |
| 5019 | VcenterSyncFailed (VCENTER_SYNC_FAILED) | Control | Minor |
| 5020 | EnodeConsoleError (ENODE_CONSOLE_ERROR) | Control | Minor |

| Event ID | Event Name | Class | Severity |
|---|---|---|---|
| 5021 | RemoveEnodeFailed (REMOVE_ENODE_FAILED) | Control | Minor |
| 5022 | SecondaryEmanageJoined (SECONDARY_EMANAGE_JOINED) | Control | Info |
| 5023 | SecondaryEmanageJoinFailed ((SECONDARY_EMANAGE_JOIN_FAILED) | Control | Major |
| 5024 | TenanatSoftQuotaExceeded (TENANT_SOFT_QUOTA_EXCEEDED) | File system | Major |
| 5025 | EmanageActivationStarted (EMANAGE_ACTIVATION_STARTED) | Platform | Info |
| 5026 | EmanageActivationCompleted (EMANAGE_ACTIVATION_FINISHED) | Platform | Info |
| 5027 | SystemReactivationStarted (SYSTEM_REACTIVATION_STARTED) | Platform | Critical |
| 5029 | DcReplicationFailed (DC_REPLICATION_FAILED) | File system | Major |
| 5031 | DcTestImageCreatedSuccessfuly (TEST_IMAGE_FINISHED) | File system | Info |
| 5032 | DcTestImageCreationFailed (TEST_IMAGE_FAILED) | File system | Minor |
| 5034 | ReplicationServiceFailed (REPLICATION_SERVICE_FAILED) | Platform | Major |
| 5037 | ExportEnableFailed (ENABLE_EXPORT_FAILED) | File system | Warning |
| 5038 | DataRebuildError | File system | Critical |
| 5038 | domain_logon_failure | File system | Critical |
| 5039 | REPLICATION_SERVICE_ALL_FAILED | Platform | Critical |