opentext™

S8000 - SIEM Storage Appliance

Appliance Version: 25.3

Administrator's Guide to the S8000 Appliance

Legal Notices

Open Text Corporation

275 Frank Tompa Drive, Waterloo, Ontario, Canada, N2L 0A1

Copyright Notice

The only warranties for products and services of Open Text and its affiliates and licensors ("Open Text") are as may be set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Open Text shall not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.

Trademark Notices

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number
- Document Release Date, which changes each time the document is updated
- Software Release Date, which indicates the release date of this version of the software

To check for recent updates or to verify that you are using the most recent edition of a document, go to:

https://www.microfocus.com/support-and-services/documentation

Support

Contact Information

Phone	A list of phone numbers is available on the Technical Support Page: https://www.microfocus.com/en-us/contact-support/stackb
Support Web Site	https://www.microfocus.com/en-us/support
ArcSight Product Documentation	https://www.microfocus.com/documentation/arcsight/#gsc.tab=0

Contents

About this Guide

This installation guide provides instructions on how to install and initialize an S8000 - SIEM Storage Appliance appliance.

For more information, see "About the S8000 - SIEM Storage Appliance " below.

Intended Audience

This guide provides information for administrators who need to install, initialize, and restore S8000 - SIEM Storage Appliances.

Additional Documentation

This documentation library includes the following resources, based on the product that you use.

ArcSight Platform

- *S8000 - SIEM Storage Appliance Release Notes*, which provide information about the appliance release.
- Deploying a W8300 Appliance in the *Administrator's Guide to the W8300 - SIEM Worker Node Appliance*, which provides instructions on how to use this storage in step 13.

For specific product issues, contact OpenText Customer Care.

About the S8000 - SIEM Storage Appliance

The S8000 - SIEM Storage Appliance provides on-premises object storage for the DB8400 - SIEM Database Appliance in your SIEM Platform environment. The S8000 - SIEM Storage Appliance is a scalable storage platform based on the most stable version of the Ceph storage system, with a Ceph management platform, deployment utilities, and support services.

> An S8000 - SIEM Storage Appliance deployment must include a minimum of four S8000 appliances set up as a cluster.

Chapter 1: Setting Up an S8000 - SIEM Storage Appliance

This section describes how to set up your S8000 - SIEM Storage Appliance.

These basic steps enable you to start using your S8000 - SIEM Storage Appliance.

	Task	See
❏	1. Power on theS8000 - SIEM Storage Appliance	"Powering On the S8000 - SIEM Storage Appliance" below
❏	2. Set up Remote Access	"Setting Up the Appliance for Remote Access" below
❏	3. First Boot of the Appliance	"First Boot Initialization of the S8000 - SIEM Storage Appliance (Bootstrapping)" on page 7
❏	4. Deploying an S8000 - SIEM Storage Appliance	"Setting up the DB8400 - SIEM Database Appliance to Use the S8000 - SIEM Storage Appliance Cluster for Storage" on page 13

Powering On the S8000 - SIEM Storage Appliance

To power on the appliance:

1. Unpack the appliance and its accompanying accessories.

   > **Note:** Read carefully through the instructions, cautions, and warnings that are packed with the appliance shipment. Failing to do so can result in bodily injury or appliance malfunction.

2. Follow the rack installation instructions to securely mount it to a suitable rack.
3. Make the front and rear panel connections.
4. Activate the power switch.

Setting Up the Appliance for Remote Access

All appliances are equipped with an `iDRAC Service Module (iSM)` for remote access. OpenText strongly recommends setting up and configuring your appliance for out-of-band remote access. Doing so ensures that you or Customer Support (with your permission and assistance) can remotely access the console of your appliance for troubleshooting, maintenance, and control over the powering on and off of the box.

Changing the iDRAC password on your Appliance

Appliance boxes come with a random iDRAC password. For information on how to locate the password, see Secure Default Password.

This is a unique password, which will be required the first time iDRAC is accessed. The appliance then will prompt for a new password to be chosen. For security reasons, OpenText recommends to change this password as soon as possible.

To set up your appliance for remote access, follow the instructions in the EMC iDRAC Service Module.

Encryption of SEDs

The S8000 Appliances support FIPS enabled self-encrypting disks (SEDs).

A SED is a data storage device with built-in cryptographic processing to encrypt and decrypt the data it contains. This process occurs within the device itself, independent of any connected information system, and it provides data protection against the loss or theft of the disks, as well as certain levels of hacking attempts.

This protection consists of setting up passphrase-access-only.

The SEDs ship without the passphrase, allowing you to chose your own. To set up a passphrase, first follow the steps to establish a security key.

The chosen passphrase can then be applied to pre-existing virtual disks by following the steps in Secure a pre-existing virtual disk.

To change or disable a security key, please follow the specific procedures listed under this section.

First Boot Initialization of the S8000 - SIEM Storage Appliance (Bootstrapping)

> **Tip:** Be aware that this process will require network information for the appliance, such as:
>
> - Static `IP` address
> - Resolvable FQDN hostname
> - NTP server that's both accessible and running
>
> All of this information must be available to successfully complete the bootstrapping.

Perform this operation on each appliance to initialize it on first boot.

1. Log into your appliance using `iDRAC` (see "Setting Up the Appliance for Remote Access" on page 5 for instructions), and launch the Virtual Console.
2. Turn on the appliance using the Power Controls option, in case the appliance is off.
3. Using the local drive (`NVMe`), select from the menu the version of `Red Hat` you want to boot from.
4. From the console, log in using your default username (`otadmin`) and password (`change_me`).
5. You will be required to change the password for `otadmin`.Enter a new password and retype it to confirm.

```
You are required to change your password immediately (administrator
enforced).
Current password:
New password:
Retype new password:
```

> **Note:** The STIG-compliant password policy rules for both the `arcsight` and the `root` password require:
>
> - A minimum of 15 characters
> - A minimum of 1 number
> - A minimum of 1 lowercase character
> - A minimum of 1 uppercase character
> - A minimum of 1 special character
> - A maximum of 2 consecutive repeating characters
> - A maximum of 4 consecutive repeating characters of the same class
> - A minimum of 8 different characters
> - To not be a word from the dictionary
> - To be different from the last seven passwords

6. The **OpenText Appliance** splash screen will display, with the **User must set 'root' password to proceed** message. You will be required to enter the `otadmin` user password you just reset to make the change to the `root` password:

```
password for otadmin
Changing password for user root
New password:
Retype new password:
passwd: all authentication tokens updated successfully
```

> Once your passwords have been set, you will need to wait for at least one day to update to a different one. And the maximum expiration period for a password is 60 days.

7.

```
========================================
User must set 'arcsight' password to proceed:
=========================================
```

8. Complete the **Network Configuration**. The screen will display a list of network interfaces and their status:

```
********************************************************************
Network Configuration
**********************************************************************
WARNING: You must specify static IP address and resolvable hostname
(FQDN).
**********************************************************************
List of network interfaces
**********************************************************************
enoxxxxnp0      UP             xx:xx:xx:xx:xx:xx       <BROADCAST, MULTICAST, UP, LOWER_
enoxxxx         DOWN            xx:xx:xx:xx:xx:xx       <NO-CARRIER, BROADCAST, , MULTI
ensxxxx         DOWN            xx:xx:xx:xx:xx:xx       <NO-CARRIER, BROADCAST, MULTICA
```

```
**********************************************************************
Select one active connection to configure:
**********************************************************************
1) enoxxxxnp0
#? 1
```

Select the active connection that you want to configure for the public network. You can configure the secondary networks using the OS network utility later.

> **Note:** You must deploy the S8000 storage cluster in the same protected subnet as the DB8400 - SIEM Database Appliance. For more information about securing your network, see " Security Hardening" on page 27.

9. Configure the network using a static IP address (FQDN) by providing this information:

```
**********************************************************************
Configure the network connection enoxxxxnp0 for device enoxxxxnp0:
**********************************************************************
Enter the hostname (FQDN) for this appliance: your_appliance_host_fqdn

Configure network using static IP address:
Enter static IPv4 address:
Enter IPv4 prefix (1-32):
Enter IPv4 gateway:
Enter IPv4 Primary DNS server:
Enter IPv4 Secondary DNS server (optional):
Enter spaced separated IPv4 DNS search domains: your_appliance_domain
```

10. Next, the NTP server must be configured:

```
**********************************************************************
NTP Server Configuration
**********************************************************************
WARNING: You must specify an accessible NTP server

Enter the NPT server for this appliance:
```

The console will display a summary of the network configuration and NTP server configuration, and will ask you to verify by entering Y:

```
Do you want to configure network settings and NTP services using above
configuration? (Y/N)
```

If you need to correct the information, enter N, and the process will ask you for each item again. If you enter Y, the process will continue:

11.
```
The node is S8000.
IMPORTANT: SIEM Storage Appliance installation can only be launched on
primary node.
Is this the primary node for SIEM Storage Appliance installation? (y/n):
```

    Answer y to only one node in the cluster that you designate as the primary node. For the other nodes, answer n. If you answer y, the process will continue:

```
Continuing setup primary node...
Generate self-signed certificate and first time login token...
```

12. If the configuration ends successfully, you will see the following message:

```
**********************************************************************
The appliance network and NTP server have been setup successfully
**********************************************************************
Go to https://<your_appliance_host_fqdn>:6443 to install S8000 product
IMPORTANT: You will need the token to login for the first time:
xxxxxxxx
**********************************************************************
```

> The console will not allow you to copy the token, which you will need for your first login to the S8000 administration console. Access the URL provided above in your browser, and type the token manually as shown in the console.

Regeneration of the First Login Token

In case there's a need to obtain a First Login Token again (other than with the preceding procedure), you can regenerate it by running the following command in the console:

```
# /var/opt/arcsight/appliance_scripts/generate_first_login_token.sh
```

You will be prompted for the `arcsight` password, and after providing it, the First Login Token will be generated again:

```
=================================================================
Go to https://<your_appliance_host_fqdn>:6443 to install S8000 product
IMPORTANT: You will need the token to login for the first time:
xxxxxxxx
=================================================================
```

Deploying the S8000 - SIEM Storage Appliance

This section will describe how to set up and deploy a S8000 - SIEM Storage Appliance.

The S8000 storage cluster and the DB8400 - SIEM Database Appliance must be deployed in the same protected subnet. Plan your cluster size for the EPS you are targeting, with a 10% buffer for possible growth. OpenText **strongly recommends** that you deploy all nodes at the same time. To deploy a *new* node after your initial deployment, contact OpenText support.

Prerequisite

For optimal cluster performance, you must obtain a DNS load balancer host name that resolves to all the nodes' IP addresses in a round-robin fashion. Instead of a DNS load balancer, you can use your own load balancer such as nginx in front of the S8000 storage cluster that distributes requests in a round-robin fashion on port 9000 to individual S8000 nodes. If you are deploying and monitoring the S8000 storage cluster from a machine outside of the protected subnet, you must expose ports 6443 (for the appliance installer web app), 8443 (for the Ceph dashboard), and 3000 (for Grafana) to the external internet.

To deploy a S8000 - SIEM Storage Appliance:

1. Log in to the appliance using your ArcSight-supplied organizational credentials (username and password) and then click **Log in.**
2. On the **OpenText EULA** page, review the OpenText End User License Agreement. Select **I have read and agree with the end user license agreement** and then click **Next**.
3. On the **RedHat EULA** page, review the RedHat End User License Agreement. Select **I have read and agree with the end user license agreement** and then click **Next**.
4. On the **Cluster Setup** screen, click the + icon.
5. To add nodes to your cluster, under **Add Nodes**, click **Add Node**.
6. In the **Add Node** dialog box, in **Application Hostname/IP address**, enter the FQDN or the network path to the new node.
7. In **ArcSight OS user password**, enter the password to the node (your ArcSight password), then click **Save**. The node is added to your list of nodes.
8. Repeat steps 6 to 9 until you have added all of your nodes.
9. On the **Appliance Config** page, enter values for the following and then click **Next**.
   a. **Ceph dashboard admin username**: Enter a username for the Ceph dashboard administrator account.
   b. **Ceph dashboard admin password**: Enter a password for the Ceph dashboard administrator.
   c. **Confirm Password**: Retype the password to confirm.
   d. **Bucket name**: Enter the name of the S3 bucket that is used to store DB8400 - SIEM Database Appliance data. The bucket name must conform to the following rules:

- Must contain only lowercase letters, numbers, or dashes
- Must start and end with a letter or number
- Must be 3 to 63 characters in length
- Cannot contain uppercase letters, underscores, periods, or special characters
- Cannot be formatted like an IP address (for example, 192.168.0.1)

e. **Erasure coding data chunks (K)**: Using the default value is strongly recommended. See the Ceph documentation on erasure to understand more before you change this value.

f. **Erasure coding coding chunks (M)**: Using the default value is strongly recommended. See the Ceph documentation on erasure to understand more before you change this value.

g. **Access host**: Specify the host name of the load balancer to access the cluster. You can contact your system administrator to configure a host name on a DNS server that resolves to all nodes in the cluster (DNS loadbalancing), or you can use your own load balancer such as nginx and configure it to route requests to all nodes in the S8000 deployment cluster in a round robin format.

h. **S3 Port**: Enter the port number on which S3 service must be exposed.

10. On the **Review Summary** page, review and verify the summary information. If the summary is correct, click **Submit**. If the information is incorrect, click **Previous** to return to the previous page containing the incorrect information. Then correct the information, return to this page, and click **Submit**.

> **Note:** You should review the summary information carefully: After you click Submit, you cannot make any changes. To make changes, you will need to re-image the appliance.

11. The appliance configuration proceeds. If the installation fails or succeeds, a message is displayed.

a. If the installation is successful, the setup is confirmed and your S3 server information is displayed.

b. If the installation fails, an error message is displayed. Click Retry to submit the information. Alternatively, click Exit to exit the wizard and review the error logs, and then re-run the setup wizard (return to step 1, above). If you get repeated errors, contact OpenText customer support.

12. Click **Exit Setup**.

Setting up the DB8400 - SIEM Database Appliance to Use the S8000 - SIEM Storage Appliance Cluster for Storage

After you deploy the S8000 appliance cluster, you can use it as communal storage for the DB8400 - SIEM Database Appliance.

Copy the following information from the final deployment screen:

- S3 Server
- S3 Server Port
- S3 Bucket URL
- Access key
- Access secret
- Click **Download Storage Appliance CA Certificate** to download the S3 server's CA certificate to your computer.

For instructions to use the information above, see step 13 of "Deploying a W8300 Appliance" in the Administrator's Guide to the W8300 - SIEM Worker Node Appliance.

Restoring an Appliance to Factory Settings

You can restore appliances to their original factory settings by using the procedures detailed here. To perform a restore procedure, you will require:

- An `.iso` image file containing the factory settings for the version of S8000 - SIEM Storage Appliance you are restoring. Find the name of the file in the **Downloading the ArcSight Platform Installation Files** section of the H.

  > ✓ Once you have acquired the image file, please refer to the signature verification instructions, and perform the verification steps before starting the procedure below

The restore procedure can be conducted in two ways:

- If you have physical access to the appliance, use the "Restoring an Appliance Using a USB Memory Stick" on the next page method
- If you have only iDRAC access to the appliance, use the "Restoring an Appliance Using iDRAC Access" on page 16 method

Multinode Clusters

In a multinode cluster configured for high-availability, the HA quorum may be spoiled by bringing down a single node for restoration. In this case, all nodes in the cluster will need to be restored to their factory settings. For example, in a 3-node highly-available S8000 - SIEM Storage Appliance cluster, to restore a single S8000 - SIEM Storage Appliance node to factory settings, the complete process for each node would include:

1. Remove the impacted node from Kubernetes.
2. Restore it to factory settings as described above.
3. Run the bootstrap installer.
4. Add the node back to the cluster.

Restoring an Appliance Using a USB Memory Stick

This method will require the following external hardware:

- A 32 GB or higher USB memory stick (the faster type available, but at least USB 2.0 or 3.x)
- A Linux machine to perform the burning of the .iso image into the USB memory stick

Image Burning

1. Connect the USB memory stick to one of the ports of the Linux machine.
2. From the command line, execute the following command to burn the `.iso` image into the USB memory stick:

   ```
   dd if=<iso_image_file_name>.iso status=progress oflag=sync of=/dev/sdb
   bs=1M
   ```

   Where `<iso_image_file_name>` is the name of the image file downloaded here.

   And wait until the progress has reached 100%.
3. Turn your appliance off and connect the bootable USB stick you just created to one of its ports. Reboot the appliance.

Restore Procedure:

1. Access the remote console of the appliance through `iDRAC`.

   If you already used the remote access, use the password you setup the first time you connected. Otherwise, for instructions see:

"Setting Up the Appliance for Remote Access" on page 5

2. From the `iDRAC` **Dashboard**, select the **Virtual Console** on the right lower corner.
3. Click the **BOOT** button on the upper right hand corner and select the **BIOS Boot Manager** option.

   A pop-up window will request to **Confirm Boot Action**, setting a new device to boot from. Select **Yes**.
4. The previous step will not initiate the reboot automatically. For that, you will need to click the **POWER** button, and from the **Power Control** pop-up window, choose the **Reset System (warm boot) option**.

   A pop-up window will request to **Confirm Power Action**. Select **Yes**.
5. The booting process will prompt a selection from the **Boot Manager**. Choose **One-shot UEFI Boot Menu**.
6. From the **Select UEFI Boot Option**, select your USB stick (its name will depend on brand and model, but it will start with **Disk connected to back USB**).
7. The appliance will boot from the selected USB stick.

   The restore process will start automatically if you allow it some time, or you can click on the **ArcSight User Image SIEM-R7615-S8000-RH96-FIPS-STIG-CIS2-25.3-1.iso** option at the top to start right away.

   Twice during this process you will receive a warning about all the data in the partition or hard disk being overwritten. You must enter Y to proceed both times:

   ```
   Are you sure you want to continue? (y/n)
   ```
8. Different screens will follow each other, some of them with progress bars, indicating the restoring progress of a specific system portion. None of these require user intervention, and the whole process takes approximately 10 minutes. Once the restore process has reached this point:

   ```
   realtime =none
   The next step: true
   Now run: true
   ```

   Your input will be required to reboot the appliance:

   ```
   reboot
   ```
9. Once the reboot process is finished, follow the instructions listed in:

   Connecting to the S8000 Appliance

Restoring an Appliance Using iDRAC Access

⚠ When using the iDRAC Remote File Share feature to perform the restore procedure, make sure there is no USB drive connected to the appliance ports, since its presence may interfere with the restore process.

This method will require the following preparation:

- Store your .iso image in a location that is accessible to the iDRAC network. For more information, see the iDRAC documentation.
- Configure the iDRAC Remote File Share option in the Virtual Media tab using shared the .iso image downloaded here.

Restore Procedure:

1. Access the remote console of the appliance through iDRAC.

   If you already used the remote access, use the password you setup the first time you connected. Otherwise, for instructions see:

   "Setting Up the Appliance for Remote Access" on page 5
2. From the iDRAC **Dashboard**, select the **Virtual Console** on the right lower corner.
3. Click the **BOOT** button on the upper right hand corner and select the **BIOS Boot Manager** option.

   A pop-up window will request to **Confirm Boot Action**, setting a new device to boot from. Select **Yes**.
4. The previous step will not initiate the reboot automatically. For that, you will need to click the **POWER** button, and from the **Power Control** pop-up window, choose the **Reset System (warm boot) option**.

   A pop-up window will request to **Confirm Power Action**. Select **Yes**.
5. The booting process will prompt a selection from the **Boot Manager**. Choose **One-shot UEFI Boot Menu**.
6. From the **Select UEFI Boot Option**, select **Virtual Optical Drive**.
7. The appliance will boot from the .iso image in the Remote File Share.

   The restore process will start automatically if you allow it some time, or you can click on the **ArcSight User Image SIEM-R7615-S8000-RH96-FIPS-STIG-CIS2-25.3-1.iso** option at the top to start right away.

   Twice during this process you will receive a warning about all the data in the partition or hard disk being overwritten. You must enter Y to proceed both times:

```
Are you sure you want to continue? (y/n)
```

8. Different screens will follow each other, some of them with progress bars, indicating the restoring progress of a specific system portion. None of these require user intervention, and the whole process takes approximately 10 minutes. Once the restore process has reached this point:

   ```
   realtime =none
   The next step: true
   Now run: true
   ```

   Your input will be required to reboot the appliance:

   ```
   reboot
   ```

9. Once the reboot process is finished, follow the instructions listed in:

   Connecting to the S8000 Appliance

Chapter 2: Managing the S8000 - SIEM Storage Appliance

Note: Contact OpenText customer support to perform other management tasks on your S8000 cluster, such as

- Use certificates signed by your Certificate Authority. Note that before you install a Vertica database, you must replace the certificates issued by your CA.
- Add a new S8000 node to the cluster.

Undestanding Ceph

The S8000 - SIEM Storage Appliance is based on Ceph 19.2.2. Ceph is a scalable, open-source, software-defined storage platform. Ceph Storage clusters consist of the following services.

Ceph OSD

The Ceph OSD is the object storage daemon for the Ceph distributed file system. Its primary function is to store data objects on physical storage devices (hard drives).

- OSDs handle the actual reading and writing of data.
- OSDs perform data replication, erasure coding, and recovery operations as directed by the Ceph cluster.
- Each OSD manages a single physical storage drive.
- OSDs are the foundation of Ceph's distributed object storage architecture.

Ceph monitor (CEPH-MON)

Ceph monitors maintain the cluster's map, which contains essential information about the cluster's state, including the location of all OSDs, the state of the cluster, and the cluster configuration.

- Ceph monitors form a distributed consensus using the Paxos algorithm to ensure data consistency.
- They provide a central source of truth for the cluster's state.
- Clients and other Ceph daemons query the monitors for cluster information.

Ceph manager (CEPH-MGR)

The Ceph manager provides monitoring and management capabilities for the Ceph cluster.

- The Ceph manager collects and exposes cluster metrics and statistics.
- It hosts various modules that extend Ceph's functionality, such as the dashboard, Prometheus exporter, and REST API.
- It removes the burden of monitoring and stat collection from the monitors.

Ceph metadata server (CEPH-MDS)

The Ceph MDS is responsible for managing the metadata of files stored in the Ceph File System (CephFS).

- It tracks the directory structure, file permissions, and other metadata associated with files.
- It allows clients to navigate and access files in CephFS.
- It does not handle the file's data, only its metadata. The file's data is handled by the OSD's.

Ceph object gateway (CEPH-RGW)

The Ceph object gateway (RGW) provides an object storage interface compatible with Amazon S3 and OpenStack Swift APIs.

- THe Ceph object gateway allows applications to store and retrieve data using standard object storage protocols.
- It supports features such as buckets, objects, and access control.
- It translates S3 and Swift API calls into Ceph native storage operations.

Monitoring the S8000 - SIEM Storage Appliance Cluster with Grafana

The S8000 - SIEM Storage Appliance provides the following dashboards that you can use to monitor the health of the cluster. You can use these dashboards to effectively monitor the health, performance, and usage patterns of your Ceph cluster from various perspectives.

Ceph - Cluster

The Ceph cluster dashboard provides a high-level, immediate overview of the entire Ceph cluster's health, capacity, and basic performance. This is often the first dashboard you check to see if any issues are impacting the cluster's performance.

Usage and insights: You use the Ceph cluster dashboard to quickly assess if the cluster is healthy, check available space, monitor overall load, and see if any core components (OSDs, MONs, PGs) are in a problematic state.

Key widgets and metrics

- Cluster Status: A prominent display showing overall health (for example, HEALTH_OK, HEALTH_WARN, HEALTH_ERR)
- Capacity Usage: Gauges or bar charts showing total and used storage capacity (raw capacity, used capacity, available capacity)
- PG Status: Counts of Placement Groups (PGs) in various states (for example, active and clean, degraded, undersized, recovering, backfilling). A high number of non-active+clean PGs indicates potential issues or ongoing recovery.
- OSD Status: Counts of OSDs that are up or down and in or out of the cluster
- MON Status: Number of Monitors in quorum
- Overall IOPS: Cluster-wide read/write operations per second
- Overall Throughput: Cluster-wide read/write bandwidth (for example, MB/s)
- Client IOPS/Throughput: Aggregated I/O from clients, including RBD, Ceph file system (CephFS), and RGW

Ceph Cluster - Advanced

The advanced cluster dashboard provides deeper insight into cluster-wide performance metrics, recovery operations, and internal Ceph processes. You can use it to diagnose performance issues or understanding background activity.

Usage and insights: Diagnose performance bottlenecks related to OSD latency, monitor the progress and impact of recovery/backfill operations, check for scrubbing issues, and investigate MON health more closely.

Key widgets and metrics:

- Detailed PG States: More granular breakdown of PG states and ongoing recovery operations (recovery or backfill throughput, objects recovered or sec)
- OSD Commit/Apply Latency: Histograms or graphs showing the distribution of latencies for OSD operations, helping identify storage bottlenecks
- Scrub Status: Information on ongoing scrubs, errors found during scrubbing
- Network Latency: (If configured) Latency metrics between cluster nodes
- MON Performance: Quorum status details, clock skew information
- Recovery or Rebalance Metrics: Throughput and progress of data movement due to OSD failures, additions, or re-weighting

Ceph Pool Details

The Ceph pool details dashboard focuses on the metrics of a single, specific Ceph pool, selected using a drop-down menu. Essential for understanding the usage and performance of data belonging to a particular application or user group that uses that pool.

Usage and insights: Understand how much capacity a specific application (using the pool) consumes, analyze its performance characteristics (IOPS, throughput, latency), and troubleshoot issues related specifically to that pool's data or PGs.

Key widgets and metrics

- Pool Usage: Capacity used (bytes, objects) within the selected pool, percentage of cluster capacity used by this pool
- Pool I/O: Read/write IOPS and Throughput specifically for this pool

- Pool Latency: Average Read/write latency for operations targeting this pool
- PG States for Pool: Counts of PGs belonging to this pool in various states
- Pool Configuration: Display of key pool settings (replication size, min_size, crush rule, application type)

Ceph Pools Overview

THe pools overview provides a comparative view of usage and performance across all Ceph pools in the cluster.

Usage and Insights: Identify the largest, fastest-growing, or busiest pools. Compare performance characteristics between different pools. Useful for capacity planning and resource allocation across different services using Ceph.

Key widgets and metrics:

- Pools Table: A table listing all pools with columns for Used Capacity, Max Available, Objects, IOPS (Read/Write), Throughput (read/write), average latency
- Comparative Graphs: Stacked bar charts or line graphs showing capacity usage, IOPS, or throughput contributions from different pools over tim.

Host Details

The host details show detailed system resource metrics for a single, specific host (server) within the Ceph cluster, selected via a dropdown menu. It is useful to diagnose issues that are tied to a specific piece of hardware.

Usage and insights: Determine if a specific host is overloaded (CPU, RAM, network, disk I/O). Identify hardware bottlenecks that might be impacting Ceph performance locally on that node. Correlate host-level issues with Ceph component behavior running on it.

Key widgets and metrics:

- CPU Usage: Overall CPU utilization, per-core usage, load average for the selected host
- Memory Usage: Total, used, free, cached, swap memory for the host
- Network I/O: Network bandwidth (bytes/sec) and packet rate (packets/sec) for relevant network interfaces on the host (e.g., public/cluster networks)

- Disk I/O: Read/write IOPS, throughput, latency, and utilization (%) for physical disks on the host, especially those backing OSDs
- Running Ceph Daemons: Indication of which Ceph services (OSD, MON, MDS, RGW) are running on this host

Host Overview

THe host overview provides a comparative view of key resource metrics across all hosts participating in the Ceph cluster.

Usage and insights: Quickly identify which hosts are under the most stress (high CPU, memory, or network load). Spot systemic issues affecting multiple hosts. Ensure load is reasonably balanced across the hardware.

Key widgets and metrics:

- Hosts Table: A table listing all hosts with key metrics like CPU Usage (%), Memory Usage (%), Network Throughput, potentially Disk Latency averages.
- Heatmaps and graphs: Visualizations showing distribution of CPU load, memory usage, or network traffic across all hosts, making outliers easy to spot.

MDS Performance

MDS performance monitors health of the Ceph Metadata Servers (MDS), which are critical components for CephFS.

Usage and insights: Diagnose CephFS performance issues (slow listings, metadata operations). Check if MDS daemons are overloaded or if cache tuning might be needed. Monitor MDS failover and cluster scaling related to metadata load. Note: Primarily relevant if you use CephFS.

Key widgets and metrics:

- MDS Load: Metrics indicating how busy the active MDS daemons are.
- MDS Cache Performance: Cache hit rates (dentry, inode caches), cache size. High hit rates are desirable.
- Client Requests: Rate of different metadata operations handled by the MDS (for example, getattr, lookup, open, setattr, readdir).
- MDS Map Status: Information about active or standby MDS daemons and their states.

- Session Count: Number of connected CephFS clients.

OSD device details

OSD device details provide in-depth performance and internal metrics for a single, specific OSD, usually selected via a dropdown. Crucial for diagnosing issues with individual storage devices.

Usage and insights: Identify slow or failing individual disks/OSDs. Diagnose performance issues originating from a specific OSD. Understand the internal behavior and efficiency of the storage backend (BlueStore/RocksDB) for that OSD.

Key widgets and metrics:

- OSD I/O Performance: Read/write IOPS, throughput, and latency specifically for this OSD
- Disk Latency: Underlying block device latency (commit/apply latency) associated with this OSD
- BlueStore/RocksDB Metrics: (If using BlueStore) Internal metrics like cache hit rates, compaction statistics, write amplification details
- OSD Resource Usage: CPU and Memory consumed by the specific OSD process
- PG Count: Number of Placement Groups hosted by this OSD
- Disk Utilization: Percentage fullness of the underlying device used by the OSD

OSD Overview

The OSD overview gives a comparative view of status, usage, and performance across all OSDs in the cluster.

Usage and insights: Quickly identify problematic OSDs (down, out, full, slow). Check for imbalances in data distribution (usage %) or load (PG count, latency). Monitor the overall health and performance of the storage layer.

Key widgets and metrics:

- OSD Status Map and Table: Visualization showing which OSDs are up or down, in or out. Table often includes usage percentage, PG count per OSD.
- OSD Usage Distribution: Histogram or heatmap showing how full each OSD is, helping to spot imbalances.

- OSD Latency Distribution: Histogram or heatmap showing commit/apply latency across all OSDs, highlighting slow outliers.
- PG Distribution: Visualization showing the number of PGs per OSD, checking for balance.
- Total OSD IOPS/Throughput: Aggregated performance across all OSDs.

RGW Instance Detail

The RGW instance detail focuses on the performance and operational metrics of a single, specific RadosGW (Ceph Object Gateway) instance or daemon, selected via a drop-down menu. RGW provides S3/Swift compatible object storage.

Usage and insights: Troubleshoot issues with a specific RGW endpoint or server. Analyze the load and performance characteristics of one gateway. Identify error patterns (e.g., high 5xx errors) specific to an instance. The RGW instance detail is useful if you use RadosGW.

Key widgets and metrics:

- HTTP Requests: Rate of GET, PUT, POST, DELETE, HEAD, OPTIONS requests handled by this instance.
- HTTP Status Codes: Counts of successful (2xx), client error (4xx), and server error (5xx) responses served by this instance.
- Request Latency: Average latency for requests processed by this RGW instance.
- Bandwidth: Data transferred (sent/received) by this instance.
- Cache Performance: (If RGW caching is enabled) Hit rates and usage of the RGW cache.
- Resource Usage: CPU/Memory consumed by the specific RGW process.

RGW Overview

Purpose: Provides an aggregated view of performance and operations across all RadosGW instances in the cluster.

Usage and insights: Understand the overall load on your object storage service. Monitor the general health and performance of the RGW layer. Identify global error trends or performance degradation affecting the S3/Swift service. The RGW overview is useful if you use RadosGW.

Key widgets and metrics:

- Total RGW Requests: Cluster-wide sum of GET, PUT, DELETE, etc., requests per second
- Total RGW Bandwidth: Cluster-wide RGW data transfer rate
- Average RGW Latency: Average request latency across all RGW instances
- Aggregated HTTP Status Codes: Cluster-wide counts of 2xx, 4xx, 5xx responses
- Active RGW Instances: Count of currently running RGW daemons
- Top Users/Buckets: Aggregated statistics that show the heaviest users or buckets by requests or bandwidth

RGW S3 Analytics

The RGW S3 analytics provides deeper insights into how the RadosGW S3 service is being used, often focusing on user activity, bucket statistics, and request types. This might require specific RGW configuration for logging or metrics collection.

Usage and insights: Capacity planning based on bucket growth. Identifying which users or applications are driving the most load. Understanding common access patterns (for example, read-heavy over write-heavy). Potentially useful for chargeback/showback reporting. Note: Availability and detail depend on RGW configuration.

Key widgets and metrics:

- Bucket Analytics: Top N buckets by size, object count, requests, or bandwidth usage
- User Analytics: Top N users by requests, bandwidth usage
- Request Type Breakdown: Pie chart or graph showing the proportion of GET, PUT, LIST, DELETE, etc., operations
- Bandwidth Usage Trends: Historical graph of bandwidth consumed via RGW
- Object Size Distribution: (Less common) Histogram showing the distribution of object sizes stored

RGW Sync Overview

Teh RGW sync overview monitors the status, performance, and health of RGW multi-site replication, if configured. This allows data to be synchronized between different Ceph clusters (zones).

Usage and insights: Ensure data is being successfully replicated to disaster recovery or geographically distributed sites. Diagnose delays or failures in multi-site sync. Monitor the bandwidth consumption and performance of the replication process. Note: Only relevant if RGW multi-site replication is set up.

Key widgets and metrics:

- Sync Status per Zone/Bucket: Indicators showing if sync is up-to-date, lagging, or encountering errors for different replication targets
- Replication Lag: Time difference between data modification on the primary site and its replication to a secondary site
- Backlog Size: Amount of data (bytes or objects) waiting to be synchronized
- Sync Throughput: Rate at which data is being successfully replicated (for example, MB/s or objects/sec)
- Sync Errors: Counts or logs of failed replication attempts

Security Hardening

You should keep the following best practices in mind for security hardening of your cluster.

Security Updates

You should have a plan to regularly test and deploy security updates.

Product Updates

OpenText recommends running product updates as they become available.

Access Management

Access management includes authentication, authorization, and accounting. When you grantsystem access to users, apply the principle of least privilege, and only grant trusted users access to the storage cluster system.

Certificate Management

OpenText recommends that organizations deploy certificates that are signed by their Certificate Authority (CA) instead of using self-signed certificates that come with the S8000 - SIEM Storage Appliance.

Security Zones

A security zone includes users, applications, servers, or networks that share common trust requirements and expectations within a system. This guide refers to three distinct security zones that are required to deploy a security-hardened Red Hat Ceph Storage cluster. These security zones are listed below from least to most trusted.

Public Security Zone

The public security zone is an entirely untrusted area of your infrastructure. It can refer to the Internet as a whole or simply to networks that are external to your S8000 - SIEM Storage Appliance deployment that you do not have no authority over. Any data with confidentiality or integrity requirements that traverse this zone should be protected using compensating controls such as encryption. The public security zone **should not** be confused with the Ceph Storage Cluster’s front- or client-side network, which is referred to as the public_network and is usually **not** part of the public security zone or the Ceph client security zone.

Ceph Client Security Zone

The Ceph client security zone refers to networks that access Ceph clients, such as the DB8400 - SIEM Database Appliance accessing the S8000 Ceph cluster over S3. You should put the Ceph client security zone (ArcSight Database) behind a firewall that separates it from the public security zone.

Storage Access Security Zone

The storage access security zone refers to internal networks that provide Ceph clients with access to the Ceph Storage Cluster. The Storage Access Security Zone must provide access to the following ports on the S8000 node for Ceph services to run. If you need to access the Ceph and Grafana dashboards from a workstation outside of theS8000 cluster subnet, expose ports 8443 and 3000 outside the subnet.

Port	Service	Description
9283	ceph-mgr	Prometheus metrics exporter endpoint
8765	ceph-mgr	Used by telemetry or local RESTAPI.

Port	Service	Description
3300	ceph-mon	Ceph Messenger v2 (default MON communication)
6789	ceph-mon	Ceph Messenger v1 (legacy MON communication)
9926	ceph-exporter	Metrics exporter for Prometheus
5000	conmon	Private registry for ceph container images
9100	node-exporter	Host metrics exporter for Prometheus
9093/9094	alertmanager	Prometheus Alertmanager ports
3000	Grafana	Grafana monitoring dashboard
9000/9001	RGW	RGW S3 endpoint instances
9095	Prometheus	Prometheus instances
8443	ceph-mgr	Ceph administration dashboard

You can create additional users in Grafana if needed for monitoring.OpenText strongly recommends that organizations have a secure access management policy and follow the principle of least privilege when they grant access. Only trusted users should have access to the Ceph administrator user and S3 user (access key/secret) credentials.

Publication Status

Released: July, 2025

Updated: Wednesday, July 16, 2025