

IBM Cloud Object Storage System™
Version 3.13.6

Manager Administration Guide



This edition applies to IBM Cloud Object Storage System™ Version 3.12.3 and is valid until replaced by new editions.

© **Copyright IBM Corporation 2017, 2018.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Document Information v

Chapter 1. Overview 1

Browser compatibility	1
Manager Web Interface navigation	1
Functional tabs	1
System links	1
Navigation tree and search	2

Chapter 2. First-time setup. 3

Configuring a system	3
Configuring a new system.	3

Chapter 3. Configuration 5

System access	5
Approve registered devices	5
Approving a single device.	5
Approving multiple devices	5
Cabinets	6
Importing a cabinet configuration	6
Exporting a cabinet configuration	7
Creating a cabinet	7
Configuring a cabinet with drag-and-drop	7
Access pools	8
Creating an Access Pool	8
Editing an Access Pool	9
Editing a Device.	11
Moving an Accesser Node	11
Configuring HTTPS certificates for Access Pools	12
Deleting an Access Pool	12
Storage pools.	13
Creating a storage pool	13
Editing a storage pool	14
Monitoring storage capacity	14
Adding storage capacity	15
Expanding a storage pool.	15
Changing the storage pool reallocation rate.	16
Pausing a storage pool reallocation	16
Resizing a storage pool	17
Replacing storage pool sets	17
Removing storage pool sets	18
Merging two storage pools	18
Delete a storage pool	20
Replace and evacuate the data from Slicestor devices	20
Replacing a Slicestor Device	21
Pause/resume data evacuation	22
Terminating data evacuation.	22
Rolling back a data evacuation	22
Changing the destination device	22
Change the rate of evacuation during evacuation	23
Troubleshooting data evacuation incidents	23
SMC pools	24
Creating an SMC Pool.	24
Editing an SMC Pool	25

Adding an SMC device to an SMC Pool.	25
Monitoring an SMC device	25
Editing an SMC device	26
Removing an SMC device from an SMC pool	26
Deleting an SMC Pool	26
Vaults	27
Overview	27
Management vaults.	28
Standard vaults	30
Container vaults.	30
Service vaults.	31
Configuring vault protection.	31
Creating vaults	31
Create vaults by using vault templates	35
Vault proxy settings	36
Vault security.	36
Tags	38
Vault data migration	39
Vault mirrors	40
Overview	40
Creating a vault mirror	41
Creating mirrors by using mirror templates.	42
Deploying a vault mirror	43
Setting vault mirror permissions	43
Editing a vault mirror	44
Deleting a vault mirror	44
Monitor a vault mirror	45
Creating a protected vault mirror	45
Creating protected mirrors by using mirror templates	46
Repairing a protected mirror	47

Chapter 4. Security 51

Security overview	51
Roles	51
Authentication and authorization	52
Creating an account	52
Editing an account	54
Deleting an account	56
Creating a group	56
Editing a group	57
Deleting a group	57
Granting access key and password authentication	58

Chapter 5. Administration 61

Overview	61
Backing up and restoring the Manager	61
Configure Manager backups.	61
Setting the retention policy	61
Set backup configuration parameters	62
Back up to a local directory	62
Backing up to a secure file transfer protocol	62
Backing up the Manager manually	63
Restoring Manager data	63
Monitor the Manager backup progress	64

Retrieving a Manager backup from a remote server	64
Configuring SMTP	65
Configuring Call Home	65
Configure alerts	66
Configuring email alert rules	66
Configuring alert forwarding	67
SNMP trap details	67
Suppressing device events	68
Configure active directory / LDAP	68
Configuring keystone authentication	69
Configure provisioning API	70
Configuring certificate authority	71
Adding a certificate authority	71
Edit a certificate authority	72
Delete a certificate authority	72
Revoke a certificate authority	72
Supply external certificates to IBM Cloud Object Storage System devices	73
Enable TLS	74
Configure drive health.	75
Configure organization	76
Configure SNMP	76
Configure Device Level API	76
Configure preferences	76
Configure custom login banner	77
Configuring system properties	77
Enabling Cross Site Request Forgery protection	79
Configure advanced system settings	80
Configuring SSH keys	80
Set license keys	81
Configuring notifications	81
Editing or disabling a notification configuration	82
Deleting a notification configuration	82
Chapter 6. Monitoring the system	85
Monitored components	85
Device health summary	85
Monitor device health	86
Monitor storage pool health	87
Monitor vault health	87
Monitor site health	88
Monitor drive health	89
Drive lifecycle states	89
Drive lifecycle state descriptions and troubleshooting	91
ONLINE	91
INIT.	91
OFFLINE	91
DIAGNOSTIC	91
MIGRATING	94
UNUSABLE	95

FOREIGN	95
UNKNOWN	95
RAID states	96
Event console.	96
Open incidents	96
Events	96
Event search	97
Audit search	99
Export	99
Performance graphs	100
Annotated graphs	107
Disabling events on a device	107
Device summary	108
Filtering devices by using tabs	108
Filtering devices by using extra filters	108
Apply filters to device list results.	108

Chapter 7. Maintenance	111
Overview.	111
Upgrade	111
Constraints to upgrades	115
Migrating devices to IPv6	116
IBM Cloud Object Storage Insight	117
Configuring IBM Cloud Object Storage Insight	119
Manually starting IBM Cloud Object Storage Insight sessions.	120
Viewing an anonymized object	120
Logs	120
Collecting logs manually	120
Collecting logs automatically	122
Redacting client IP addresses	123
Collecting log status	124
Configuring device logs	124
Troubleshooting console	125
Changing the device local password.	128
Automatic report emailing	128
Reporting	129
Generating disk drive and device reports	130
Generating storage pool capacity and disk report	131
Generating a vault usage report	132
Vault summary report	133
Reporting device summary.	134
Failed FRU report	134
Reporting event information	135
Firmware report	136
Redaction status report	136

Notices	137
Trademarks	139
Homologation statement	139

Document Information

Intended Purpose and Audience

This *IBM Cloud Object Storage Manager™ Administration Guide* describes how to use the Manager Web Interface to configure, secure, monitor, maintain, and administer the system. The audience is those responsible for data storage administration. This guide applies to both the Manager device and the Manager application.

CAUTION:

If the devices have been installed and a more recent version of the software is available, the devices should be upgraded to the newer version. Contact IBM® Customer Support for assistance with upgrades.

Chapter 1. Overview

Browser compatibility

The Manager Web Interface supports the following browsers:

- Microsoft Internet Explorer® 11 through current
- Microsoft Edge®
- Firefox 4.0 through current
- Chrome versions through current
- Safari versions through current

Attention:

Microsoft ended support of IE9 and IE10 on 1/12/2016. Only the most current Microsoft browser versions will be supported: IE11 and Microsoft Edge, or use the most recent Firefox, Chrome, or Safari browsers. In general, users of older browser versions should upgrade to the most recent version. IE browsers require special settings when used for upgrade. See “Upgrade” on page 111 for details.

Note: In systems with many vaults, the performance of Chrome on MacOS might be slower than Firefox or Safari on MacOS.

If your browser is not listed, most functions work. Many problems can be resolved by turning on your browser's compatibility function.

Manager Web Interface navigation

The Manager Web Interface is structured to facilitate optimum setup, configuration, monitoring, and administration of the system.

There are three main components for navigating the Manager Web Interface:

- “Functional tabs”
- “System links”
- “Navigation tree and search” on page 2

Functional tabs

Functions are organized in tabs: **Monitor**, **Configure**, **Security**, **Maintenance**, and **Administration**.

The Manager Web Interface function can be mapped to Information Technology Infrastructure Library (ITIL) and Fault, Configuration, Accounting or Administration, Performance and Security management (FCAPS) processes and workflow:

Table 1. Manager Web Interface function mapping

	Monitor	Configure	Security	Administration	Maintenance
ITIL	Run	Build/Plan	Manage	Manage	Build
FCAPS	FM/PM	CM	SM	AM	CM

System links

The upper right corner has links to account information, help, and log out.

My account

Account information can be accessed by clicking the account name in that area.

- Click **Change** to change the name that is displayed on the account, email address, and the time zone.
- Click **Change Password** to change the password for the current account.

Note: The default account, admin, is assigned super user access to all Manager Web Interface functions. It allows the admin account to complete all initial setup without requiring extra accounts.

Help

The **Help** link provides help specific to the current page.

The **Help Index** provides links to all product documentation.

Navigation tree and search

The navigation tree provides search and browse functions and can be found on the left side of the Manager Web Interface.

Navigation tree

Click **Open All** to expand or **Close All** to collapse the list of devices.

Note: If fewer than 13 Vaults or Vault Mirrors, or 250 Devices, exist, the navigation tree displays a list of those components in the navigation tree. If 13 or more Vaults or Vault Mirrors, or 250 Devices, exist, the tree is replaced with a link that goes to a separate page as described in “Vaults” on page 27. To clarify this change, a pop-up window explains the change.

Search functions

Autocomplete and standard search functions are available. When a suggestion is selected, the associated page is displayed.

For standard search, click **Go**. The results are provided on the right.

Autocomplete cannot be used for the following items:

Table 2. Items that cannot be auto-completed.

Page	Items
Monitor or Configure	<ul style="list-style-type: none">• site abbreviation• device IP addresses• device type• device alias• device serial number
Security	<ul style="list-style-type: none">• account user name• email

Tip: Standard search can be used for these items.

Chapter 2. First-time setup

Configuring a system

When all devices are physically installed and individually configured, further configuration must be done at the system level.

Before you begin

Open the Manager Web Interface in a web browser, by using `https://{Manager_IP}` where {Manager_IP} is the IP address that is created during the installation of the Manager device.

When first logging in to the Manager Web Interface, the user needs to accept IBM standard and non-IBM End User License Agreements (EULA). The user is needs to complete the Print Name (License Acceptor) field and check the appropriate box and click **Accept IBM & non-IBM Licenses** to accept the EULA and then continue to configure the System. Declining the agreement takes the user to the **Decline** page.

After you accept the EULA, you can create a new system or restore a previous setup.

Procedure

1. Click either **Create new** or **Restore this Manager Web Interface**.
2. Click the **Begin** link.

Note: For more information on restoring a previously created system, see “Restoring Manager data” on page 63.

Configuring a new system

Before you begin

Note: If the initial session ends before the setup is complete, the next session will begin at the next step.

Procedure

1. Enter and confirm a new password for the admin user name.
2. Click **Save and Continue** to continue the initial setup process.
3. Enter the name of the first site at the prompt. The default is My Site. More site information can be added but is not needed.
4. Click + **Add Additional Site** to add sites. It is possible to create, modify, and delete sites after this step, but one site will be created by default.

Note: If the Slicestor[®] devices are to be deployed to more than one site, create the sites even if the devices are initially staged in one location. By defining the sites and assigning Slicestor devices to those sites, data is written equally across devices and sites. It ensures that the reliability and availability benefits of using multiple locations are realized.

5. Click **Finish** to show the **Dashboard**.

Devices that are completed with physical configuration are shown on here, although it can take a few minutes for the Manager Web Interface to see the devices to be approved. Continue with configuration tasks here or move to set up more accounts so that other users might potentially take on some or all of these tasks.

To return to the **Dashboard** page, click the logo in the upper left corner of the page.

Chapter 3. Configuration

System access

Attention: Configuration tasks require an account that has access rights to the **Configure** page.

Approve registered devices

Devices register themselves with the Manager Web Interface as part of the initial configuration. Before devices can be used, they must be approved by using one of the following methods. Devices to be approved are listed on the **Configure** page under **Devices Pending Approval**.

Devices can be approved individually or in bulk. Bulk selection can be achieved by clicking multiple check boxes or all devices. The workflow is slightly different between the two approval methods. The fingerprint is needed from the initial configuration form to verify the key that is shown as part of the device approval. Go to **External CA Device Approval** if you defined external CAs for the devices.

Note: During device approval, SliceStor devices can be assigned to sites. If SliceStor devices are to be deployed to more than one site, create the sites even if the devices are initially staged in one location. Defining the sites and assigning SliceStor devices to those sites ensures that data is written equally across devices and sites. It guarantees that you realize the reliability and availability benefits when using multiple locations.

Approving a single device

Procedure

1. From the list of **Devices Pending Approval** on the bottom of the **Configure** page, click a single device to display the devices information page.
2. Verify that the key fingerprint is correct and click **Approve** or **Deny** from the action bar.
3. Click **Deny** to remove the device from the pending list.
 - To add the device back to the pending approval list, rerun **manager ip {ip-address}** (or **manager ip_ipv6 {ip-address}**), where {ip-address} is the IP address from the denied device.
 - A single device approval gives the option to add the device to an existing site or create a new site and set an optional alias.
4. Verify that the devices are added on the Summary section under the Devices tab and that the device summary count is updated accurately.

Approving multiple devices

Devices must be approved by the manager before they are added to the system.

Before you begin

Attention: Any time a device is added or a vault, site, cabinet, or an administration configuration is changed, the manager must be backed up by using the **Backup and Restore** utility. Permanent data loss can occur if the manager database becomes corrupted. Periodic backups must also be performed to preserve historical statistics and log information. See **Backup** settings on the **Administration** menu.

Procedure

1. Click the check boxes next to the device **Hostname** column to approve that device, from the list of **Devices Pending Approval** on the bottom of the **Configure** page.

2. Click **Bulk Approve/Deny**. After the devices are selected for approval, the device registration screen appears.
3. Review the security key fingerprint information for the devices to ensure that the registration is from a trusted source. For bulk approvals, you can select individual or all devices.
4. Click **Approve**, **Deny**, or **Cancel** from either action bar. After the devices are approved, a menu for site assignment appears.
5. Click **Deny** to remove the device from the pending list.

If a Device is Denied, the device must be reinitialized from the IBM Appliance Configuration utility before it can be Approved. (For more information, see Configuring appliances).

It is necessary to rerun **manager ip {ip-address}** (or **manager ip_ipv6 {ip-address}**), where {ip-address} is the IP address from the denied device.

Note: The only choice is to approve all or deny all. If only a subset of the devices needs approval, click **Cancel**.

6. Associate devices with sites by checking one or more devices and selecting the site to which they belong. During bulk operations, the approved devices are displayed at the top of the **Bulk Edit Device Site** page, and a list of sites appear at the bottom.

Note: If multi-node devices need approval, selecting a node automatically selects all other nodes in the same chassis.

7. Click **Save** to confirm the mapping of devices to sites. As devices are associated with sites, they are removed from the device list at the top of the page.
8. Continue associating devices with sites until complete.

Multiple devices can be selected by checking the check boxes to the left of each site.

After all bulk devices are assigned to a site, an alias can be set for each device in its adjacent text box.

9. Verify that the devices are added on the summary section under the devices tab and that the device summary count is updated accurately.
10. Click **Save**.

Setting an alias for each device is optional. For either approval method, the device displays an inactive state until polling is completed. It takes less than 1 minute. The **Bulk Edit Device Alias** page is shown.

Cabinets

Importing a cabinet configuration

The Cabinet application can be used to group the devices by cabinets to facilitate maintenance and inventory tracking operations.

About this task

The configuration of devices within cabinets can be imported from a cabinet configuration .csv (comma-separated values) file. The file format can be found in the Manager Reference Guide.

Procedure

1. Log in to the Manager Web Interface with your user name and password.
2. Click **Configure** tab.
3. Click **Import** in the **Import/Export Configure Cabinet Description File** bar.
4. Select the file.
5. Enter your password.

6. Click **Import** to import the configuration.

Attention: Importing a file deletes the current configuration and cannot be reversed. Export the current configuration as a backup before you import the new file.

Exporting a cabinet configuration

The configuration of devices within cabinets can be exported as a .csv file.

Procedure

1. Log in to the Manager Web Interface with your user name and password.
2. Click the **Configure** tab.
3. Click **Export** in the **Import/Export Cabinet Description File Configure Cabinet** action bar.
4. Save the .csv file to a location on your computer.

Creating a cabinet

Use this procedure as an alternative to importing a cabinet.

Procedure

1. Log in to the Manager Web Interface with your user name and password.
2. Click the **Configure** tab.
3. Click the **Create Cabinet** link in the **Summary** section.
4. Enter the information in the form.
5. Click **Save** to complete the creation of the new cabinet.

Configuring a cabinet with drag-and-drop

As an alternative to editing and importing the cabinet configuration file, the configuration of devices within cabinets can also be modified by dragging devices into and out of the cabinet on the **Edit Cabinet** page.

About this task

In the cabinet view, as part of a multi-node Slicestor device configuration (see “Approve registered devices” on page 5), all nodes that belong to the same chassis are presented together. In the UI, a single 4U instance appears that can be placed in the cabinet.

Note: Virtual appliances can't be added to a cabinet and do not appear in the **Devices Not In Cabinet** column.

Procedure

1. Log in to the Manager Web Interface with your user name and password.
2. Click the **Configure** tab.
3. Click the arrow next to name of a site from the left navigation tree under **Sites - <site name>** to expand it.
4. Click the name of the cabinet you want to configure.
5. In the Edit Cabinet page, click **Change** in the action bar.
6. Drag a device from the **Devices Not In Cabinet** column into a slot in the **Cabinet** column to add it to a cabinet.
7. Drag a device out of the **Cabinet** column to remove it from that cabinet.
8. Add generic devices by clicking **Add Generic Device** and edit the description in place by clicking the text.
9. Click **Update** to save the configuration changes.

Access pools

An access pool is a logical set of zero or more Accesser[®] nodes.

Access pools can be assigned the same set of attributes, such as programming interface and access service port numbers, in aggregate instead of per device. An access pool can be deployed to a selection of vaults and vault mirrors in a system. An access pool can be removed from the system, but affects all deployed vaults and mirrors. A password prompt is required since delete is a non-recoverable action. However, it is possible to re-create and re-deploy the access pool if deleted.

Creating an Access Pool

These steps allow you to set a management vault at the access pool level.

Procedure

1. Click the **Create Access Pool** link under the **Summary** heading on the home page to display the **Create New Access Pool** page.
2. If changes to the Access Pool are needed:
 - a. Enter a name for the new Access Pool in the **Name** field.

Note: The maximum length of the **Name** field is 255 characters.

- b. Enter a description for the new Access Pool in the **Description** field.
- c. From the **API Type** drop-down menu, select the wanted programming interface:
 - **Cloud Storage Object**
 - **Simple Object**
 - **OpenStack Object Storage**

Note: When the system is in container mode, the Simple Object API Type, Protected Vaults, and Protected Objects are not supported.

Note: Operations on Protected Vaults and Protected Objects are only supported via Cloud Object Storage.

If you plan to deploy this access pool to a protected vault, the Simple Object and OpenStack Object API Types are not supported.

- d. Check the check boxes in the **Service Type** list for the ports that should listen for requests on the Accesser nodes in the Access Pool.

Table 3. Service Type ports

Port	Secure	Listening Service
80	No	HTTP port for specified API Type
8080	No	HTTP alternative port for specified API Type
443	Yes	Secure HTTP port for specified API Type
8443	Yes	Secure HTTP alternative port for specified API Type
8337	No	HTTP port for Service API
8338	Yes	HTTPS port for Service API

Note: By default all the ports for the specified API Type are checked. The Service API ports are not checked by default as it is usually not necessary to open the Service API on all access pools. However, Service API port 8338 must be open if the access pool is being used for file system vaults.

The Service API is used to create or configure storage accounts and is normally used by storage-as-a-service portals. You must deploy a container vault to the access pool before you can use the Service API. Accounts that use the Service API must have the Service Account role.

Attention: At least one service must remain checked to maintain I/O operations to the attached Vaults and Accesser nodes.

- e. Specify the **S3 Virtual Host Suffix**, if virtual host styled addressing is wanted. To use this feature:
 - **Cloud Storage Object** must be selected from the **API Type** drop-down menu.
 - Vaults must have DNS-compliant names.
 - The provided name must also be set up in the DNS servers to route to Accesser nodes that are part of this access pool.
 - f. Specify any needed **Additional Subject Alternative Names** that need to be added into the Accesser nodes certificates. By default, the certificates include the devices hostname and IPs. If the **External PKI** feature is being used, this subject alternative name is added to the CSR.
 - g. If **External PKI** feature is being used and the certificate is being signed by an external certificate authority, it might not be possible to include the Accesser node IPs in the certificate. In this case, clear **Include default IPs in Subject Alternative Names** to avoid failures when the certificate is loaded into the Manager Web Interface.
 - h. If Container Mode is enabled, a **Default Container Vault** drop-down appears. Any deployed container vault can be selected. In which case, if a LocationConstraint is not specified during a PUT bucket requests, accessers in this access pool creates the container in this container vault. If no default container vault is selected, the LocationConstraint is a required parameter in the PUT bucket request.
 - i. Under the **Access Devices** heading, check the check box to the left of each Accesser node from the available list that should be added to the Access Pool. Click the **Select All** link to add all of the Accesser nodes.
3. In the **Deployment** section, there are three filters to limit the display of vaults to select:
 - **Storage Pool**
 - **Item Type**
 - **Text Search**
 4. Check the check box to the left of each vault to be deployed from the **Deployment** list.

Note:

- Click the **Select All** link to select all Vaults.
 - Click one vault, hold down the **Shift** key, then click another vault farther down the list to select a range of Vaults.
5. Click **Save** to create the Access Pool.

Editing an Access Pool

These steps allow you to edit the management vault at the access pool level.

Procedure

1. Click the **Configure** tab.
2. Click the **Open All** link above the Left Navigation Tree.
3. Click the link of an Access Pool to configure to display the **Access Pool** page.
4. Click **Change** to display the **Editing: {Access Pool}** page.
5. If changes to the Access Pool are required:
 - a. Change the name for Access Pool in the **Name** field.
 - b. Change the description for Access Pool in the **Description** field.

- c. From the **API Type** drop-down menu, select the wanted programming interface:
 - **Cloud Storage Object**
 - **Simple Object**
 - **OpenStack Object Storage**

Note: When the system is in container mode, the Simple Object API Type, Protected Vaults, and Protected Objects are not supported.

Note: Operations on Protected Vaults and Protected Objects are only supported via Cloud Object Storage.

- d. Check the check boxes in the **Service Type** list for the ports that should listen for requests on the Accesser nodes in the Access Pool.

Table 4. Service Type ports

Port	Secure	Listening Service
80	No	HTTP port for specified API Type
8080	No	HTTP alternative port for specified API Type
443	Yes	Secure HTTP port for specified API Type
8443	Yes	Secure HTTP alternative port for specified API Type
8337	No	HTTP port for Service API
8338	Yes	HTTPS port for Service API

Note: By default all the ports for the specified API Type are checked. The Service API ports are not checked by default as it is usually not necessary to open the Service API on all access pools. However, Service API port 8338 must be open if the access pool is being used for file system vaults.

The Service API is used to create or configure storage accounts and is normally used by storage-as-a-service portals. You must deploy a container vault to the access pool before you can use the Service API. Accounts that use the Service API must have the Service Account role.

Attention: At least one service must remained checked to maintain I/O operations to the attached Vaults and Accesser nodes.

- e. Specify the **S3 Virtual Host Suffix** if virtual host styled addressing is wanted. To use this feature:
 - **Cloud Storage Object** must be selected from the **API Type** drop-down menu.
 - Vaults must have DNS-compliant names.
 - The provided name must also be set up in the DNS servers to route to Accesser nodes that are part of this access pool.
- f. Specify any required **Additional Subject Alternative Names**, which need to be added into the Accesser nodes certificates. By default, the certificates include the devices host name and IPs. If the **External PKI** feature is being used, this subject alternative name is added to the CSR.
- g. If **External PKI** feature is being used and the certificate is being signed by an external certificate authority, it might not be possible to include the Accesser node IPs in the certificate. In this case, clear **Include default IPs in Subject Alternative Names** to avoid failures when the certificate is loaded into the Manager Web Interface.
- h. If container mode is enabled, a **Default Container Vault** drop-down appears. Any deployed container vault can be selected. In which case, if a LocationConstraint is not specified during a PUT bucket requests, Accessers in this access pool creates the container in this container vault. If no default container vault is selected, the LocationConstraint is a required parameter in the PUT bucket request.

- i. Under the **Access Devices** heading, check the check box to the left of each Accesser node from the available list that should be added to the Access Pool.
Click the **Select All** link to add all of the Accesser nodes.
 - j. In the **Deployment** section, there are three filters to limit the display of vaults to select:
 - **Storage Pool**
 - **Item Type**
 - **Text Search**
 6. Check or clear the check box to the left of each vault to be deployed or removed from the **Deployment** list.
- Note:**
- Click the **Select All** link to select all Vaults.
 - Click a vault, hold down the Shift key, then click another vault farther down the list to select a range of Vaults.
7. The **Submit Confirmation** dialog box displays if:
 - Accesser nodes were added or removed.
 - Vaults or Vault Mirrors were added or removed.
 The dialog box shows the count of vaults and mirrors that are deployed or removed.
 8. Choose an action to create the new Access Pool.
 - Click **Submit** to create the new Access Pool.
 - Click **Review Changes** to return to the **Create** page and make changes (returning to Step 3).

Editing a Device

About this task

Note: Devices in storage pools cannot alter their management vault.

Procedure

1. Click the **Configure** tab.
2. Click the **Open All** link above the Left Navigation Tree.
3. Click the link of any Accesser node to display the **Accesser: {device-name}** page.
4. Click **Change** on the **Configure Device** page to display the **Editing: {Device}** page.
5. If changes to the node are required:
 - a. Change the alias for the Accesser node in the **Alias** field.
 - b. Change the description for the Accesser node in the **Description** field.
 - c. Select a different Access Pool from the **Access Pool** drop-down menu to change the Access Pool of the Accesser node.

Note: The drop-down menu shows the current Access Pool to which the Accesser node belongs, or **No Access Pool** if the device is not in any Access Pool.

The drop-down menu is available for an Accesser node only.

6. Choose an action to create the new Access Pool:
 - Click **Update** to finalize the changes to the Accesser node.
 - Click **Cancel** to cancel any changes to the Accesser node.

Moving an Accesser[®] Node

Access Pool devices can be moved to a different Access Pool.

Before you begin

Note: Moving an Access Pool affects vault and mirror deployments.

Procedure

1. Click **Configure**.
2. Click **Open All** above the Left Navigation.
3. Click the link of an Access Pool to display the **Access Pool: {access-pool-name}** page.
4. Click **Move Device** to display the **Access Device Selection to Move from Access Pool: {access-pool-name}** page.
5. Click the check box to the left of the Accesser[®] Nodes to move.
6. Click the radio button of the destination Access Pool to which the Accesser[®] Nodes are to be moved to.
7. Click **Submit** to move the Accesser[®] Nodes.

Configuring HTTPS certificates for Access Pools

You can use certificates for HTTPS access that are trusted inside your organization instead of the default Manager signed certificates.

Procedure

1. Click the **Configure** tab.
2. Click the **Open All** link above the Left Navigation.
3. Click the link of an Access Pool to display the **Access Pool: {access-pool-name}** page.
4. In the **Access Pool HTTPS Certificate** section, click **Configure** to display the Editing Access Pool HTTPS Certificate page.
5. Paste PEM-formatted private key and certificate text into the corresponding **Private Key PEM** and **Certificate PEM** fields.
 - To add more certificates, paste one after another.
 - To remove a single certificate, delete on the text for that certificate.
 - To remove all certificates, delete all contents.
6. Click **Update** to update the certificates for the Access Pool.

Deleting an Access Pool

Procedure

1. Click the **Configure** tab.
2. Click the **Open All** link above the Left Navigation Tree.
3. Click the link of an Access Pool to display the **Access Pool: {access-pool-name}** page.
4. Click **Delete Access Pool** to display the **Delete Access Pool** page with the count of vaults that are to be deleted.
5. Enter your password in the **Password** field.
6. Click **Delete** to delete the Access Pool.

Note: The selected Access Pools are immediately deleted. No confirmation dialog displays when **Delete** is clicked.

Storage pools

A storage pool is defined by a logical grouping of Slicestor devices that are used to store vault data.

A vault is created on a storage pool. These are a few rules to keep in mind:

- Storage pools must be defined before vault creation. If pools are not defined, vault creation is redirected to the Create Storage Pool page.
- Multiple vaults can be created on a storage pool.
- A Slicestor device can be a member of a single storage pool. A Slicestor device can be replaced and the data is evacuated to another device. Network throughput that is allocated for this operation can be controlled via a data evacuation rate limit parameter.
- Storage Pools can be created with packed storage enabled. Packed storage improves small object performance.

Note: Packed storage can be enabled when all devices are upgraded to ClevOS 3.4.0.

Creating a storage pool

A storage pool is defined by a logical grouping of (Slicestor) devices used to store vault data. A vault is initially created on a storage pool, and can be expanded by either using an existing storage pool or by creating a new one.

Before you begin

Note: A Slicestor device can only be assigned to one storage pool. Likewise, a new storage pool can only be created from unassigned devices. Once created, a storage pool cannot be expanded, but additional pools can be created and merged to expand a vault.

When creating a storage pool, consider the following items:

- Slicestor devices can be selected from any number of sites.
- Each pool width (for a given vault) must be a multiple of the vault width.
- Devices in the pool can be mixed with different model types with different capacities.
- Each device can only be allocated to a single pool.
- Each pool can only be used by a single vault type.
- A minimum number of three Slicestor devices is required to create a storage pool.
- The smallest drive count must be at least half of the largest drive count. If necessary, you can modify this constraint through advanced configuration by contacting IBM Customer Support.

When creating a storage pool using Concentrated Dispersal, the following additional items apply:

- Each pool must have a minimum of three devices and a maximum of seven devices (3 devices \leq storage pool width \leq 7 devices).
- The vault width is a restricted multiple of the pool width (for a given vault). Example: Storage pools with three devices can have vault widths of 18 and 36.
- All devices in a pool must have the same number of drives.

Procedure

1. Right-click on any text in the left navigation tree on the **Monitor**, **Configure**, and **Dashboard** pages.
2. Click the link **Create Storage Pool** in the **Summary** section of the **Dashboard** and **Configure** pages.

Storage pools can have different capacity nodes. A model group, consisting of nodes that can be included in the same storage pool, must be selected. If nodes across multiple sites are available, the required site must be selected.

Note: If a storage pool has nodes that span multiple sites, it is recommended that the nodes be balanced across the sites. Not balancing the nodes can introduce read or write availability issues if a single site outage occurs.

3. Type a name for this storage pool in the **Name** field.
4. Select an IDA width for this storage pool from the **Width** drop-down box. To use Concentrated Dispersal vaults in this pool, select a small width of 3 - 7 SliceStor® Devices.
5. Check the **Packed Storage** check box to enable packed slice storage for this storage pool.

CAUTION:

Packed storage can be enabled when all SliceStor nodes are upgraded to ClevOS 3.4 or newer.

Note: Ask IBM support if enabling packed slice storage benefits current storage needs.

6. Select nodes either automatically using **Suggest Devices** per the above rules, or manually. **Suggest Devices** selects the width number of devices by using the selected check boxes that list devices by drive count. More than one drive count check box can be selected.

Note: If you reimaged a SliceStor appliance to for use in a storage pool, the device might not be available for use immediately after approval. The SliceStor appliance can be used when the following event is generated in the console:

SliceStor started for the first time.

7. After the devices are identified, click **Save** to create the Storage Pool.

What to do next

Attention:

If there is not an even distribution of SliceStor nodes across the available sites where the loss of one site would make Vaults either unusable or read-only, the Manager Web Interface displays a confirmation dialog box that asks the operator if they accept the settings with the risks they present:

The selected devices are not balanced evenly across sites. This could lead to read and write availability issues in the case of a site outage. Do you still wish to continue?

The operator can click **Cancel** or **OK** to change or keep the settings.

From the **Configure Storage Pool** page, click **Change** to rename this storage pool or **Monitor** to shortcut to the monitor function for this storage pool. A different storage pool can be selected from the navigation tree.

Editing a storage pool

Allows you to edit the management vault at the storage pool level.

Monitoring storage capacity

Capacity is reported after SliceStor® Devices are approved.

Allocated, unallocated, and raw capacity can be found in three locations:

Table 5. Capacity locations

Page	Description
Monitor	Overall and individual capacity
Configure	Summary of overall unallocated, raw, and allocated capacity

Table 5. Capacity locations (continued)

Page	Description
Dashboard	Summary of overall unallocated, raw, and allocated capacity

In the **Monitor** tab, when you click a Vault in the left navigation tree, Raw and Usable used space and free space are shown. Raw space is a precise calculation. Usable space is an estimate because actual usage depends on the size of the files stored. Estimates are indicated by ~. Vault Capacity view is selected by default instead of Storage Pool Capacity. You can toggle to the **Storage Pool Capacity** view to display **Other Used**, which consists of space that is taken by other vaults and incompressible overhead size.

Note: Vault capacity is impacted if a Slicestor® Device is down.

In situations where data is missing, for example, due to replacement of failed drives, vault and storage pool capacity reporting anomalies might arise. The value that is being reported on the Manager Web Interface can give the impression that more space is available for writing data than exists. When the storage pool is full, attempts to write new objects fail. After rebuilding completes for the data on these drives, the values that are reported are correct.

When drives are quarantined in a Slicestor® Device, the allocated and unallocated numbers that are associated with the **Device Capacity** section in the **Monitor Device** page of the Manager Web Interface are temporarily inaccurate. Within a few minutes, the correct values are displayed.

Adding storage capacity

Storage capacity can be expanded in three ways.

Procedure

Add a storage pool that consists of a fresh collection of devices.

- “Expanding a storage pool”
- “Merging two storage pools” on page 18

CAUTION:

In some cases, it is best to add a storage pool instead of merging storage pools, due to restrictions on vault creation.

What to do next

Note: Contact IBM Customer Support for any questions about merge storage pool usage and limitations.

Expanding a storage pool

Follow these steps to expand a Storage Pool.

Before you begin

Before you expand a Storage Pool, the following preconditions must be met:

1. Upgrade all Accesser nodes and Slicestor nodes to release 3.6.0 or later.
2. Approve a number of Slicestor nodes equal to a multiple of the least common multiple of all the vault IDA widths for the vaults that use the Storage Pool to be expanded. If you are using Concentrated Dispersal, you can expand by using either the Concentrated Dispersal set size or an integer multiple of the full IDA Width.

Note: If you reimaged a SliceStor appliance to expand a storage pool, the device might not be available for use immediately after approval. The SliceStor appliance can be used when the following event is generated in the Event Console:

SliceStor started for the first time.

Procedure

1. Under the **Configure** tab, click the **Storage Pool** to be expanded.
2. Expand the list under Storage Pools on the left navigation tree.
3. Click the **Storage Pool** to replace a device. The **Storage Pool: <Pool Name>** page displays.
4. In the SliceStor Devices section, click **Change Sets and Devices**.
5. In the Expand Storage Pool section, click **Configure Storage Pool Expansion**.
6. From the **Width** drop-down menu under the **General** bar, select the **IDA Width** for the expanded Storage Pool.
7. Select the SliceStor nodes to be used to expand this pool. These devices are used to form a new **Device Set** within this Storage Pool.
8. If the nodes to be added are decided, select the SliceStor nodes for the new Device Set by 1 of three options:
 - Drive count by using the check boxes to filter SliceStor devices by number of drives configured.
 - Sites by using the check boxes for one or more Sites.
 - Individual nodes via the check boxes in the **Devices** section.
9. If the nodes to be added are undecided, click **Suggest Devices** to allow the Manager application to select SliceStor nodes on your behalf.
10. Click **Continue**. A confirmation page appears.
11. Click **Save** to save your changes.
12. Monitor the progress of the reallocation and ensure it continues to completion.
13. Change the Storage Pool Reallocation Rate.

Changing the storage pool reallocation rate

Storage Pool expansion can be throttled.

Procedure

1. Click the **Storage Pool** from the **Monitor** tab.
2. Click **Change** under the **Data Reallocation in Progress** window. The **Edit Data Reallocation** dialog box displays.
3. For all device sets:
 - a. Check the **Enable** check box in the **Bulk Change** section.
 - b. Type a transfer limit in MB per second in the **MB/s** field.
4. For one device set:
 - a. Check the **Per-Device Rate Limiting** check box for that Device Set.
 - b. Type a transfer limit in MB per second in the **MB/s** field.
5. Click **Submit** to accept these changes.

Pausing a storage pool reallocation

A Storage Pool expansion can be paused.

Procedure

1. Click the **Storage Pool** from the **Monitor** tab.

2. Click **Change** on the **Data Reallocation in Progress** window. The **Edit Data Reallocation** dialog box displays.
3. To pause a reallocation:
 - a. For all device sets, check the **Pause** check box in the **Bulk Change** section and click **Apply to All Sets**.
 - b. For one device set, check the **Pause** check box for that Device Set.
4. Click **Submit** to accept these changes.

Note: When a Device Set's reallocation is paused, it appears with a **-Paused-** indicator next to the Device Sets name.

Resizing a storage pool

How to resize a storage pool.

About this task

You can resize an existing partially populated storage pool to take advantage of extra capacity.

Note: To resize a storage pool, the storage pool must contain multiple sets.

Procedure

1. Install more drives in each Slicestor® Node in the storage pool.
2. When most of the devices have the additional capacity, click the **Configure** tab.
3. In the navigation tree, click the storage pool that you want to resize. The Configure Storage Pool page displays a notification that the system identified more capacity.
4. When you finish installing extra drives, click **Done Adding Capacity** in the notification. The notification displays the estimated capacity before resize and the project capacity after resize.
5. Confirm the capacity values and click **Approve & Start Data Reallocation**. The resize process begins.
6. Monitor the resize process on the Monitor Storage Pool page.

Replacing storage pool sets

How to replace storage pool sets.

Before you begin

Before you replace Storage Pool sets, you must perform the following preconditions:

1. Upgrade all Accesser® Nodes and Slicestor® Nodes to ClevOS 3.10.0 or newer.
2. Approve an amount of Slicestor® Nodes equal to a multiple of the least common multiple of all the vault IDA widths for the vaults that use the Storage Pool that you want to replace. If you are using Concentrated Dispersal, you can replace sets by using either the Concentrated Dispersal set size or an integer multiple of the full IDA Width.

To replace one or more Storage Pool sets:

Procedure

1. In the **Configure** tab, click the **Storage Pool** you want to replace.
2. In the Slicestor Devices section, click **Change Sets and Devices**.
3. Click **Configure Set Replacement**.
4. In the **Sets to Replace** section, select the **Storage Pool** sets you want to replace.
5. Select the Slicestor® Nodes you want to use to replace the Storage Pool sets. The devices are used to form a new Device Set within this Storage Pool.

6. Select the Slicestor[®] Nodes you want to add for the new Device Set in one of the following ways:
 - Select the check boxes for choosing Slicestor models with similar drive counts
 - Check the check boxes for one or more Sites to add nodes by Site.
 - Check the check boxes in the Devices section to add individual nodes.
 - Click **Suggest Devices** to allow the Manager application to select Slicestor[®] Nodes on your behalf.
7. Click **Continue**. A confirmation page appears.
8. Click **Save**.
9. Monitor the progress of the reallocation and ensure it continues to completion.

Removing storage pool sets

To remove one or more Storage Pool sets, follow these instructions.

Before you begin

Before you can remove Storage Pool sets, you must upgrade all Accesser[®] Nodes and Slicestor[®] Nodes to release 3.10.0 or newer.

Procedure

1. In the **Configure** tab, click the **Storage Pool** that you want to replace.
2. In the Slicestor Devices section, click **Change Sets and Devices**.
3. Click **Configure Set Removal**.
4. Select the **Storage Pool** sets you want to remove and click **Continue**. A confirmation page appears.
5. To remove the Storage Pool sets, click **Save**.
6. Monitor the progress of the reallocation and ensure it continues to completion.

Merging two storage pools

The **Merge Storage Pool** method expands capacity by "combining" a heavily used storage pool with either a new storage pool or an existing, less used storage pool.

Before you begin

Attention: This method does not work in ClevOS 3.8.0 or newer unless the Storage Pool was merged in a previous version.

CAUTION:

Consider the following issues when you merge storage pools:

- Only two storage pools can be merged at a time.
- Storage pools can be merged regardless of the Slicestor device models that are used in the pools.
- After the storage pools are merged, the new pool capacity is the sum of the capacities that are associated with the original storage pools.
- All vaults that are associated with the original storage pools now become visible to the new merged storage pool.
- Storage pools of any size can be merged when no vaults exist on either pool. However, after the storage pools are merged, only vaults that are a divisor of both the storage pool widths can be created. It is advantageous in many circumstances, but exceptions exist.

For instance, Pool1 (Width = 8) and Pool2 (Width = 9), where both Pool1 and Pool2 do not have any vaults. After the pools are merged, a user can create vaults of size 1.

- Storage pools of the same width can be merged irrespective of the number and size of vaults on either pool.
- If the storage pool widths are different and vaults exist, the width of each pool must be an exact multiple of the width of all the vaults that exist on either pool. This scenario allows larger or smaller capacity storage pools to be merged.

Given Pool1 (Width = 8) with vault11 (Width = 8) and vault 12 (Width = 4) and Pool2 (Width = 16) with vault21 (Width = 8), Pool1, and Pool2 can be merged.

- Some storage pools cannot be merged.

Given Pool1 (Width = 8) with vault1 (Width = 8) and Pool2 (Width = 16) with vault2 (Width = 16), Pool1, and Pool2 cannot be merged since vault2 cannot be associated with Pool1 after the merge. Such options are not presented on the Merge Storage Pool page in the Manager Web Interface.

Follow these steps when you merge storage pools. Storage pools can be merged in any order.

- **Pool_test1** is the original storage pool.
- **Pool_test2** is either a new storage pool or another storage pool that has lower usage when compared to **Pool_test1**.

About this task

Merge Storage Pools is a new approach for performing capacity expansion, replacing individual vault expansion. It can be used to expand storage capacity with a new or less usage storage pool. With this approach, two storage pools, along with all vaults across these storage pools, are merged together. Only two storage pools can be merged at a time. More storage pools can be merged with the newly merged storage pool over time as needed. After the pools are merged, the new pool capacity is the sum of the capacity that is associated with the original storage pools. All vaults that are associated with the original storage pools now become visible to the merged storage pool.

When two storage pools are merged, two sets are formed. Set 1 contains the devices that are associated with the first storage pool, and Set 2 contains the devices that are associated with the second storage pool. The set names (Set 1, Set 2) cannot be edited. When a new vault is created, its width must be the same as a single set. Also, when writes are performed to a vault in the merged storage pool, the objective is to fill the storage pools evenly. Thus, writes are adjusted based on the capacity available within individual sets. As an example, if both sets have equal capacity available, writes are ~50/50 across the sets. However, if one set is full and the other is empty, writes are 0/100.

Procedure

1. Log in to the Manager Web Interface through a browser.
2. Click **Create a Storage Pool** on the **Configure** page.

Note: If an existing storage pool exists for the merge operation, skip this step.

Pool_test2 is a newly created storage pool.

3. Go to the **Configure** page of the original storage pool.
4. In the Slicestor Devices section, click **Change Sets and Devices**.
5. In the Merge Storage Pool section, click **Configure Storage Pool Merge**.
6. Select the storage pool to be merged (**Pool_test2**). All allowable storage pool options are presented.
7. Click **Save**. Two Sets, **Set 1** and **Set 2**, are created, and each Set corresponds to the devices in the respective, original storage pools. Each Set has three devices. The storage pools are merged.

What to do next

The storage pool name that is retained after the merge depends on the storage pool from which the **Merge Storage Pool** operation was initiated (Steps 3 and 4).

In this example, **Pool_test1** was used. As a result, the name of the new merged storage pool (after Step 6) is **Pool_test1**.

Note: If vault expansion was performed before upgrade, storage pools that are associated with any expanded vaults are merged automatically.

Delete a storage pool

A Storage Pool can be removed from the system if no vaults are attached to this pool. A password prompt is needed for this unrecoverable action.

The devices and vaults that are associated with this storage pool are also shown. Click **Change** in the device box to replace a specific device in the storage pool.

Replace and evacuate the data from Slicestor devices

A Slicestor device can be replaced and its data can be evacuated to another device.

The device that is moving out of the storage pool is called the source device, and the device that is moving in to the storage pool is called the destination device. Before a source device is replaced, the following internal checks must pass:

- The storage pools that contain the source and destination devices must be updated to ClevOS 3.4.0.0 or later.
- The destination device must have a capacity equal to or greater than the source device.
- The destination device must be reimaged if it previously contained data.

Note: When a Slicestor appliance is reimaged and approved, it takes some time before device is available to be used. The Slicestor appliance can be used when the following event is generated in event console:

Slicestor started for the first time.

- Both the source and destination device must not already be a part of an ongoing device replacement.
- The replacement process must not impact vault health.
- When evacuating data from more than one device, the storage pool must be healthy. The maximum number of devices you can evacuate in parallel is equal to the difference between the vault Width and Alert Level, and is also dependent on the current health of the pool.

For example, if a storage pool has 16 devices and a 16/9/11/13 vault, you can evacuate up to three devices at a time without affecting the pool's health. Attempting to evacuate a fourth device brings the vault below the Alert Level and is considered unsafe. Similarly, if two of those devices are already unhealthy, you can evacuate only one device.

Attention: Data cannot be evacuated during data reallocation. Execute any planned data evacuations before expanding storage pools. The system should be in as healthy a state as before performing expansion: replace any bad drives, perform any evacuations, and so on.

After data reallocation starts, drives can be replaced, but devices that are replaced during data reallocation will not evacuate data. Do not replace devices during data reallocation; all data slices on the old device are lost.

Replacing a Slicestor[®] Device

Before you begin

CAUTION:

This procedure assumes that a new Slicestor[®] Device is properly installed and configured to serve as the destination Slicestor[®] Device.

Procedure

1. Click the **Configure** tab.
2. Click a **Storage Pool** in the left navigation tree to view the **Storage Pool: {storage-pool-name}** page.
3. In the **Slicestor Devices** section, click **Change Sets and Devices**.
4. In the Replace Individual Device section, click **Configure Device Replacement**.
5. From the list that is provided, select a source device.
6. Click **Next**.
7. If wanted, the network throughput that is allocated for data evacuation can be set:
 - a. Check the **Data Evacuation Rate Limit** check box.
 - b. Enter a throughput limit in MBps. The limit can be changed later from the top-level **Configure** page.
8. On the right side, select the destination device.
9. Click **Confirm Replacement**. A pop-up asking for confirmation appears.
10. Click **OK**. The data replacement and evacuation starts. Evacuation progress can be tracked on the **Monitor** page of the destination device.

Note:

All data evacuation events are visible only on the Source Slicestor device.

The graphs that are found on the **Monitor** page for the Source Slicestor device are as follows.

Scanning Rate

Shows **Graph Data Not Available** instead of any rate data.

Outgoing Rebuild

Shows no **Rebuild Bytes Sent**.

Incoming Rebuild

Shows **Rebuild Slices Received**.

Pause/resume data evacuation

- To pause a data evacuation, click **Pause Evacuation** from the source devices **Device Monitor** page.
- To resume a data evacuation, click **Resume Evacuation** from the source devices **Device Monitor** page.

Terminating data evacuation

Data evacuation can be terminated, if necessary, before it is completed. It can be necessary due to hardware failure on the source or destination devices or if the operator chooses to terminate the data evacuation.

About this task

Alternately, just reimaging the Source device also terminates data evacuation.

Attention: Terminating data evacuation before completion results in loss of the remaining data slices that are not yet copied to the destination. These data slices must be rebuilt later.

Procedure

1. Click **Pause Evacuation** from the source devices **Device Monitor** page.
2. Remove the Source device from the Managers **Device Configure** page.
The **Enter Password** page appears. Enter your password to confirm the device removal.

Rolling back a data evacuation

If the need to return the evacuated data to the original source device (rollback) arises during evacuation.

About this task

Note: You can only rollback a data evacuation for individual devices.

Procedure

1. Pause the evacuation on the Source device.
2. Click the **Configure** tab.
3. Expand the list under Storage Pools on the left navigation tree.
4. Click the **Storage Pool** to replace a device. The **Storage Pool: <Pool Name>** page displays.
5. In the Slicestor Devices section, click **Change Sets and Devices**.
6. In the Replace Individual Device section, click **Configure Device Replacement**.
7. In the Current Slicestor Devices list, select the current destination device. The device actively shows as being in evacuation.
8. Click **Next**.
9. In the Replacement Slicestor Devices list, select the current source device. The device actively shows as being in evacuation.
10. Click **Confirm Replacement** to confirm the swap of source and destination devices. The data evacuates from the original destination device to the original source device.

Changing the destination device

Follow these steps to change the destination device.

About this task

While data evacuation is taking place, it might be necessary to change the destination device (for example, if hardware problems exist on the current one, or if it was incorrectly chosen). No

straightforward way exists to change the destination device besides rolling back and starting a new evacuation.

Procedure

1. Roll back evacuation per “Rolling back a data evacuation” on page 22.
2. Wait for the evacuation to roll back. If it is necessary to stop the evacuation, reimage the original destination device (new source device). It stops the rollback.
3. Start a new evacuation with the correct destination device.

Change the rate of evacuation during evacuation

A user can enable/disable the data evacuation rate limit on all source devices on the top level Configure page.

Click **Configure** on the **Data Evacuation Rate Limit Configuration** action bar and complete the form that is provided.

Note: On the **Configure** page of a source device, the **Remove** button is disabled while data evacuation is in progress.

Attention: A user is allowed to pause evacuation and remove the source device. It is a destructive action and can result in losing all data that is not copied.

Troubleshooting data evacuation incidents

Incident	Action
Error occurred during data evacuation.	<p>It is caused by not being able to persist data evacuation progress tracker to the Source device OS disk. If the problem persists or causes evacuation to halt, the user might need to replace the OS drive on the Source device to continue data evacuation.</p> <p>See instructions for replacing the OS drive on:</p> <ul style="list-style-type: none">• SliceStor 2210 device• SliceStor 2212 device• SliceStor 1440 device• SliceStor 2440 device• SliceStor 4100 device• HP ProLiant SL4500 Series Quick Setup Instructions (2-node)• Quick Setup Instructions HP ProLiant SL4500 Series (3-node)• HP SL4540 User Guide• HP ProLiant Gen8 Server with HP iLO Management Engine Setup Guide

Incident	Action
Data evacuation is reporting too many I/O errors. Make sure both the source and destination devices are healthy and accessible.	<p>This incident might be seen due to I/O errors on data disks on either the Source or the Destination device.</p> <ul style="list-style-type: none"> • Check to see whether the source and destination devices are accessible over the network. • Check to see whether disks exist in the DIAGNOSTIC state on either of these devices. <p>The disks can be resumed first and if they get quarantined again, they might need to be failed. See the “DIAGNOSTIC” on page 91 section on how to resume or fail a Quarantined drive.</p> <p>If a quarantined disk is failed on the source device, the data slices on this disk can never be evacuated. These data slices must be rebuilt later.</p> <p>If a quarantined disk is failed on the destination device, the data slices that are already evacuated to it might be lost. These data slices must be rebuilt later.</p> <p>If no action is taken, data evacuation continues to until no further progress can be made due to errors. In this case, the next incident will appear.</p> <p>If the problem persists, contact IBM customer support.</p>
Data evacuation is not progressing due to I/O errors. Make sure both the source and destination devices are healthy and accessible.	<p>Follow all the steps from the previous incident.</p> <p>If this incident appears without any quarantined disks on either the Source or the Destination device, it can be due to low error ratio that doesn’t cause disks to become quarantined. In this case, data evacuation continues to until no further progress can be made. It can result in halting the evacuation. The operator can terminate the data evacuation and let Rebuilder rebuild the remaining data slices.</p> <p>Attention: Terminating the data evacuation results in a loss of any data slices that are not evacuated. If an evacuation is terminated, wait for the lost data slices to be rebuilt before you attempt another data evacuation.</p>

SMC pools

Storage Metrics Center Pools allow the user to process, calculate, and publish usage data for Service Providers offering IBM COS STaaS. The metrics can be aggregated on a per Container or Storage Account basis. The tenant usage metrics are collected from the Access Layer in the form of access logs, which are processed and made available to Service Providers via a Cassandra Cluster. The SMC Controller is a stateless entity but relies on the distributed Cassandra Database for workflow management around fault tolerance, error handling, and high availability.

Creating an SMC Pool

Procedure

1. Enable Container Mode.
 - a. Click the **Configure** tab.
 - b. In the **Configure Container Mode** section, click **Configure**.
2. In the SMC Pool section, click **Create SMC Pool**.

3. In the General section, enter a **Name**, **Report Granularity**, **Replication Factor**, and **Access log rotation time**.
4. In the Devices section, select the devices for the SMC Pool.
5. Click **Save**.

Editing an SMC Pool

View detailed information or change configurable SMC Pool parameters.

About this task

Report Granularity is the time duration over which SMC rolls up and aggregates the billing metrics.

Replication Factor is the number of copies of billing metrics that are spread across the devices in this Pool (default = 3). The value cannot exceed the number of devices in the pool.

Access Log Rotation Time is how often access logs are rotated. If "Not Set", the rotation occurs when the access log reaches 500 MB. Otherwise, the logs rotate based on the selected time interval. The upload of the access log occurs within 1 hour of the log rotation.

Procedure

1. Click the **Configure** tab.
2. Click the arrow beside **SMC Pools** in the **Navigation Tree** to expand the SMC Pools subtree.
3. Click an SMC Pool.
4. Click **Change**.
5. Edit the SMC Pool name, description, or access log rotation time and click **Update**.
6. Click **Delete SMC Pool** to remove the pool from the system. A password prompt is needed since delete is an unrecoverable action.
Attention: Each SMC device must be reimaged after the pool is deleted. Deleting an SMC pool is a permanent action with no possibility of a reversal.
7. Click **Add Device** or **Remove Device** to modify number of SMC devices.

Adding an SMC device to an SMC Pool

Procedure

1. Click the **Configure** tab.
2. In the Navigation Tree, click the arrow beside **SMC Pools** to expand the SMC Pools subtree.
3. Click an SMC Pool.
4. In the Devices section, click **Add Device**.
5. Enter the system administrator password.
6. In the Devices section, select the device that you want to add from the list of SMC devices that are not already in the SMC Pool.

Note: It is imperative that the device that is added has never been in an SMC Pool. Adding a device alters the distribution of metrics data.

7. Click **Add Device**.

Monitoring an SMC device

Procedure

1. Click the **Monitor** tab.
2. In the Navigation Tree, click the arrow beside **Devices** to expand the Devices subtree.

3. Click the arrow next to **SMC** to expand the SMC Devices subtree.
4. Click an SMC device.
5. Under **Performance**, in the **Other** section, view the performance graphs for the SMC device:
 - Access Log Transaction Processed
 - Access Log Queue Length
 - Cassandra Disk Space Utilization

Editing an SMC device

Procedure

1. Click the **Configure** tab.
2. In the Navigation Tree, click the arrow beside **Devices** to expand the Devices subtree.
3. Click the arrow next to **SMC** to expand the SMC Devices subtree.
4. Click an SMC device.
5. Click **Change**.
6. Edit the alias and description, or assign a new site or tags.

Note: To change the seed node status of the device, contact IBM Customer Support.

7. Click **Update**.

Removing an SMC device from an SMC pool

Procedure

1. Click the **Configure** tab.
2. In the Navigation Tree, click the arrow beside **SMC Pools** to expand the SMC Pools subtree.
3. Click an SMC Pool.
4. In the Devices section, click **Remove Device**.
5. Ensure the **nodetool decommission** command is completed successfully on the device you want to remove.
6. Enter the system administrator password.
7. In the Devices section, select the device that you want to remove.
8. Click **Remove Device**.

Deleting an SMC Pool

Procedure

1. Click the **Configure** tab.
2. In the Navigation Tree, click the arrow next to **SMC Pools** to expand the SMC Pools subtree.
3. Click an SMC Pool.
4. Click **Delete SMC Pool**.
5. Ensure all the devices in the SMC Pool were reimaged before you add them to any new or different SMC Pools.
6. Enter the system administrator password.
7. Click **Delete**.

Vaults

Overview

Four types of vaults can be created: *management*, *standard*, *container*, and *service* vaults.

- Management vaults retain statistics data that is generated within the system.
- Standard vaults are used to store user data.
- Container vaults are used when the system is in container mode. Containers are created in container vaults. User data is stored within containers.
- Service vaults are needed when the system is operating in container mode. The service vault is used internally by the system to store container metadata, storage accounts, and access keys.

CAUTION:

It is recommended that periodic backups be performed, particularly any time a device is added or a vault is changed. Otherwise, permanent data loss can occur if the Manager database becomes corrupted.

Vault limits

There is a limit to the number of vaults the system can contain. Standard and container vaults count towards this limit, but management and service vaults do not.

The maximum number of vaults is limited to 1000 by default. In some cases, certain system configurations can support up to 1500 vaults. For more information on increasing the vault limit beyond 1000 vaults, contact IBM Customer Support.

When the vault count exceeds 95 percent of the vault limit, a warning message appears on the home page with the current number of vaults.

The system is nearing the maximum number of 1000 standard vaults (996 are present).
Once the limit is reached, standard vault creation requests will be rejected.

When the vault limit is reached, more vault creation requests are rejected.

Filtering vaults in the vault summary section

On the **Monitor** and **Configure Vault Summary** pages, a gray box encapsulates a collection of filters by which the list of Vaults can be limited.

Note: When Container Mode is enabled, container vaults and service vaults do not appear in the list of all vaults, as access cannot be set on these vaults through the Manager Web Interface.

Vault filters include the following selection options:

- Check boxes for Health and Conversion Compatibility.
- Drop-down menus for **Storage Pools**, **Vault Types**, **IDAs**, **Access Pools**, **Tags**, **Versioning**, and **Protection Level**.

Standard vaults come in two flavors as determined by the **Protection Level** setting:

Retention

The vault is a protected vault and objects stored within that vault will have retention properties that will dictate when these objects can be deleted.

- Drop-down menus for Provisioning Codes, Regions, and Storage Classes if Container vaults exist.

Below this gray box a plain text **Search Results** field exists.

To filter the list of Vaults, you can select one or more of any of the options in the drop-down menus or check boxes.

The **Search Results** field accepts one or more terms that are separated with a space. It does not support **AND**, **OR**, or **None** logical operators or keywords.

All filters are combined when you filter the list ((logical)**AND**). Within each filter, if more than one option is selected, any of those selected options can be included when you filter the list ((logical)**OR**).

To limit the Vaults that are displayed to the ones with:

- Healthy or Warning states.
- Storage Pools **ABC** or **DEF**.
- Versioning enabled.
- The terms **Seattle** and **S3** in their names.

Perform the following steps:

1. Check **Healthy** and **Warning** in the **Health** check boxes.
2. Select **ABC** and **DEF** in the **Storage Pool** drop-down menu.
3. Check **enabled Versioning** check box.
4. Type **Seattle S3** in the **Search Results** field.

It provides the wanted filtered result.

Management vaults

A Management Vault stores statistics data a device collects locally. The statistics that are collected to populate the graphs within the Manager Web Interface are important to identify long-term trends in system performance and planning.

The Management Vaults also store Access Logs. It enables the Access Log to be stored indefinitely whereas the number of stored Access Logs is limited based on OS disk size. By default, Access Logs are copied to the Management Vault. It is recommend that this setting not be changed for Accessor devices that host Protected Vaults.

Within 60 minutes of the current Access Log being rotated, it is uploaded to the Management Vault. When the Management Vault is configured initially, the backlog of rotated Access Logs is uploaded slowly into the Management Vault. It is expected to take a number of hours to complete.

The Access logs can be queried from the Management Vault by using the Cloud Storage Object API. If the system is in container mode, these requests must be executed through the Service API. The access logs are stored at <device_uuid>/accessLogs/. The <device_uuid> can be located in configureDevice.adm of the appropriate device. The individual Access Logs are named access.log-YYYY-MM-DDTHHMMSSMMM.gz. The time stamp indicates the time that this Access Log was rotated.

A log file that is named access.log-2014-12-08T182701330.gz was from 08 December 2014 at 18:27:01.330.

In releases prior to ClevOS 3.10.0, management vaults are optional. In a new installation of ClevOs 3.10.0 or newer, management vaults are enabled and set to automatic configuration at installation. If you are upgrading to 3.10.0 or newer from a previous version of ClevOS, the system preserves the management vault settings (enabled or disabled) of the previous release.

Note: Before you enable a Management Vault, checks should be performed on storage pools. Do not enable a management vault if the following conditions are true for storage pools:

1. Contain many vaults (hundreds or more).
2. Are in operation for a long time (> 2 years).
3. Are near 100% usage.
4. Have sensitive performance concerns.
5. Are overly expensive (might want to use a cheaper storage pool).

If unsure, contact IBM Customer Support.

Management Vaults can be configured through the **Configure** tab of the Manager Web Interface. From the **Configure Management Vault** page, a user can:

- Enable or disable the ability to configure Management Vaults in the system.
- Configure the access log rotation period. If **Not Set**, access logs rotate based on file size. Otherwise, access logs rotate based on the time period selected.
- Choose between automatic and manual management configuration.

Note: Automatic configuration is recommended.

The Device Management Vault Configuration displays storage pools what contain devices.

You can choose what content is stored in a Management Vault. The Management Vault options are:

- System configuration
- Backup of HTTP access logs.
 - Remove client IP addresses. If **Backup HTTP access logs** is selected, **Remove client IP addresses** can also be selected. If **Remove client IP addresses** is selected, a hyperlink to view the redaction progress displays.
- Backup of device statistics

Note: When automatic configuration is selected, all Management Vault options are selected.

When automatic configuration is selected, the Manager application performs the following items:

- Creates a single Management Vault for each storage pool in the system that has a vault.
- Updates the Management Vault of all Slicestor devices in a storage pool to the vault created in that pool.
- Assigns Accesser devices to a Management Vault.
- Assigns Manager device to the Management Vault with the least number of devices.

CAUTION:

If a Slicestor device is replaced while in a pool, then the original device no longer has an assigned Management Vault and the new device points to the Management Vault.

Automatic configuration follows these rules:

- New Slicestor devices are not assigned a Management Vault.
- New Accesser devices are not assigned a Management Vault until they are deployed to a vault. At that point, the Management Vault on that pool is assigned.
- New Management Vaults are created on a storage pool when the first vault on that pool is created. The Management Vault that is created has the same configuration (width, threshold, secure slice, and so on) as the new vault on the pool.
- A storage pool with Management Vaults can be deleted if the pool contains Management Vaults.
- Management Vaults cannot be created, edited, or deleted. Management Vaults for a device cannot be assigned or removed manually.

Note: A Management Vault is not created on a storage pool that has no vaults.

The name and description for the Management Vault is generated automatically. The name is in the format `csinternal-mgmt-poolName` with a description of "Stores backups of internal statistic data".

Manual configuration has different restrictions:

- Management Vaults must be manually created/edited/deleted and assigned to devices.
- Management Vaults are created through the API or the `/configureManagementVault` page in the Manager Web Interface.
- Device management vaults can be managed by using a new **Device Management Vault Configuration** form.

At any time, a user can disable management vaults in the Manager. It does not delete any pre-existing Management Vaults but prevents any new data from being written to them.

CAUTION:

For a device that is running 3.4, with Management Vault enabled, statistics data will be retained after an OS disk replacement. If the device is imaged to an earlier release, like 3.2.x, the statistics for that device are lost.

Standard vaults

After the system is set up and a storage pool is available, standard vaults can be created. To use Standard vaults, the system must be operating in vault mode.

As part of the planning phase, a determination needs to be made on the vault width and threshold. These decisions affect the availability, resiliency, performance, and storage capacity of the vault. These settings should be used to adjust the characteristics of each vault as wanted for the specific application.

Vaults can be created by:

1. Customizing a configuration
2. Leveraging a vault template

In addition, vault access can be restricted based on IP addresses.

Note: It is recommended that periodic backups be performed, particularly any time a device is added or a vault is changed. Otherwise, permanent data loss can occur if the Manager database becomes corrupted.

Related concepts:

"Vault limits" on page 27

There is a limit to the number of vaults the system can contain. Standard and container vaults count towards this limit, but management and service vaults do not.

Container vaults

Container vaults are used instead of standard vaults when the system is running in container mode. As such, container vaults have the same restrictions as standard vaults and they can be created by using the same method.

Containers are created on container vaults. Object I/O is executed on objects that are stored within containers. Direct I/O to container vaults is not permitted.

Related concepts:

"Vault limits" on page 27

There is a limit to the number of vaults the system can contain. Standard and container vaults count towards this limit, but management and service vaults do not.

Service vaults

The service vault is needed when the system is operating in container mode. The service vault is created while container mode is being enabled and cannot be deleted until container mode is disabled. Only a single service vault can be created.

The service vault stores a mapping of container names to container vault, storage accounts, and access keys. Therefore, all container I/O fails when the service vault is not accessible from the access devices. If the content of the service vault is lost, the containers are lost as well.

Configuring vault protection

Vault Protection allows objects stored in vaults to have associated deletion protection. Protected objects cannot be deleted until the associated data retention duration expires and all legal holds associated with the object are been removed.

Procedure

1. Click the **Configure** tab.
2. In the **Configure Vault Protection** section, click **Configure**.
3. Click the **Enable Vault Protection** check box.
4. Click **Update**.

Creating vaults

To begin vault creation, follow these steps.

Procedure

1. In the **Configure** tab, click **Create Vault** in the **Summary** section.
Create Vault is also available by right-clicking on any item in the navigation tree.
2. If a storage pool was not created already, the **Create New Standard Vault** page appears.
3. In the **General** section, complete these fields:

Field Label	Acceptable Field Value
Name	Each vault must be uniquely named (maximum of 255 characters); this name is used by the Manager for all references to this vault. Vault names can include underscores and alphanumeric characters. The vault name can also contain periods (.), but the name cannot start or end with a period or contain more than one period in a row. The first character of the name must be a letter, underscore, or number.
Description	An optional free-form description can also be entered. Information that you might include in the description field might be initiator host name and IP address, names and phone numbers of administrators, and key users of the vault.
Tags	Tags can be created and or assigned to a vault before the vault is created. For more information, see the Tags section.
Provisioning Code	During container creation, specify a unique provisioning code with a maximum of 255 characters to indicate in which container vault the container should be created. The default value is the vault name, once entered. The value is usually the locationConstraint or region . If the provisioning code is not set, containers can be created in this new container vault if it is configured as the default container vault for an access pool.

Field Label	Acceptable Field Value
Region	Optional. A region can be provided with a maximum of 255 characters to indicate where the contents of this vault reside. The locationConstraint parameter, which is shown for containers that are associated with this vault in the S3 GET Service Extended and S3 GET Bucket Location APIs, are populated with the value set for region.
Storage Class	Optional. A storage class can be provided with a maximum of 255 characters to assign a classification to all objects stored within this vault. The header x-amz-storage-class shown in the S3 GET/HEAD Object and the StorageClass in the response body of the S3 GET Bucket are populated with the value set for storage class.

4. In the **Configuration** section, several options display.

- When the width of the pool for this vault is greater than 6, complete the following fields:

Field Label	Acceptable Field Value
Width	<p>This setting is referred to as the width of the vault and corresponds to the number of slices into which all data in the vault is split.</p> <p>Vault width must be a factor of the storage pool width. The Manager Web Interface allows any vault width greater than or equal to 6 and less than or equal to 60.</p>
Threshold	<p>The minimum number of slices that must be available to perform a read. Pre-defined, supported thresholds are presented when the drop-down list is clicked. The vault threshold, always less than the width, determines the reliability of the vault. If the set of available Slicestor devices is such that the number of slices falls below this threshold, the vault content cannot be read, and the vault appears as red in the Monitor application.</p> <p>The Manager Web Interface allows any value between 1 and Vault Width, inclusive.</p> <p>If the vault is on a storage pool that spans multiple sites, the Manager Web Interface warns the user if the selected threshold is high enough such that a single site outage affects read and write availability.</p>
Write Threshold	<p>The Manager Web Interface allows any value such that all the following are true:</p> <ul style="list-style-type: none"> • Write Threshold > Threshold. <p>CAUTION: Write Threshold = Threshold is allowed if Threshold = Vault Width or if Vault Width < 6.</p> <ul style="list-style-type: none"> • Write Threshold ≤ Vault Width. • (Write Threshold + Threshold) > Vault Width. <p>Write Threshold defaults to Threshold + 2, if that is within the allowed range. Otherwise, the selected Write Threshold is the halfway point between the minimum allowed Write Threshold and Vault Width, rounded up. This value is selected by default in the Write Threshold drop-down when Threshold is selected. This value is also used as the Write Threshold when a vault is created through the Manager REST API and a Write Threshold is not specified.</p> <p>If the vault is on a storage pool that spans multiple sites, the Manager Web Interface warns the user if the selected write threshold is high enough such that a single site outage affects write availability.</p>

Field Label	Acceptable Field Value
Alert Level	Optional. If the set of available Slicestor devices is such that the number of slices is between the write threshold and the alert level exclusive, the vault icon is yellow in the Monitor application. In this case, the vault is still fully functional.

Attention: If the **Threshold** is set such that the loss of one site would make Vaults either unusable or read-only, the Manager Web Interface displays a confirmation dialog box that asks the operator if they accept the settings with the risks they present.

If only the **Threshold** causes an issue:

Warning: This IDA configuration is susceptible to read availability issues during a single site outage. Do you still wish to continue?

If the **Threshold** and **Write Threshold** cause issues:

Warning: This IDA configuration is susceptible to read and write availability issues during a single site outage. Do you still wish to continue?

Click **Cancel** or **OK** to change or keep the settings.

- When the width of the pool for this vault is 3 - 7, select a vault optimization to create a Concentrated Dispersal vault:

Table 6. Concentrated Dispersal vault optimization configuration

Field Label	Width	Description
Storage Efficiency	3 - 7 Note: Contact Customer Support to enable the creation of a 7-wide Concentrated Dispersal device set.	More usable capacity with reasonable performance.
Performance	3 - 6	Better performance with less usable capacity.

Note: When you choose a vault optimization, it cannot be changed later.

5. If you enabled vault protection on the system, choose a **Protection** setting.

Note: This section displays only if Vault Protection Configuration is enabled in the **Configure** tab.

- **Disabled.** The vault does not have a protection levelRetention set.
- **Retention:** When you select the Retention protection level, data is retained for a default duration of time, unless you specify a custom duration. After you create the vault, you can modify the retention time settings, but you cannot change the protection level.

In the **Data Retention Durations** section, specify the following values or accept the default values:

- **Default Duration:** The default retention period (in days) for an object in this vault. The minimum supported value is 0 days, and the maximum is 36159 days. The default value is 730 days.
- **Minimum Duration:** The minimum retention period (in days) for an object in this vault. The minimum supported value is 0 days, and the maximum is 36159 days. The default value is 365 days.
- **Maximum Duration:** The maximum retention period (in days) for an object in this vault. The minimum supported value is 0 days, and the maximum is 36159 days. The default value is 7305 days.

6. In the **Options** section, complete these fields:

Field Label	Description
Enable SecureSlice™ Technology.	Optional. SecureSlice™ provides extra encryption benefits that are combined with dispersal. This box is checked by default for newly created vaults. This feature can be cleared, although it is not recommended. If it is cleared, a warning message appears, and a confirmation is needed before proceeding. When a vault is created, the SecureSlice™ option cannot be changed.
Enable Versioning.	Check to enable versioning on this vault. Note: Versioning cannot be enabled if the Protection Level is set to Retention .
Delete Restricted.	This feature allows Security Officers to restrict vault access permissions such that users with write access to the vault are not able to delete objects from the vault. Additionally, object versioning is enabled by default so that existing content is preserved upon overwrite when using a write-by-name interface. Users that are granted owner permissions on a Delete Restricted vault are allowed to delete objects and versions.

7. In the **Quotas** section, complete these optional fields if wanted:

Field	Value
Soft Quota	Optional. If wanted, select a value for a soft quota. A notification is sent to the Event Console , if the soft quota setting is exceeded. It does not cause restrictions to usage. Setting the quota higher than the total space available in one or more storage pools that are associated with this vault has no effect.
Hard Quota	Optional. If wanted, enter a hard quota value. The Accesser device (or application) does not permit the user to exceed the hard quota value for this vault. A notification is also sent to the Event Console if the hard quota setting is exceeded. Setting the quota higher than the total space available in one or more storage pools that are associated with this vault has no effect.

8. In the **Advanced Index Settings** section, **Name Index Enabled** is checked by default for Standard vaults and you can enable Recovery Listing.

Field	Value
Name Index Enabled	<p>Enabled by default. When enabled, Name Index allows a user to list contents of a vault in lexicographical order based on the object's name, or key. The Name Index is updated whenever objects are added or removed from a vault.</p> <p>The Name Index must be enabled to provide prefix-based listing and sorted listing results for named object vaults. Changing this option requires service to restart Accesser devices before release 3.4 to take effect.</p> <p>If you disable Name Index, you can re-enable it only by contacting Customer Support.</p> <p>Note: Name Index cannot be disabled for Protected Vaults.</p>
Recovery Listing Enabled	Recovery Listing allows for limited listing capability even when the contents of a vault are not indexed. When enabled, Recovery Listing lists the SourceNames of the metadata headers. Recovery Listing is slower than the Name Index listing and the results are not sorted. Recovery Listing can be used to list contents of a vault for which Name Index is corrupted or not enabled.

CAUTION:

If both Name Index and Recovery Listing are enabled, the Recovery Listing settings take precedence over Name Index settings. It means that a user receives a Recovery Listing response for a listing request. Applications that expect a Name Index listing might produce errors.

9. Optional: In the **Notification Service** section, select a **Configuration**. For more information, see “Configuring notifications” on page 81.

Note: You cannot enable notifications on container vaults, mirrored vaults, vault proxies, or vaults that are migrating data. Once notifications are enabled, this vault cannot be used in a mirror, for data migration, or with a vault proxy.

Note: Notifications are sent only for new operations that occur after the vault is assigned to the configuration.

- a. Select the topic to which you want to send notifications.
 - **Default:** The default topic specified in the configuration.
 - **Custom:** Enter a topic name to override the default topic specified in the configuration.
10. The Access Pools available for deployment are displayed on the Vault Summary section. This step is not necessary for Simple Object vaults that are accessed through the Accesser application.
11. Click **Save**.

Create vaults by using vault templates

An alternative approach to vault creation is based on the use of vault templates.

These templates allow a user to create multiple vaults with the same parameters quickly and enable common vault configurations to be leveraged across multiple users. A vault template is created on a storage pool and can then be used during the creation of a vault. All parameters that are set in the vault template apply to the vault. When a storage pool is created, the following steps can be performed to create a vault template:

1. Click the **Configure** tab and select a storage pool.
2. From the **Vault Templates** section, click **Create Vault Template**. A new page appears.
3. A new vault template can be created by either importing an existing template or customizing a new template.
4. When customizing a new template, you can enter an optional provisioning code (for example, US East). This code determines where data is stored and which vaults are written when using vault provisioning. The provisioning code must be unique for each vault (255 characters max) and defaults to the vault name, once entered. This parameter is typically the S3 Put Bucket **locationConstraint** or **region**. In container mode, PUT bucket requests create containers instead of vaults. Therefore, the provisioning code on the vault template is only supported in vault mode.
 - In the Provisioning Code field, you can enter a region (255 characters max). This region is used to indicate where the contents of the vault reside. The **locationConstraint** shown for the container(s) associated with the vault in the S3 GET Service Extended and S3 GET Bucket Location APIs are populated with this value. The region code on the vault template is only supported in vault mode.
 - In the Provisioning Code field, you can enter a storage class (255 characters max). The storage class is used as a classification assigned to all objects stored within the vault. The header **x-amz-storage-class** shown in the S3 GET/HEAD object and the **storageClass** in the response body of the S3 GET Bucket are populated with this value. The storage class on the vault template is only supported in vault mode.
5. In addition to providing vault configuration parameters, Accesser devices can be selected for deployment and access to vaults can be controlled by specifying IP addresses. For more information on these configuration parameters, see “Creating vaults” on page 31
6. When complete, click **Save**.

Note:

(Vault Mode Only) If one or more templates exist or are created, it is possible to specify a default template (**Configure > Default Vault Template Configuration**). This default template is used when the Provisioning API is enabled and no provisioning code is specified. Vault templates with **Retention** enabled cannot be used as a default template.

A vault can be created by using the newly created template.

1. On the **Storage Pool: {storage-pool-name}** page, the new template appears under the **Vault Templates** section.
2. Click **Create Vault** next to the template to be used.
3. Provide the needed information. The newly created vault inherits the SecureSlice state (enabled or disabled) from the **Vault Template**. It cannot be changed for the vault after it is created.
4. Click **Save**.

Vault templates can be edited and deleted as well:

1. On the **Storage Pool: {storage-pool-name}** page, in the **Vault Templates** section, select the vault template to be edited/deleted, and a new page appears.
2. Click **Change** or **Delete Vault Template** from the action bar at the top of the page.
3. Edit the template, make any revisions, and click **Save**.

Vault proxy settings

The S3 proxy allows transparent access to data stored in a S3 bucket as though it was stored in an IBM vault.

Only supported in vault mode. Protected Vaults do not support Vault Proxy.

By clicking **Change** on the action bar, a user can configure proxy settings on a vault. Multiple proxy types are available to select.

A proxy redirect may be configured for S3 data migration or for another vault on this system:

1. Select the **Proxy Type**, either **S3** or **Internal**.
2. Enter the **Endpoint URL** and the **Bucket Name**.
An optional **Access Key ID** and **Secret Access Key** may also be required.
3. Select the vault to proxy.

Ensure that the vault is deployed to one or more access pools and there is at least one common authorized user or group between the proxied vault and the vault configuring the proxy.

Note: When using an internal proxy, failure to configure common access permissions for both vaults could result in I/O errors.

The system will attempt to read the object from this vault, and will redirect the Read to the proxy location if the object is not found (i.e. it has not been migrated). This feature does not migrate data.

At any point, you can disable the proxy setting by clicking the disable radio button on the Manager Web Interface.

Vault security

Vaults can be protected from unauthorized mounts and accessed through an IP filtering access control list and assigned vault users.

Limit vault access by IP

CAUTION:

For unrestricted access to a vault, do not perform these steps.

1. Click **Change** in the **Access Control** section of the **Vault Configuration** page.
2. Enter the clients that are allowed access to the mirror in the **Authorized IP Addresses** text field. These values, which are separated by a space or a comma, can be either:
 - Individual host names
 - Individual IP addresses (in IPv4 or IPv6 format)
 - Classless InterDomain Routing (CIDR) Notation for a range of clients
3. Click **Update** to apply the changes.

Note: When you remove an IP address from the list, changes do not take effect until a new session is started.

Classless InterDomain Routing (CIDR) Notation

CIDR allocates IP addresses and routes Internet Protocol packets.

The notation shows a compact representation of an IP address and its routing prefix.

The CIDR number comes from the number of 1s in the subnet mask when converted to binary.

192.168.100.0/24 represents the IPv4 address and its associated routing prefix 192.168.100.0. Its subnet mask is 255.255.255.0 or 11111111.11111111.11111111.00000000 in binary. It equals 24 ones, or "/24 (pronounced "slash twenty-four"). fe00:d7::6/64 represents an IPv6 address and its associated routing prefix fe00:d7::6. Its subnet mask is /64.

Limit access by user

This method is applicable for vault mode only. In container mode, user access to containers is controlled via container configuration, which can be modified via Cloud Object Storage requests.

If you are logged in with **Super User** or **Security** roles, you can assign vault access here or from the **Security** tab.

Note: Vaults can be protected from unauthorized mounts and accessed through an IP filtering access control list and assigned vault users.

Vault access permissions must be entered for each user of an object vault. These permissions can be changed later.

Anonymous users can access a vault by using 1 of three levels of permission:

- R/W (Read/Write/Delete)

Note: Anonymous R/W/D is not supported on Protected Vaults.

- R (Read Only)
- None (no access, the default)

Specified users can access a vault by using 1 of three levels of permission:

- Owner (Read/Write/Delete)
- R/W (Read/Write/Delete)
- R (Read Only)
- None (no access, default), unless created via API

The Owner option is inherited when a vault is created through the (provisioning API). It can also be assigned by an account that has the Security Officer role. The Owner option must be set to delete a vault.

Note: For more information, see “Configure provisioning API” on page 70.

Tags

To bucket objects of similar purpose, one can assign a specific *tag* to one or more vaults or devices.

Creating a tag

Follow these steps to create a tag.

Procedure

1. Click the **Configure** tab to display the top-level **Configure** page.
2. Click **Configure Tags**.
3. Enter the name of the new tag in the **Tag Quick Creation** field.
4. Click **Create Tag** in the upper right corner of the window.

Note: Tags can also be created from when you Create Vaults, Edit Vaults, or Edit Devices.

Assigning a tag

Follow these steps to assign a tag to a vault/device.

Procedure

1. Click the **Configure** tab to display the top-level **Configure** page.
2. Click **Configure Tags**.
3. Click the tag that you want to assign to a vault/device.
4. Search for the vaults/devices that have the tag that is assigned to it, under the **Tag Association** sub section. You can search for specific vaults/devices by using the sorting options.
5. Select one or more vaults/devices to be assigned to the tag and click **Update** on the upper right portion of the window.

Editing a tag

Follow these steps to edit a tag.

Procedure

1. Click the **Configure** tab to display the top-level **Configure** page.
2. Click **Configure Tags**.
3. Click the tag that you want to modify.
4. Change the name of the tag if wanted.
5. Change the description of the tag if wanted.
6. Click **Update** in the upper right corner of the window.

Deleting a tag

Follow these steps to delete a tag.

Procedure

1. Click the **Configure** tab to display the top-level **Configure** page.
2. Click **Configure Tags**.
3. Click the tag that you want to delete.
4. Click **Delete Tag** in the upper right corner of the window. The Manager Web Interface then asks for confirmation.
5. Click **Delete** in the upper right corner of the window to delete the tag.

Vault data migration

Data can be copied from one standard vault to another standard vault by using vault migration. Vault migration is not supported on vaults with the Retention protection level enabled, or on service, container, and management vaults.

Configuring a vault migration

Follow these steps to configure a Vault data migration.

Before you begin

Before you configure a Vault data migration, an operator must perform the following item:

- Upgrade all Slicestor[®] Nodes in the specified Vaults Storage Pools to 3.7.1 or later.

Procedure

1. Navigate to a Vault to be migrated.
2. Click **Setup Vault Migration** in the **Vault** action bar.
3. Select the appropriate option for this migration:
 - a. Migrate data from this specified Vault to another Vault.
 - b. Migrate data to this specified Vault from another Vault.
4. Make a choice, depending on the selection made in the previous step:
 - a. If this Vault is intended as the source, select whether to create the destination Vault or choose an existing Vault from the drop-down.
 - b. If this Vault is intended as the destination, choose a source Vault from the drop-down.
5. Check the **Rename Destination Vault** check box, if applicable to this migration. Selecting this option renames the destination Vault with the source Vaults name. It also applies the source Vaults Access Pool deployments, allowed IP addresses, and authorized user permissions to the destination. This option is used when the destination Vault is intended to replace the source Vault and a continuous I/O to the source vault exists.
6. If you are configuring a migration between two existing Vaults, click **Finish Setup** to save your changes. If you are configuring a migration to a new Vault, click **Proceed to Create Vault**.
7. If it is a migration between two existing Vaults, skip to the “Monitoring a vault migration” section. Otherwise continue with the following steps.
8. Create the destination Vault by using the same procedure as standard Vault creation.

Note: If the **Rename Destination Vault** option was checked, the **Vault Name** field is read-only and completed with the source Vault's name.

Monitoring a vault migration

Follow these steps to execute a Vault data migration:

About this task

After a migration is configured, the next step is to start and monitor the migration.

Procedure

1. Navigate to the destination Vault configured in the migration. For the first few minutes, **Start** is unavailable. It is done to ensure that the necessary configuration changes are propagated to all devices.
2. When **Start** is available, click it to begin the migration.
3. During the migration the following options are available:
 - **Pause** the migration.

- **Change** the migration throttle. The migration rate can be throttled by MB/sec and objects/sec.
 - **Abort** the migration. It halts progress and leave all data copied to the destination vault in its current location. Abort can be executed at any point during the migration before completion.
4. Upon start, the migration begins scanning all objects on the source Vault to determine how many objects to migrate.
 5. After the migration is finished scanning, a progress bar showing the percentage of objects that are migrated appears.
 6. The migration is finished when the progress bar reaches 100%. The UI reflects this completion.

What to do next

The migration might encounter a scenario where scanning or migration is blocked. It can occur for various reasons (for example, one of the Vaults is below threshold).

When blocked, the migration pauses and waits for user intervention. Click **Resume** to continue the migration.

During the migration, some objects can fail to migrate. The migration attempts to migrate all objects and track the ones that fail to copy. If failed objects exist at the end of the migration, the UI presents the user with two options:

- **Retry** the failed objects. It attempts to retry any objects that failed. If any objects fail again, the same UI widget appears.
- **Complete with Failures.** Click to mark the migration as done and no attempt is made to migrate the failed objects again.

All active vault migrations can be viewed on the **Vault Summary** page. A **Vault Migrations** tab appears when active migrations exist in the system.

The migrations can be filtered by Migration Status, source and destination vault health, and source and destination Storage Pool. After a migration is complete, the record appears in the summary table for 72 hours. After that time, the migration will not be shown.

Vault mirrors

Overview

To support two site configurations, the concept of a mirror is introduced, which links a vault at each site together.

Note: Only standard vaults can be part of mirrors, and only retention vaults can be part of protected mirrors.

CAUTION:

The number of vault mirrors supported on a system is limited to the number of vaults present when the vault mirror is created. If the system has a vault limit of 1000 and currently contains 999 vaults, then the mirror can be created if an initial vault is provided. If the system has a vault limit of 1000 and currently contains 1000 vaults, then the vault mirror creation is rejected regardless of whether an initial vault is provided.

Filter vault mirrors in the vault summary section

On the **Monitor** and **Configure Mirror Summary** pages, a gray box encapsulates a collection of filters by which the list of Vault Mirrors can be limited.

These filters include:

- Drop-down menus for Access Pools and Storage Pools.

- Check boxes for Health.

Below this gray box also exists a plain text **Search Results** field.

To filter the list of Vault Mirrors, you can select one or more of any of the options in the drop-down menus or check boxes.

The **Search Results** field accepts one or more terms that are separated with a space. It does not support AND, OR, or None logical operators or keywords.

All filters are combined when you filter the list ((logical)AND). Within each filter, if more than one option is selected, any of those selected options can be included when you filter the list ((logical)OR).

To limit the Vault Mirrors displayed to the ones with:

- Healthy or Warning states.
- Storage Pools ABC or DEF.
- The terms Seattle and S3 in their names.

Perform the following steps:

- Check **Healthy** and **Warning** in the **Health** check boxes.
- Select **ABC** and **DEF** in the **Storage Pool** drop-down menu.
- Type **Seattle S3** in the **Search Results** field.

The wanted filtered result is provided.

Creating a vault mirror

About this task

CAUTION:

The individual vaults can differ in settings (such as IDA and SecureSlice). The mirror displays data that is stored in both underlying vaults with two caveats:

- Usable Capacity is limited to the size of the smaller vault.
- Used Capacity is reported as the greater used capacity of the vaults.

It is highly recommended that the storage pools are equal usable capacity, as any "extra" capacity in one or the other storage pools is unusable by the vault mirror.

Attention: Vault mirrors cannot be created if the system has reached the vault limit.

If you want to create a protected mirror, see “Creating a protected vault mirror” on page 45.

Procedure

1. Navigate to the **Configure** tab.
2. Click **Create Mirror** under the **Summary** heading. The Create Mirror page appears.
3. Enter a name for the vault mirror in the **Name** field.
4. Enter an optional description in the **Description** field.
5. Click the **Write** check box, if you would like to enable the Synchronous Write Mode feature.
6. If a vault exists that should serve as the first vault in the vault mirror, select that vault from the **Vault** drop-down menu.

There are two primary use cases associated with seeding a vault mirror with an existing vault:

- a. Migrate an existing vault to be part of a persistent vault mirror. This workflow uses the vault mirror synchronization capability to migrate from a single vault to a vault mirror.
- b. Migrate existing vaults to a new vault. This workflow uses the vault mirror synchronization capability to migrate data. The vault mirror is temporary and is deleted after the data is migrated to the new vault.

CAUTION:

To make this transition transparent to the client, a short outage is needed.

- 1) In both use cases, perform the following steps:
 - a) Rename the existing vault to something new.
 - b) Create a vault mirror with the existing vaults original name.
 - c) Create second vault.
 - d) Perform an I/O operation on the vault mirror to initialize the synchronization.
- 2) With the second use case, also perform the following steps:
 - a) Allow the synchronization process to complete.
 - b) Delete the vault mirror.
 - c) Delete the original vault.
 - d) Rename the new vault with the original vaults name.
7. Click **Save** to create the vault mirror. The **Mirror:<Mirror Name>** page displays with an alert that indicates:
Mirror creation successful
8. If an existing vault was used as the basis of a mirror, skip to Step 9, otherwise follow these steps:
 - a. Click **Create Vault 1** to add the first vault to the vault mirror. The Create New Standard Vault page displays.
 - b. Click **Continue** to begin creating the vault.
 - c. Follow the steps in “Creating vaults” on page 31 to create the standard vault.
When the width of the pool for this vault is 3 - 6, you can also choose to enable Extended Availability for a mirror with Concentrated Dispersal vaults. Enable Extended Availability to ensure that you can read and write to the mirror if a Slicestor device outage occurs on one side and the other side is unavailable. However, this reduces usable capacity. The **Mirror:<Mirror Name>** page displays.
9. Click **Create Vault 2** to add the second vault to the vault mirror. The **Create New Standard Vault** page displays.
10. Click **Continue** to begin creating the vault.
11. Follow the steps in “Creating vaults” on page 31 to create the standard vault.
When the width of the pool for this vault is 3 - 6, you can also choose to enable Extended Availability for a mirror with Concentrated Dispersal vaults. Enable Extended Availability to ensure that you can read and write to the mirror if a Slicestor device outage occurs on one side and the other side is unavailable. However, this reduces usable capacity. The vault mirror is now complete.

Creating mirrors by using mirror templates

An alternative approach to mirror creation is based on the use of mirror templates. These templates allow a user to create multiple mirrors with the same parameters quickly and enable common mirror configurations to be leveraged across multiple users.

About this task

A mirror template is created and can then be used during the creation of a mirror and both the vaults that are mirrored. All parameters set in the mirror template apply to the mirror and to the vaults that were created with the mirror. The following steps can be performed to create a mirror template.

If you want to create a protected mirror by using a template, see “Creating protected mirrors by using mirror templates” on page 46.

Procedure

1. Navigate to the top-level **Configure** tab.
2. Select **Configure**, from the **Template Management** section.
3. Click **Create Mirror Template**, from the **Mirror Template** section. A new page appears where a new mirror template and two vault templates can be created.

When customizing, an optional provisioning code can be entered (for example, US East). This code determines where data is stored and which vaults are written when you are using vault provisioning. In addition to providing mirror configuration parameters, Accesser devices can be selected for deployment and access to mirrors can be controlled by specifying IP addresses.

4. Click **Save**.
5. To create a mirror from the template, click **Create Mirror from Template**.
6. Provide the necessary information and click **Save**.

What to do next

Note: It is possible to specify a default template, which can be either a mirror template or a vault template. This default template is used when the Provisioning API is enabled and no provisioning code is specified.

Mirror templates can be edited and deleted as well.

1. From the **Template Management** page, in the **Mirror Templates** section, select the mirror template to be edited/deleted, and a new page appears.
2. Click **Change** or **Delete Mirror Template** from the action bar at the top of the page.
3. Edit the template, providing any revisions, and click **Save**.

Deploying a vault mirror

Writing and reading objects from the mirror is not possible until the mirror is deployed to one or more Accesser devices.

About this task

If you are using Concentrated Dispersal vaults in the mirror, both vaults must be created before the mirror can be deployed. Both vaults must be Concentrated Dispersal vaults with the same Vault Optimization. For more information, see “Creating vaults” on page 31.

Procedure

1. Click **Change** on the **Deployment** action bar.
2. Check the Accesser devices to which the mirror should be deployed.

Note: Protected mirrors cannot be deployed with just one protected vault in the mirror. The primary and secondary vaults cannot be independently deployed to an Accesser device while they are part of a protected mirror.

3. Click **Update**.

Setting vault mirror permissions

Access to the mirror can be restricted to one or more hosts or IP addresses.

See “Vault security” on page 36 for instructions on how to restrict access.

Editing a vault mirror

About this task

Click **Change** at the top of the **Configure Mirror** page to change the mirror name and description. Changes to these values have no effect on the underlying mirrored data. If it is a protected mirror, you can also change the protection level if there are no vaults in the mirror.

In addition, six vault-specific settings can be modified:

- Versioning.
- Delete Restricted.
- Soft Quota.
- Hard Quota.
- Name Index Enabled.
- Recovery Listing Enabled.

Note: In a protected mirror, only Soft Quota, Hard Quota, and the Description can be modified. The protected mirror's retention settings or protection level can only be updated if it is empty and does not have any vaults associated to it.

Changes to these values apply to both vaults in the mirror. To change these values, perform the following steps:

Procedure

1. Click **Change** on the **Vault Settings** action bar.
2. Click **Update** after you make the wanted changes.

Deleting a vault mirror

Follow these steps to remove a mirrored vault configuration.

Procedure

1. Click **Delete Mirror** at the top of the **Configure Mirror** page.
2. Enter the password.
3. Optional: From the **Select a vault** menu, choose one of the mirror's vaults to be renamed and redeployed to match the mirror.
4. Click **Delete**.

Note: Deleting a mirror does not delete the two vaults. Both vaults retain existing objects and their respective current configuration. However, the system no longer synchronizes the content automatically.

CAUTION:

If the mirror is using **Concentrated Dispersal**, proceed with caution when deleting the mirror. The vault IDAs in the mirror will provide a lower level of data reliability and availability when used in a non-mirrored setup.

Results

If you chose a vault to rename and redeploy, the selected vault is renamed to have the same name as the mirror, and is then deployed to the same access pools as the mirror. The vault you did not select still exists, but is not deployed to any access pools.

If you did not choose a vault to rename and redeploy, then both vaults still exist but are not deployed to any access pools.

Monitor a vault mirror

An authorized user can monitor the status of a mirror at any time by going to the **Monitor** tab and clicking the mirror via the left navigation tree.

Creating a protected vault mirror

Before you begin

To create a protected mirror, the system must have Protection enabled. For more information, see "Configuring vault protection" on page 31.

About this task

CAUTION:

The individual vaults can differ in settings (such as IDA and SecureSlice). The mirror displays data that is stored in both underlying vaults with two caveats:

- Usable Capacity is limited to the size of the smaller vault.
- Used Capacity is reported as the greater used capacity of the vaults.

It is highly recommended that the storage pools are equal usable capacity, as any "extra" capacity in one or the other storage pools is unusable by the vault mirror.

Attention: Vault mirrors cannot be created if the system has reached the vault limit.

Versioning, Delete Restricted, Recovery Listing, SSE-C settings, and synchronous action do not apply for mirrored protected vaults.

Procedure

1. Navigate to the **Configure** tab.
2. Click **Create Mirror** under the **Summary** heading. The Create Mirror page appears.
3. Click **Protected Mirror** and then click **Next**.
4. Click **Continue without using Template**.
5. Enter a name for the protected mirror in the **Name** field.
Attention: Once you name the protected mirror, it cannot be renamed later.
6. Enter a description in the **Description** field.
7. Optional: Enable **Restrictive Access Control**. When Restrictive Access Control is enabled, vault access permissions do not automatically provide equivalent object access permissions. Object read, metadata write, and access control updates can only be performed by the owner of the object in a protected vault or any user authorized by the owner. If this setting is not enabled, users with vault permissions inherit equivalent object permissions such as the ability to modify object protection. Once enabled, you cannot disable Restrictive Access Control.
8. If a protected vault exists that should serve as the primary vault in the protected mirror, select that vault from the **Seed vault** menu. You cannot use an existing protected vault that is deployed to an access pool as the primary vault. You must first undeploy the access pool before you can use it as the primary vault.

There are two primary use cases associated with seeding a vault mirror with an existing vault:

- Migrate an existing vault to be part of a persistent vault mirror. This workflow uses the vault mirror synchronization capability to migrate from a single vault to a vault mirror.
- Migrate existing vaults to a new vault. This workflow uses the vault mirror synchronization capability to migrate data. The vault mirror is temporary and is deleted after the data is migrated to the new vault.

CAUTION:

To make this transition transparent to the client, a short outage is needed.

For both use cases, the following generic procedure applies:

- a. Rename the existing vault to something new.
- b. Create a vault mirror with the existing vault's original name.
- c. Create a second vault.
- d. Perform an I/O operation on the vault mirror to initialize the synchronization. Allow the synchronization process to complete.

To migrate existing vaults to a new vault, perform the following additional steps:

- e. Delete the vault mirror.
 - f. Delete the original vault.
 - g. Rename the new vault with the original vault's name.
9. In the Protection section, enter **Data Retention Durations** for the mirror or accept the default values. If you are creating a new protected vault as the primary vault, then the vault inherits the protection level and retention settings of the mirror. If you selected a seed vault, then the seed vault's Data Retention Duration values appear and cannot be changed.
10. Click **Save** to create the protected vault mirror. The **Mirror:<Mirror Name>** page displays with an alert that indicates:
- Mirror setup successful.
11. If an existing protected vault was used as the basis of the protected mirror, skip to step 12, otherwise follow these steps:
- a. Click **Create Vault 1** to add the primary vault to the vault mirror. The Create New Vault page displays.
 - b. Select a storage pool and click **Continue** to begin creating the vault.
 - c. Follow the steps in "Creating vaults" on page 31 to create the vault.

Note: You cannot create mirrored protected vaults using a protected vault template. The **Mirror:<Mirror Name>** page displays.

12. Click **Create Vault 2** to create the second vault of the protected mirror. You cannot add an existing vault as the secondary vault. The **Create New Vault** page displays.
13. Select a storage pool and click **Continue** to begin creating the vault.

Note: The secondary vault must be in a different storage pool than the primary vault.

14. Follow the steps in "Creating vaults" on page 31 to create the vault.

Note: You cannot create mirrored protected vaults using a protected vault template.

15. Click **Save**.

Creating protected mirrors by using mirror templates

An alternative approach to protected mirror creation is based on the use of mirror templates. These templates allow a user to create multiple protected mirrors with the same parameters quickly and enable common mirror configurations to be leveraged across multiple users.

Before you begin

To create a protected mirror from a template, the system must have Protection enabled. For more information, see "Configuring vault protection" on page 31.

About this task

You can create a protected mirror template and then use it to create a protected mirror and both the protected vaults that are mirrored. All parameters set in the protected mirror template apply to the mirror and to the vaults that are created with the mirror. You can update the protected mirror template at any time.

Protected mirrors cannot be set as the default mirror template for a system. Versioning, Delete Restricted, Recovery Listing, SSE-C settings, and synchronous action do not apply for mirrored protected vaults.

Note: You cannot create mirrored protected vaults using a protected vault template.

Procedure

1. Navigate to the **Configure** tab.
2. Select **Configure** from the **Template Management** section. The Template Management page appears.
3. In the Mirror Template section, click **Create Mirror Template**. The Create Mirror Template page appears.
4. Select **Protected Mirror Template** and click **Next**. A new page appears where a new mirror template and two vault templates can be created.

When customizing, an optional provisioning code can be entered (for example, US East). This code determines where data is stored and which vaults are written when you are using vault provisioning. In addition to providing mirror configuration parameters, Accesser devices can be selected for deployment and access to mirrors can be controlled by specifying IP addresses.

5. Click **Save**. The Mirror Template:<Mirror Name> page appears.
6. To create a mirror from the template, click **Create Mirror from Template**.
7. Provide the necessary information and click **Save**.

What to do next

Mirror templates can be edited and deleted as well.

1. From the **Template Management** page, in the **Mirror Templates** section, select the mirror template to be edited/deleted, and a new page appears.
2. Click **Change** or **Delete Mirror Template** from the action bar at the top of the page.
3. Edit the template, providing any revisions, and click **Save**.

Repairing a protected mirror

Repairing a failed primary vault

About this task

If one or more storage pools on the primary vault fails, you must delete the mirror, which results in a data outage until a new protected mirror is deployed. The following procedure allows the system to maintain compliance in the face of a temporary or permanent failure on the primary vault.

Procedure

1. Get a snapshot of the data on the secondary vault. A vault listing is ideal.

Note: Data on the primary vault that was not synchronized to the secondary vault might not be recoverable.

2. Select the mirror of interest from the navigation tree on the left side of the screen.
3. Click **Delete Mirror**.

4. Enter the password and click **Delete**.

Note: Do not select one of this mirror's vaults to be renamed and redeployed to match the mirror if you plan to use one of the vaults as part of a new protected mirror.

5. Right click an item in the navigation tree and select **Create Mirror** from the context menu.
6. Select **Protected Mirror** and then click **Next**.
7. Click **Continue without using Template**.
8. Name the protected mirror and choose the secondary vault from the old mirror as the seed vault.

Note: Use the name of the deleted mirror for the new mirror so that the new mirror can be accessed using the same endpoint as the one used to access the deleted mirror.

Note: The secondary vault from the old mirror now becomes the primary vault for the new mirror.

Note: Concentrated Dispersal vaults cannot be used as seed vaults.

9. Click **Save**.
10. Choose the storage pool for the new secondary vault.
11. Click **Continue**.
12. Name the new secondary vault and select the configuration.

Note: It is recommended that the secondary vault has the same IDA as the primary vault.

13. Click **Save**.
14. In the Deployment section of the Mirror Configure page, click **Change**.
15. Select the access pool on which to deploy the mirror and click **Update**.

Note: Allow at least five minutes to elapse before writing new data to the mirror.

16. Get a snapshot of the data on the primary vault. A vault listing is ideal. The original data should be present.
17. Periodically get a snapshot of the data on the secondary vault.

Note: Initial snapshots do not have all the data listed. However, more and more data appears in subsequent snapshots.

Note: You might see additional operational latency when using the new protected mirror if the secondary and primary vaults from the deleted protected mirror were geographically dispersed.

18. Contact IBM Customer Support to recover the storage used by the primary vault on the old mirror.

CAUTION:

During the repair process, I/O to the mirror might experience higher than normal latency. Run a low rate of I/O operations to the mirror until the secondary vault is back online and has synchronized 10 to 25 percent of the data from the primary vault.

Repairing a failed secondary vault

About this task

Procedure

1. Re-imaging Slicestor devices

If the secondary vault has one or more degraded Slicestor devices, use the following procedure.

- a. Get a snapshot of the data on the primary vault. A vault listing is ideal.
- b. Contact IBM Customer Support to re-image the degraded Slicestor devices.
- c. Accept the devices on the Configure page of the Manager Web Interface.

Note: At this time, the secondary vault should come back online and automatically start synchronization with the primary vault.

- d. Periodically get a snapshot of the data on the secondary vault.

Note: Initial snapshots do not have all the data listed. However, more and more data appears in subsequent snapshots.

CAUTION:

During the repair process, I/O to the mirror might experience higher than normal latency. Run a low rate of I/O operations to the mirror until the secondary vault is back online and has synchronized 10 to 25 percent of the data from the primary vault.

2. Replacing sets

If the Slicestor devices on the secondary vault are not recoverable by re-imaging and additional Slicestor devices are available, the secondary vault can be recovered through the replacement of sets.

- a. Get a snapshot of the data on the primary vault. A vault listing is ideal.
- b. Select the storage pool of interest from the navigation tree on the left side of the screen.
- c. Click **Configure**.
- d. From the Slicestor device's section, click **Change Sets and Devices**.
- e. From the Replace Storage Pool section, click **Configure Set Replacement**.
- f. Select the sets to replace. When the replacement is complete, the storage pool automatically comes online.
- g. Periodically get a snapshot of the data on the secondary vault.

Note: Initial snapshots do not have all the data listed. However, more and more data appears in subsequent snapshots.

CAUTION:

During the repair process, I/O to the mirror might experience higher than normal latency. Run a low rate of I/O operations to the mirror until the secondary vault is back online and has synchronized 10 to 25 percent of the data from the primary vault.

3. Assigning a new vault as the secondary vault

You can assign a new, empty vault as the secondary vault. This process involves breaking the mirror and results in data outage until the new mirror is redeployed. This assumes that the storage pool to which the new secondary vault will be deployed exists. If a new storage pool needs to be created for the secondary vault, see the *IBM Cloud Object Storage Manager Administration Guide* and create the storage pool before you proceed.

- a. Get a snapshot of the data on the secondary vault. A vault listing is ideal.

Note: Data on the primary vault that was not synchronized to the secondary vault might not be recoverable.

- b. Select the mirror of interest from the navigation tree on the left side of the screen.
- c. Click **Delete Mirror**.
- d. Enter the password and click **Delete**.
- e. Right click an item in the navigation tree and select **Create Mirror** from the context menu.
- f. Select **Protected Mirror** and then click **Next**.
- g. Click **Continue without using Template**.
- h. Name the protected mirror and choose the secondary vault from the old mirror as the seed vault.

Note: Concentrated Dispersal vaults cannot be used as seed vaults.

- i. Click **Save**.
- j. Choose the storage pool for the new secondary vault.
- k. Click **Continue**.

- l. Name the new secondary vault and select the configuration.

Note: It is recommended that the secondary vault has the same IDA as the primary vault.

- m. Click **Save**.
- n. In the Deployment section of the Mirror Configure page, click **Change**.
- o. Select the access pool on which to deploy the mirror and click **Update**.
- p. Get a snapshot of the data on the primary vault. A vault listing is ideal. The original data should be present.
- q. Periodically get a snapshot of the data on the secondary vault.

Note: Initial snapshots do not have all the data listed. However, more and more data appears in subsequent snapshots.

CAUTION:

During the repair process, I/O to the mirror might experience higher than normal latency. Run a low rate of I/O operations to the mirror until the secondary vault is back online and has synchronized 10 to 25 percent of the data from the primary vault.

Chapter 4. Security

Security overview

The **Security** tab, accessible by the initial admin account, provides a means to create, delete, and maintain accounts within the Manager. Records are kept of actions that are taken on the Manager Web Interface.

Separate accounts for each administrator can be created to identify which system administrator performs a specific task. Users are only able to perform tasks that are associated with the roles that are assigned to their account. You can assign multiple roles to a user, except you cannot assign a read/write and read-only version of the same role to a single user. For example, you cannot assign both the Read Only system administrator role and the system administrator role to the same user.

The **Security** page is only accessible by accounts with Super User or Security Officer roles.

Roles

Users and groups that are created in the Manager can be assigned roles during the creation process or later.

Every time a user attempts to access a URL, the Manager first checks to see whether the current user has a role with the privilege to view that page. Additionally, certain elements on a page might be hidden if the current user does not have the privilege to view them.

The following roles can be assigned to a user or a group.

Table 7. User or group roles

Roles	Description
Super User	Can perform any action within the Manager except vault read/write. All tabs are accessible within the application.
System Administrator	<p>Can perform any action within the Manager except security, account management, and vault read/write. Monitor, Configure, Maintenance, and Administration tabs are available.</p> <p>This role has two Access Control List options: Read/write and read-only. By default, the role has read/write access. To limit the role to read-only access, check the box in the Read Only column when assigning the role.</p>
Security Officer	<p>Can perform security and account management actions within the Manager Web Interface. A Security Officer cannot do the following items:</p> <ul style="list-style-type: none">• Change the roles of any Super User account.• Change the roles of their own account.• Change the vault privileges of their own account. <p>This role provides access to the Security tab.</p> <p>This role has two Access Control List options: Read/write and read-only. By default, the role has read/write access. To limit the role to read-only access, check the box in the Read Only column when assigning the role.</p>
Operator	Can perform monitoring actions within the Manager. Only the Monitor tab is available.

Table 7. User or group roles (continued)

Roles	Description
Vault Provisioner	(Vault Mode Only) Can create / delete vaults by using the Provisioning API. This role alone does not grant access to the Manager Web Interface and is only visible on the UI if the Provisioning API is enabled. Note: See Configure Provisioning API.

A Vault User (vault mode only) has read/write or read-only access to Object vaults. This role alone does not grant access to the Manager Web Interface. This role is specified on a per-vault basis and is assigned by granting read or read/write access to a vault. Vault User is only applied to object vaults.

A Service Account (container mode only) must be assigned to all accounts that interact with the Service API.

Multiple roles can be assigned to a single account or group, except you cannot assign a read/write and read-only version of the same role to a single account or group. For example, you cannot assign both the Read Only system administrator role and the system administrator role to the same user. Accounts within a group inherit all roles of the group. If a group is created with a system administrator role assigned and a member of that group is assigned an Operator role, they have the larger set of privileges of the group rather than the **Monitor Only** view of the Manager.

By clicking the **Security** page, you can set up accounts and assign roles to those accounts. The Navigation tree shows accounts and roles and allows for new accounts to be created via right click.

When a user logs in to the Manager Web Interface, they see a view that is associated with the roles that are assigned.

Authentication and authorization

Four methods of authentication for the Manager are available.

- Via the Manager, by using a password
Password authentication is always available by default.
- Via the Manager, by using a personal certificate
Certificate-based authentication is available if an external CA was configured in the **Administration** tab. For more information, see “Configuring certificate authority” on page 71.
- Via an external directory service, by using an Active Directory or LDAP server
Directory service-based authentication is available if an external directory service was configured within the **Administration** tab. For more information, see “Configure active directory / LDAP” on page 68.
- Via an external Keystone server
Users and Groups can be created and assigned roles to enable access to Vaults and Vault Mirrors in the system.

Creating an account Procedure

1. Log in to the Manager Web Interface.
2. Click the **Security** tab in the main menu to display the **Security** page.
3. Click **Create Account** from the **Accounts and Groups** section to display the **Create New Account** page.
4. Enter the user's name in the **Name** field.

5. Enter the users email address in the **Email** field.
6. Select the method by which this user is authenticated.
 - a. If this user is authenticated against the Manager Web Interface itself, click the **Local** radio button.
 - 1) If this user is authenticated with a password, check on the **Allow authentication with a username and password maintained within the Manager**.
 - a) Enter the user's local user name in the **Username** field.
 - b) Enter the user's password in the **Password** field.
 - c) Enter the user's password again in the **Confirm Password** field to validate the new password.
 - 2) If this user is authenticated with a certificate, check on the **Allow authentication with PKI using a certificate from an external certificate authority**.
 - a) Enter the Certificate Subject Distinguished Name in the **Subject DN** field in (RFC2253) format.
 - b. If this user is authenticated against an external directory service, click the **Active Directory** or **LDAP** radio button.

Note: A system can use either Active Directory or LDAP as a directory service.

- If Active Directory is the enabled directory service, the second radio button shows **Active Directory**.
- If LDAP is the enabled directory service, the second radio button shows **LDAP**.
- 1) Enter the user's principal name in the **User Principal Name** field.

Note: The user principal name usually is the users email address. Check with your directory service administrator for this information.

- c. If this user is authenticated against a Keystone Server, click the **Keystone** radio button.

Note:

A system can use Keystone as an external authentication service.

- When Keystone is configured, the Keystone option is shown.
- If Keystone is configured but pointing to a Keystone v2.0 server, a Keystone domain **cannot** be set.
- 1) Enter the Keystone user name in the **Username** field.
- 2) Enter the Keystone domain in the **Domain** field.

7. Check the check boxes in the **Assign Role** column that correspond to the role you want to assign to the user. To grant read-only permissions, check the check box in the **Read Only** column.

Note: See “Roles” on page 51 for specifics on what capabilities each role possesses.

8. To grant the user permissions on standard vaults and management vaults (access to containers is managed via Service API or Cloud Object Storage requests), perform the following steps under the **Vault Access** heading:
 - a. Click the tab that has the current permission for the Vault for which you want to change permissions.
 - b. Check the check boxes to the left of each Vault for which you want to change the user's permissions.

Note:

- Click the **Select All** link to select all Vaults.
- Click one vault, hold down the Shift key, then click another vault farther down the list to select a range of Vaults.

- c. Depending upon which tab has focus, you can change permissions for the selected Vaults:
 - To grant the user owner permission, click the **Move to Owner**.
 - To grant the user read-only permission, click the **Move to Read-Only**.
 - To grant the user read and write permission, click the **Move to Read/Write**.
 - To revoke all permissions, click the **Move to No Access**.
 - d. To change which permission is granted, click the appropriate tab to find the wanted Vault, then repeat these few steps to give the user the correct permission for the Vault.
9. Click the **Save** from either of the taskbars.

What to do next

Filtering Vault Lists

The displayed Vaults can be focused by filtering the list.

Use one or more of the provided filters to reduce the Vault list:

- Select one or more Storage Pools from the **Storage Pool** list menu to display Vaults only found in the selected Storage Pools.
- Select one or more Vault types from the **Vault Type** list menu to display only Vaults of the specific types:
 - **Standard**
 - **Management**
 - **Mirror**
- Select one or more Tags from the **Tags** list menu to display Vaults with the applied Tags.
- Enter the text of the names of specific Vaults in the **Text** field to display Vaults with that text in their names.

Filters are additive: If the Management Vault Type is selected and S0 is entered into the **Text** field, only Management Vaults with S0 in the name display.

The filters take effect immediately. You do not need to accept or enable the filters specifically. To remove all filters, click the **Clear filters**.

Editing an account

Procedure

1. Log in to the Manager Web Interface.
2. Click the **Security** tab in the main menu.
3. Click the name of the user under the **Accounts and Groups** section to display the **Account: {User}** page.
4. To change details for the User, click **Change** to display the **Edit Account: {User}** page.
5. Change the details that need revision.
 - a. Enter a new name for the user in the **Name** field.
 - b. Enter a new email address for the user in the **Email** field.
 - c. Change the users authentication method:
 - 1) If this user is authenticated against the Manager Web Interface itself, click the **Local** radio button.
 - a) If this user is authenticated with a password, check on **Allow authentication with a username and password maintained within the Manager**.
 - i. Enter the users local user name in the **Username** field.

- ii. Enter the users password in the **Password** field.
 - iii. Enter the users password again in the **Confirm Password** field to validate the new password.
- b) If this user is authenticated with a certificate, check on **Allow authentication with PKI using a certificate from an external certificate authority**.
 - i. Enter the Certificate Subject Distinguished Name in the **Subject DN** field in (RFC2253) format.
- 2) If this user is authenticated against an external directory service, click the **Active Directory** or **LDAP** radio button.

Note:

A system can use either Active Directory or LDAP as a directory service.

- If Active Directory is the enabled directory service, the second radio button shows **Active Directory**.
- If LDAP is the enabled directory service, the second radio button shows **LDAP**.

- a) Enter the user's principal name in the **User Principal Name** field.

Note: The user principal name usually is the users email address. Check with your directory service administrator for this information.

- 3) If this user is authenticated against a Keystone Server, click the **Keystone** radio button.
 - a) Enter the Keystone user name in the **Username** field.
 - b) Enter the Keystone domain in the **Domain** field.

Note: If Keystone v2.0 is in use, a Keystone domain cannot be set.

- d. To change the users time zone, do one of the following tasks:
 - Click the **Use the default manager time zone (GMT)** radio button.
 - Click **Use a custom time zone for this account**.
 - Select a time zone from the **Select a time zone** list menu.
- e. Check the check boxes in the **Assign Role** column that correspond to the role you want to assign to the user. To grant read-only permissions, check the check box in the **Read Only** column.

Note: See “Roles” on page 51 for specifics on what capabilities each role possesses.

- f. To grant the user permissions on vaults, perform the following tasks under the **Vault Access** heading:
 - 1) Click the tab that has the current permission for the Vault for which you want to change permissions.
 - 2) Check the check boxes to the left of each Vault for which you want to change the users permissions.

Note:

- Click the **Select All** link to select all Vaults.
- Click one vault, hold down the **Shift** key, then click another vault farther down the list to select a range of Vaults.
- 3) Depending upon which tab has focus, you can change permissions for the selected Vaults:
 - To grant the user owner permission, click **Move to Owner**.
 - To grant the user read-only permission, click **Move to Read-Only**.
 - To grant the user read and write permission, click **Move to Read/Write**.
 - To revoke all permissions, click the **Move to No Access**.

- 4) To change which permission is granted, click the appropriate tab to find the wanted Vault, then repeat these few steps to give the user the correct permission for the Vault.
- g. Click **Update** from either of the taskbars.

Deleting an account

Before you begin

Attention: Only users with Security Officer privileges can remove a User.

Note: The admin account cannot be deleted.

Procedure

1. Log in to the Manager Web Interface.
2. Click the **Security** tab from the main menu.
3. Click the user to be deleted from the list of accounts under the **Accounts and Groups** list. The **Account: {User}** window appears.
4. Click **Delete Account**. The **Delete Account: {User}** window appears.
5. Enter your password to confirm the deletion.
6. Determine which action you want to take.
 - Click **Delete** to remove the account.
 - Click **Cancel** to leave the account untouched.

Creating a group

Before you begin

Attention: If **Create Group** does not appear, it means the Active Directory / LDAP or Keystone server are not configured properly (see “Configure active directory / LDAP” on page 68).

Procedure

1. Log in to the Manager Web Interface.
2. Click the **Security** tab in the main menu.
3. Click **Create Group** from the **Accounts and Groups** heading to display the **Create New Group** page.
4. Select the method by which this group is authenticated.
 - a. If this group is authenticated against an external directory service, click the **Active Directory** or **LDAP** radio button.
 - 1) Enter the distinguished name for the group in the **Distinguished Name** field.
 - b. If this group is authenticated against a Keystone Server, click the **Keystone** radio button.
 - 1) Enter the Keystone group name in the **Name** field.
 - 2) Enter the Keystone domain in the **Domain** field.
 - 3) Select either **Project** or **Group** from the **Type** radio button set.

Note: If Keystone v2.0 is in use, a Keystone domain cannot be set, and you cannot set a Keystone type.

5. Enter the alias for the new group in the **Alias** field.
6. Check the check boxes in the **Assign Role** column that correspond to the role you want to assign to the user. Check the check box in the **Read Only** column to grant read-only permissions.

Note: See “Roles” on page 51 for specifics on what capabilities each role possesses.

7. Perform the following steps under the **Vault Access** heading to grant the group permissions on standard vaults and management vaults (access to containers is managed via Service API or Cloud Object Storage requests).
 - a. Check the check boxes to the left of each Vault for the groups that you want to grant access.

Note: Click **Select All** to select all Vaults.

- b. To grant the group owner permission for the selected Vaults, click **Move to Owner**.
 - c. To grant the group read-only permission for the selected Vaults, click **Move to Read-Only**.
 - d. To grant the group read and write permission for the selected Vaults, click **Move to Read/Write**.
 - e. To change which permissions are granted, click the appropriate tab to find the wanted Vault, then repeat these few steps to give the group the correct permission for the Vault.
8. Click **Save** from either of the taskbars.

Editing a group

Before you begin

Attention: If **Create Group** does not appear, it means the Active Directory / LDAP server is not configured properly (see “Configure active directory / LDAP” on page 68).

Procedure

1. Log in to the Manager Web Interface.
2. Click the **Security** tab in the main menu.
3. Click the group that you want to edit, in the **Accounts and Groups** section.
4. Click **Change** and display the **Edit Account: {Group}** page to change details for the User. Regardless of how this group is authenticated, the only property that can be changed is the alias of the group in the **Alias** field.
 - For an Active Directory or LDAP Server, the **Distinguished Name** cannot be changed.
 - For a Keystone Server, **Name**, **Domain**, **Alias**, and **UUID** cannot be changed.
5. Check the check boxes in the **Assign Role** column that correspond to the role you want to assign to the user. To grant read-only permissions, check the check box in the **Read Only** column.

Note: See “Roles” on page 51 for specifics on what capabilities each role possesses.

6. Perform the following steps under the **Vault Access** heading to grant the group permissions on standard vaults and management vaults (access to containers is managed via Service API or Cloud Object Storage requests).
 - a. Check the check boxes to the left of each Vault the group should have some form of access.

Note: Click the **Select All** link to select all Vaults.

- b. To grant the group owner permission for the selected Vaults, click **Move to Owner**.
 - c. To grant the group read-only permission for the selected Vaults, click **Move to Read-Only**.
 - d. To grant the group read and write permission for the selected Vaults, click **Move to Read/Write**.
 - e. To change which permission is granted, click the appropriate tab to find the wanted Vault, then repeat these few steps to give the group the correct permission for the Vault.
7. Click **Update** from either of the taskbars.

Deleting a group

Before you begin

Attention: If **Delete Group** does not appear, it means the Active Directory / LDAP server is not configured properly (see “Configure active directory / LDAP” on page 68).

Attention: Only users with Security Officer privileges can remove a Group.

Procedure

1. Log in to the Manager Web Interface.
2. Select the **Security** tab in the main menu.
3. Select the group to be deleted from the list of groups that appears.
4. Click **Delete Group**. The **Delete Group** window appears.
5. Enter your password to confirm the deletion.
6. Determine which action you want to take:
 - Click **Delete** to remove the group.
 - Click **Cancel** to leave the group untouched.

Granting access key and password authentication

Access Key Authentication enables the generation of AWS-style credentials for user accounts. These credentials can be used to perform AWS authentication for S3 requests. Password Authentication is also supported.

About this task

This task applies to accessing standard vaults and management vaults. Access keys for container access must be managed via the Service API or Cloud Storage Object API.

CAUTION:

It is recommended that at least one option is selected (Access Key Authentication or Password Authentication).

Procedure

1. Navigate to **Security > Enable/Disable Authentication Mechanisms**
2. Enable Access Key, Password Authentication, or both by clicking **Configure**. The target account must exist or be created. Accounts that are created while Access Key Authentication is enabled no longer require a username/password to be set.
3. Take the following steps for Access Key Authentication.
 - a. Create access keys for the target account (Security | Account | Access Key Authentication).
 - b. Click **Change**.
 - c. Click **Generate New Access Key** to create new access keys (credentials). A maximum of 10 access keys can be created.

What to do next

Upon enabling **Access Key Authentication**, the Access Key Authentication section appears on the account page, and the Create Account flow allows an account to be created without requiring a user name and password.

A vault user with Access Key Authentication can log in to **My Account** and change **Name, Email, Timezone, Username, Password**, and change **Access Keys**. A vault user with Password Authentication does not have **Access Keys**.

An account can be disabled by someone with Security Officer privileges, which prevents vault access and the ability to log in to the Manager Web Interface. Individuals with these roles can also re-enable the account.

Hide Secret Key can be enabled to enhance security measures for managing the Secret Key portion of Access Keys. Enabling the feature makes the Secret Key portion of Access Keys one-time accessible at the time of creation and inaccessible thereafter. If Access Keys exist before you enable the feature, those Secret Keys are no longer accessible. It is advised to keep a record of necessary Access Keys before you enable the feature. When you create new Access Keys with hiding enabled, Secret Key appears in a modal window along with an option to download a .csv file that contains both the Access Key ID and Secret Key token. After the modal is closed, the secret portion cannot be retrieved again.

Note: AWS authentication for S3 requests **should not** be used along with groups. Access Key authentication requests ignore group permissions.

Chapter 5. Administration

Overview

The **Administration** tab provides the following functions.

- Protect Manager data.
- Configure alerts.
- Specify external authentication by Active Directory / LDAP.
- Enable TLS and set Manager preferences.
- Other tasks.

Backing up and restoring the Manager

A utility for backing up and restoring the Manager database can be found under the **Administration** tab. This database contains information about the storage configuration (devices, vaults, configurations, security certificates, and so on) and event history. It also contains users, passwords, and configuration parameters (email, alerts, and so on). Performance and device statistics are not currently part of the database backup.

About this task

The Manager data must be backed up to recover the system configuration. The Manager database backup file is encrypted for extra security since this database contains sensitive configuration information. This password is needed to restore the database. While this password can be retrieved from an existing Manager, a replacement device cannot retrieve it. It is important to keep this password in a safe place.

Attention: When a Manager device needs to be replaced, and the data was not backed up, the vault information and the keys that are used in the communication between the Manager device and Slicestor / Accesser devices / network will be lost. A manual backup is required those situations.

From the **Administration** page.

Procedure

Click **Backup** or **Restore** for the activity wanted. After selecting **Backup**, the Manager database can be backed up either automatically or manually.

Configure Manager backups

The backup location and schedule must be configured for automatic backups. Automatic backup requires the use of the local directory (/var/lib/dsnet-manager/backups), remote SFTP server, or a remote FTP server. Click **Configure** to enter these settings.

Note: You can also perform a manual backup.

Setting the retention policy

A retention policy is in place for backups to a local directory.

About this task

By default, the five most recent backups are saved in the configured location. The number of backup files can be changed in the Manager configuration file.

Procedure

1. Set the **Local Directory Path** in the `managerConfiguration.properties` file.
`backup.localDirectoryPath={userDefinedPath}`
2. Set the **Maximum Value for Backups** in the `managerConfiguration.properties` file.
`backup.maxCount={value}`
3. Set the permissions for the newly created folder to 755.
`chmod 755 {foldername}`
4. Set the owner to `dsnet-manager`.
`chown dsnet-manager:dsnet-manger {foldername}`
5. Restart the Manager Web Interface to activate these values.
`service dsnet-manager restart`

Set backup configuration parameters

When using the automatic backup option, the Manager sends the backup file to the local directory, remote FTP server, or remote SFTP server.

Table 8. Backup configuration parameters

Parameter	Description
Encryption Password	The Manager database backup file is encrypted for extra security. This database contains sensitive configuration information that might be used to obtain unauthorized access to the system. Enter up to 255 characters and reconfirm as indicated. This password is used to restore the database. A blank field is not allowed.
Enable Secure Backup Password Recovery	Click this option to allow IBM to recover lost passwords.

Attention: The encryption password must be remembered or kept in a safe place. Always store passwords in a secure place away from the information that they help protect.

Back up to a local directory

By default, the local directory path for the backup files is `/var/lib/dsnet-manager/backups`. The path can be changed by creating the Manager Web Interface configuration file, `managerConfiguration.properties`, in `/var/lib/dsnet-manager`.

Backing up to a secure file transfer protocol

About this task

You can authenticate to a remote SFTP server that uses SSH keys for a manager backup file.

Procedure

1. From the **Protocol** menu, select **SFTP**. FTP is an available option but it is not a secure form of file transfer.
2. Type the SFTP server host name or IP address in the **Hostname** field.
3. Type the SFTP user name in the **Username** field.
4. Type the SFTP password for that user name in the **Password** field.
5. Type the SFTP directory path in the **Path** field.
6. Enter the backup frequency and time, either Daily or Weekly automated backup is available.

Note: The automatic backup time shows the (local time zone).

7. Click **Update**. The Backup Configuration page displays an SSH Public Key.

8. Copy the key and paste it into the `~/.ssh/authorized_keys` file of the target user on the SFTP server file system.
9. Click **Configure** to enable strict validation of remote SSH keys.
10. In the Backup Destination section, check the **Check to enable the validation of SSH Known Hosts** check box. Upon first successful connect and authentication to the SFTP server, the server's SSH host key is captured and stored on the Manager. The SSH host key is verified against the key that is stored on the Manager during future attempts to connect to the same host.

Regenerating an SSH key Procedure

1. Click the **Administration** tab.
2. Click **Backup** in the Manager Backup & Restore section.
3. Click **Configure**.
4. In the Backup Destination section, temporarily change the **Protocol** from **SFTP** to **FTP**.
5. Click **Update**. The Backup Configuration page displays with the SSH key pair and SSH known hosts values cleared.
6. Click **Configure**.
7. In the Backup Destination section, change the **Protocol** from **FTP** to **SFTP**.
8. Click **Update**. The Manager generates a new SSH key pair.

Backing up the Manager manually Before you begin

CAUTION:

The FTP server must be configured in advance.

Procedure

1. Go to the **Administration** page.
2. Click **Backup** to show the **Backup Configuration** page.
3. Click **Backup Manually** on the **Backup Manager Manually** action bar.
4. Select the location to which you want to back up the Manager database:
 - Desktop (default)
 - The location that is defined as the automatic backup location.
5. Click **Backup** to begin the backup.

Restoring Manager data

The manager database can be fully restored from backup in a manager failure event.

Before you begin

Attention: After the database is restored, the previous version cannot be recovered or restored except from a backup copy. Deleted vaults can never be restored. Any device or vault change (including new vaults) made to the system since the last backup is lost.

About this task

The database backup contains information about the storage configuration (devices, vaults, configurations, security certificates), and event history. It also contains the Manager users, passwords, and configuration parameters (email, alerts). Performance and device statistics are also part of the database backup.

The Manager database can be fully restored from backup in a Manager device failure. If a new Manager device is detected, the initial configuration screen is shown. You can restore the database at this screen. Alternatively, it is possible to choose **New** and **Restore** from the **Administration** page.

If the Manager was configured to use an external CA, a new certificate needs to be issued if the Manager is replaced or re-imaged.

If an external agent was configured, the agent might need to be reinstalled after the restore. For more information, contact Customer Support.

It is possible to restore the Manager database from at most two major releases back.

Procedure

1. Enter the **Restore** page from the **Administration** page by clicking **Restore**.
2. Enter your account password.
3. Select the backup file by clicking **Choose File**.

Note: If the backup was to an FTP server, the file must be transferred from the remote server to a local disk using binary mode.

4. Set the FTP transfer mode to binary.
5. Select the backup encryption password that was configured during the initial backup.
6. Click **Restore** from the action bar.

Monitor the Manager backup progress

The backup progress is reported on the lower left corner.

Table 9. Backup progress states

State	Description
Generating	The backup file is being generated.
Transferring	The backup file is being saved to the backup destination.
Backup completed successfully.	The backup file was successfully saved to the backup destination.
Failed	The backup failed to generate correctly or file might not be saved to the backup destination.

Retrieving a Manager backup from a remote server

If backups are stored locally on a Manager device, it is possible to retrieve backup files from that device through methods like secure file transfer (sftp) or secure copy (scp).

Procedure

1. Retrieve files from the Manager device by using secure file transfer with the localadmin account.
`sftp localadmin@[ipaddress]:/var/lib/dsnet-manager/backups/* .`
2. Provide SFTP with the password when prompted.
`localadmin@[ipaddress]'s password:`

Note: SSH Key can be used instead of a password to authenticate with a Manager device.

Configuring SMTP

Before you begin

The administrator sets the SMTP Host to the appropriate email server address. Email notifications can then be sent from the Manager device when a potential problem occurs.

Procedure

1. Using the Manager, from the main menu, click the **Administration** tab.
2. Click **Configure** on the **SMTP Configuration** action bar.
3. Enter the information for your email relay/system as indicated in the **SMTP Configuration** screen.

From

The sender of the email. Some email systems require the sender to be a user on the system to relay.

Host

The outgoing SMTP server. Check with email administrator to be sure that the Manager is configured as an allowable SMTP relay.

Use SMTP over SSL.

Check whether this is needed for SMTP communications.

Use a custom Port.

The default SMTP port is 25. This port needs to be changed if SMTP is network-address translated or SSL is configured (port 465).

SMTP Server Requires Authentication

Some email systems require the user to log in to relay mail. Enter the user ID and password that was created for the Manager.

4. Specify one email address (it can be an email list) for a daily email that is sent at midnight in the configured Manager time zone. The time zone is specified in the Preferences section on the **Administration** tab. This feature allows an administrator to monitor the email server path. The feature is enabled after you provide an email address and click **Update**.
5. Click **Update**.
6. Click **Test** on the next page to test SMTP communications. Adjustments can be made by entering the **SMTP Configuration** through the **Administration** tab.

Configuring Call Home

When Call Home is enabled and a new incident opens in the Manager for an IBM hardware device, the Manager automatically sends an email to open a problem report ticket with IBM Customer Support.

Before you begin

You must configure SMTP to send outbound Call Home notifications. For more information, see "Configuring SMTP."

Procedure

1. Click the **Administration** tab.
2. In the **Call Home Configuration** section, click **Configure**. The Call Home Configuration page displays.
3. Select **Enable Call Home to notify IBM Customer Support of open incidents**. The configuration fields are available for input.

4. Enter the **IBM Customer Number**. The IBM Customer Number is a 6- or 7-digit number that is given to the customer for use with IBM Customer Support. Tickets that are generated by Call Home are opened under this customer number.
 5. Select the **Support Area** where the customer's account is managed:
 - **North or South America**. This area includes Central America and the Caribbean Islands.
 - **Africa, Asia, Australia, or Europe**. This area includes the Middle East and the Pacific Islands.
 - **Other - Email address required**. Notifications are sent only to the email addresses entered in the **Email Addresses to Copy** field and are not sent to IBM Customer Support. At least one email address is required.
- For systems that span across hemispheres, contact IBM Customer Support to determine which area to select.
6. Enter any **Email Addresses to Copy**. These email addresses receive the same notifications that are sent to IBM Customer Support. Enter email addresses in a comma-separated list or on separate lines.
 7. Click **Update**.
-

Configure alerts

Email alerts and alert forwarding can be set up. Use the action bar options to either **Configure Email Alert Rules** or **Configure Alert Forwarding** for third-party applications.

Configuring email alert rules

The Manager Web Interface can configure sending an email based on various conditions, like when an event with a particular severity is encountered.

Procedure

1. Click the **Administration** tab.
2. Click **Configure Email Alert Rules** on the **Alert Configuration** action bar. The **Email Alert Rule Configuration** page opens.

The page displays a list of rules that are available, including **Edit**, **Delete**, and **Disable** interfaces.
3. Click + (plus) to expand the rule and show details. Click - (minus) to collapse the rule.

Note:

Alerts for Vaults can be set only at the global level, not for individual vaults.

Alerts for Devices can be set only at the global device type level, not for individual devices.

The icon next to the rule name indicates the criticality level that is associated with it (information, warning, error, or critical).

4. Click **Create** in the **Email Alert Rule Configuration** action bar.
5. In the **Create Email Alert Rule** section, enter the email addresses for the recipients of the email (comma-separated).
6. Filter by level, event sources, tags, and event categories.

Note: For Chrome on a MacOS, the scroll bars for filter event sources and event categories might not show up. In **System Preferences** on the Mac, go to **General** and select **Always** for **Show scroll bars**.

Note: To avoid excessive email alerts, set the Level to critical. Additionally, filter event sources and categories as needed.

Note: If one or more event categories are selected, then email messages about events that are not categorized are not sent.

The number of alerts are displayed on the **Administration** page.

Configuring alert forwarding

Procedure

1. Click **Configure Alert Forwarding** from the **Alert Configuration** action bar. The **Alert Forwarding Configuration** page opens.
2. Check **Enable Alert Forwarding** and select the device types for which alert forwarding is wanted.
 - For Traps, check **Enable Trap Forwarding** and configure the Community String. Enter the IP address or host name for the external network management server (for example, 192.168.5.90:162). The default port is 162.
If HTTP and access logs are wanted, check this box.
 - For Syslog, check **Enable Syslog Forwarding**.
 - Enter the IP address or host name for the syslog server (for example, 192.168.5.90:514). The default port is 514.
 - Select a Facility from the menu.
 - If HTTP and access logs are wanted, check this box.
3. Click **Update**.

SNMP trap details

Device Uptime

sysUpTimeInstance = <TimeTicks value>

OID of the trap and depends on the type of event.

Trap structure for unstructured events

sysTrapOID = 1.3.6.1.4.1.28129.1.3.5.1.1 <csDefaultTrap>

Event message for unstructured events

1.3.6.1.4.1.28129.1.3.5.2.1 <csTrapMessage> = "Event Message"

Trap structure for structured events

sysTrapOID = <csTrapEventType>

A unique OID that identifies an event type as the second variable binding.

One variable binding exists for every event parameter and applies only to structured events.

The access log entries (also called notices) follow the same general pattern:

sysUpTime = TimeTicks

sysTrapOID = <csTrapDefaultNotice>

(If not recognized) or

<csTrapOperationComplete>,<csTrapRequestNotPerformed>,
<csTrapRequestIncomplete>,<csTrapSecurityError>,<csTrapIoFailure>,
<csTrapServerError>

These parameters are used to build the NCSA-compliant message.

- csTrapRemoteAddress
- csTrapForwardedFor
- csTrapRemoteUser
- csTrapTimeStart

- csTrapRequestMethod
- csTrapRequestUri
- csTrapProtocol
- csTrapStatus
- csTrapResponseLength
- csTrapReferer
- csTrapUserAgent
- csTrapRequestLatency

Suppressing device events

This capability suppresses events that are generated by a device and can be useful when, for example, a bad/misbehaving device is generating numerous events and email alerts.

About this task

CAUTION:

Enabling this feature does not suppress any events and must be set on each device.

Note: See “Disabling events on a device” on page 107 for more on suppressing events on a specific device.

Examples of alerts that are suppressed by this feature include system command checks, SNMP checks, process checks, and CPU temperature checks. Email alerts are also deactivated.

The ability to suppress events on devices can be enabled/disabled:

Procedure

1. When on the **Administration** page, click **Configure** on the **Event Suppression Configuration** action bar. The **Event Suppression Configuration** page opens.
2. Find the check box labeled **Enable Event Suppression**. The check box can either be checked or cleared depending on if event suppression is wanted.
3. Click **Update**.

Configure active directory / LDAP

The system supports authentication against an Active Directory or LDAP server. Detailed steps for configuring Active Directory or an LDAP server are provided.

Note: In general, although particularly for an LDAP server, an administrator must carefully review the suggested configuration information that is auto-populated by the Manager and adjust or validate. See examples that are displayed to the right of the text boxes in the configuration section.

Begin the configuration:

1. Click the **Administration** tab.
2. In the **Active Directory / LDAP Configuration** section, click the **Configure**.
 - a. For Active Directory, enter the domain in the Discovery box; the Manager attempts to automatically discover details about the server. The Manager discovers details in these ways:
 - DNS SRV records. Entering xyz.com as the domain triggers a search of DNS SRV records for ldap.tcp.xyz.com against each of the configured DNS servers. If successful, the resolution provides the name of the AD server and any backup servers.
 - If the first method fails, then resolve the domain name through DNS to the IP of xyz.com.

- b. If you run into problems with automatic discovery, confirm the following items:
 - DNS servers are configured on the Manager; at least 2 DNS servers are specified.
 - DNS port (UDP 53) is open on the firewall between the Manager and the DNS server.
 - LDAPS/LDAP ports (TCP 636 and 389) are open between the Manager and the AD server.

Note: The Manager also supports assigning roles to Active Directory groups within the **Security** tab. After the AD server is specified, Create Group will now appear on the **Security** tab.

- c. If the discovery succeeds, the fields that follow the Discovery box populate (such as enabling authentication, Domain, LDAP URL, User Authentication Model, Certificate PEM).

Note: The discovery populates information for only the top-level domain in a nested Active Directory structure. If you use a domain structure with parent-child relationships, add any child domains to the **Domain** field separated by a comma or space. For example, where mydomain.com is the parent domain and maple.mydomain.com and oak.mydomain.com are children, enter the following text in the **Domain** field: mydomain.com,maple.mydomain.com,oak.mydomain.com.

- d. If your DNS infrastructure supports SRV lookups and you want to dynamically determine the LDAP URLs at authentication time, check the **Dynamically identify LDAP controller names from DNS SRV records** box. By default, the box is not checked the software uses the static list of configured LDAP URLs.
 - When an authentication attempt is made that uses an LDAP user name and password against either a Manager or Accesser, the device dynamically discovers the LDAP service host names from DNS then. For example, a DNS SRV lookup is done for _ldap._tcp.<domainName> and a controller is selected by using standard SRV priority and weighting rules.
 - For a GDG environment, you can ensure that the device always uses a local LDAP service by configuring the LDAP Service Name on each site in the Manager. For example, a DNS SRV lookup is done for _ldap._tcp.<siteLdapServiceName>._sites.<domainName> and a controller is selected by using standard SRV priority and weighting rules.
- e. Depending on whether an Active Directory or an LDAP authentication model is being used:

Model	Extra Configuration
Active Directory	If everything appears okay, no further configuration is needed. Click Update .
LDAP server	<ul style="list-style-type: none"> • Select LDAP Bind+Search Model in the User Authentication Model section. • Review auto-populated fields and adjust as needed. • Click Update. • If Group Support is enabled for LDAP, a new action appears for accounts under the Security tab.

Note: For account or group creation or deletion, see “Authentication and authorization” on page 52 .

Configuring keystone authentication

The Manager supports authentication against a Keystone server. Detailed steps for configuring a Keystone server with the system are provided.

Procedure

1. Click the **Administration** tab.
2. In the **Keystone Auth Configuration** section, click **Configure** to display the **Keystone Auth Configuration** page.

3. Check the **Enable Keystone HTTP Authentication** check box to enable the Keystone authentication and the corresponding options on the **Keystone Auth Configuration** page.
4. Type the host name of the Keystone authentication server in the **Hostname** field. 192.168.14.63
my.keystoneserver.com

Note: Do not include http:// or https://.

5. Check the **Use SSL/TLS** check box to use HTTPS.
6. Check the **Use a custom port** check box.
 - a. If the **Use a custom port** check box is checked, the custom port field activates. Enter the **admin** port to be used for Keystone authentication.
 - The default port is 35357.
7. Choose the **Keystone Protocol Version** to be used for authentication.
 - a. Click the **v2.0** for Keystone v2.0.
 - b. Click the **v3** for Keystone v3.
 - Because the concept of a Domain was introduced in Keystone v3, the domain options are not required for v3.
8. Enter the character to be used to separate user name from domain in the **Domain Separator** field.

Note: The domain separator must be a character that is not used in any Keystone user name (for example, @). Keystone credentials for basic authentication against an Accesser must have a user name that is provided along with the domain, unless you are using Keystone v2.0 or when a default domain is set.

9. Check **Use a default domain** to allow the use of a default name.

Note: Keystone users that authenticate against Accesser devices without explicitly providing a domain are scoped by the default domain.

10. Enter the shared secret with the Keystone configuration in the **Admin Token** field.

Note: In many Keystone server configurations, the admin token is located in the /etc/keystone/ directory.

11. Paste a PEM file into the **Certificate PEM** field.

Note: A certificate is needed when the Keystone server is using SSL with self signed certificates. The certificate for many Keystone server configurations is at /etc/keystone/ssl/certs/ca.pem.

12. Click **Update**.

What to do next

Note: For account or group creation or deletion, see “Authentication and authorization” on page 52.

Configure provisioning API

The Provisioning API Configuration allows an administrator to control the type of vault provisioning requests available to users through the storage APIs (SOH, S3, and so on).

Provisioning user actions are disabled by default but can be set to Create Only (Delete by system administrator via Manager Web Interface or Manager REST API) or Create and Delete.

Use of the Provisioning API to create a vault requires the following.

- Provisioning API is enabled (**Administration > Provisioning API**) - Create Only or Create and Delete.
- Target user account is given the Vault Provisioner role (**Security > Account > Edit Account**) or (**Security > Create Account**). See “Roles” on page 51.

- The system must be in vault mode. In container mode, the Provisioning API is enabled by default and is used for creating/deleting containers instead of vaults.

Note: Vault creation provisioning requests create a vault according to the specified vault template, which is identified in the request by the vault template provisioning code. If no provisioning code is supplied, the default template is used. See “Vaults” on page 27. Compliance templates cannot be set as the default template.

The Owner option for vault authorization (“Vaults” on page 27) is inherited when a vault is created through the provisioning API. It can also be assigned by the Security Officer. This option must be set for deletion by the user.

If the Provisioning API is disabled and never enabled, the Owner option still exists for the vault-account permission association. If the Provisioning API is disabled, no difference between the Owner option and the read/write option.

Configuring certificate authority

An external certificate authority (CA) can be used to authenticate devices and users. You can specify CAs for all system devices or just the Manager device.

Procedure

1. Click **Add CA** to enter either a Device or User CA. For more information, see *** and ***.
2. Copy and paste the certificate information into the text box.
3. Click **Save**.

Note: The Online Certificate Status Protocol (OCSP) responder configured in the certificates and the Certificate Revocation List (CRL) distribution point for the CRLs issued by the CA must be reachable by all appliances over the necessary protocols (HTTP/HTTPS) so that the OSCP can be queried and CRLs can be downloaded.

Adding a certificate authority

An external user certificate authority (CA) can be entered for PKI authentication.

About this task

Note: The PKI Authentication option requires a **License Key** to be enabled. Contact IBM Sales or IBM Customer Support for additional information.

Note: Every CA is associated with a particular domain of authority that is called a realm. A CA belonging to one realm is not trusted to issue certificates for devices or users in other realms to minimize damage when a CA is compromised. The security of users or devices that belong to other realms is not impacted.

Within the system, all devices must belong to a realm that uses the reserved name network. When a CA is selected to be used for issuing device certificates, the CA is automatically associated with the network realm. When a CA is used only for issuing user certificates, the name of the realm is arbitrary and up to the Administrator to select.

Procedure

1. Select the **Device** or **User Certificates** to be entered. For a User CA, also enter the **Realm**, which must be provided by group administrator. A Realm allows multiple independent logical PKIs and provides for an extra layer of security.
2. Paste the PEM-encoded X.509 certificate into the prompt field and click **Save / Finish**.

Two check boxes are seen here. One allows the CA to issue certificates for all devices (SliceStor, Accesser, and Manager devices). The second check box allows the CA to issue certificates for users. Use format that is shown here.

```
-----BEGIN CERTIFICATE-----
MIICgjCCAeugAwIBAgIBBDANBgkqhkiG9w0BAQQFADBTMQswCQYDVQQGEwJVUzEc
MBoGA1UEChMTRXF1aWZheCBTZWN1cmUgSW5jLjEmMCQGA1UEAxMdrXFlaWZheCBT
. . . .
zfmpTMANmvPMZWnmJXbMWbfWMMdzZmsGd20hdXgPfxiIKES1h18eL51SE/9dR+
WB5Hh1Q+WKG1tfgq73HnvMP2sU1G4tega+VWeponmHxGYhTnyfxuAxJ5gDgdSIKN
/Bf+KpYrtWKmpj29f5JZzVoqgrI3eQ==
-----END CERTIFICATE-----
```

Multiple certificates can be entered by using the Begin and End delimiters for each certificate separately.

Edit a certificate authority

After a CA is entered, its realm (ability to issue certificates for devices or users) can be changed.

Note: Editing a CA does not enable the certificate to be changed; certificates are immutable objects. If the intention is to renew a CA that is about to expire, you can add an additional CA to the same realm. After it is added, that CA can be used to issue new certificates for users or devices in that realm.

Errors are shown after you click **Update**.

Note: It is not possible to clear the option Trust as a device certificate authority when existing devices in the system are using certificates that are issued by this CA. Likewise, the Trust as a user certificate authority option cannot be cleared nor the realm name be changed if the CA is the only remaining issuer for its particular realm. Instead, those users or devices must be configured to a different realm or to use a different certificate before such a change is accepted by the system.

Delete a certificate authority

Click **Delete** on the **Certificate Authority Configuration** page to delete a CA.

CAUTION:

The same restrictions that apply to editing a CA apply to deleting one. If a certificate authority is associated with any user or device in the system, it cannot be deleted or untrusted. Before you delete or untrust a user certificate authority, confirm that no accounts are associated with that CA on the Security tab. Similarly, to delete or untrust a device certificate authority, confirm that no devices are associated with that CA.

Revoke a certificate authority

The Manager supports both Online Certificate Status Protocol (OCSP) and dynamic Certificate Revocation Lists (CRL) as mechanisms to determine whether a signed and otherwise valid client certificate is revoked.

OCSP is an online protocol for doing revocation checking of a single certificate at a time. Conceptually, upon receiving the client certificate as part of the SSL handshake, the Manager determines the OCSP responder URL from the certificate and then query the target OCSP Responder to ask "is this certificate revoked?", to which the OCSP responder responds with either a "yes" or "no".

CRLs are typically published once per day or once per week by the certificate authority. When the Manager receives a client certificate as part of an SSL handshake, it determines the CRL distribution point's URL from the certificate and downloads the CRL from the URL. The CRL contains a list of all the revoked certificates and the manager checks to see whether the current client certificate is in this list. If the certificate is found to be within that CRL, then validation of the certificate fails. It can take up to 1 hour before revoked certificate information is available in the Manager.

Note: This check occurs during the initial handshake of the TLS connection setup or authentication. If the certificate is revoked later during the life of the connection, this revocation is not detected. However, if you configured the Accesser application library to use TLS, and devices are configured to use TLS for outgoing connections via the configuration option within the Manager, then for the life of the TLS connection, the certificate revocation list is retrieved on a periodic basis and the certificate of the remote party is revalidated. If the other party's certificate is found to be revoked, then the TLS connection is disconnected.

The main advantage of OCSP over CRL is in the efficiency of being to get the status of a single certificate at a time. For PKI systems that contain many revoked certificates, the size of a CRL can grow large resulting in inefficiencies when it comes to transferring over the network, holding in memory, or iterating through to determine whether a certificate is in the list of revoked certificates.

Both OCSP and CRL are supported for users who are logged in to the Manager UI / API. The Accesser HTTPS interface supports client certificate authentication as well, but is not a standard authentication mechanism in either S3 or OpenStack. The Accesser HTTPS interface supports revocation checks that use CRLs but does not support OCSP.

Note: OCSP Stapling is not supported in the current implementation.

Supply external certificates to IBM Cloud Object Storage System™ devices

Once custom CAs are configured for use, each device can either be given a certificate by one of these CAs or continue to use the internal CA.

Use these three steps to configuring a device to use an externally generated certificate:

1. The device must be approved through the Manager Web Interface to accept it into the system.
2. The CSR from the device must be extracted from the Manager Web Interface and sent to a CA.
3. The resulting PEM-encoded certificate for that device must be entered into the Manager Web Interface.

Accepting a device into the system can be accomplished from the home page. Devices pending approvals into the system appear in the **Event Console**.

Once they are approved into the system, the CSR can be extracted by either of the following methods:

- Click **Configure** > on the device's page.
- Use the **Bulk Device Certificate Configuration** feature, which becomes available on the **Configure** tab after a CA is configured for issuing device certificates.

Configuring a single device certificate

To access the CSR by the first method, follow these steps.

Procedure

1. Click a device from the left navigation bar.
2. Click **Configure**.
3. Click **Change** for the **Signed Device Certificate** section at the bottom of the page.

What to do next

After **Change** is clicked, the CSR is displayed in a manner that enables it to be copied out of the Manager Web Interface and submitted to a CA. The certificate that is issued by the CA must contain the Common Name (CN) field of the Subject Distinguished Name and the content of the Subject Alternative Name as indicated in the CSR.

These parameters are validated by the Manager Web Interface when the certificate is pasted back into the interface. Once updated, the certificate is published to the device and uses it. Another restriction is that if the external CA sets the key usage field of the device certificate, then both SSL_CLIENT and SSL_SERVER key usage fields must be set.

Note: When pasting a certificate chain rather than a single certificate, all certificates up to, but not necessarily including, a configured CA must be included. If unspecified intermediate CAs are involved in issuing the certificate, and they are not included, the input fails.

Note: If you want to change from an external certificate to a signed internal device certificate, delete all certificate contents within the text box on the **Edit Device Certificate** page and click **Update**.

Configure multiple device certificates

The second option for exporting CSRs and importing certificates is useful for handling many devices. To access this option, click **Configure** next to **Bulk Device Certificate Configuration** on the **Configure** tab.

The page provides various options for selecting devices of interest, such as devices that do not yet have a certificate, or devices with certificates that are about to expire. Further refinement of selected devices is possibly by checking or clearing each device by using the adjacent check box. Below the check boxes are the CSRs for each of the selected devices. Hovering over a CSR causes a tooltip to display the host name with which the device the CSR is associated. By highlighting all of the CSRs available, they can be copied at once and submitted to a CA.

At the bottom of the page is an input field for entering device certificates.

Note:

The certificates can be entered in any order; certificates are automatically correlated with the appropriate CSR by comparing the public key in the certificate to the public key in the CSR.

Entering certificates through this form enables all issued certificates to be imported into the Manager in a single step.

It can take up to 1 hour before revoked certificate information is available in the Manager.

Enable TLS

Enable Transport Layer Security (TLS) for data communications.

When TLS is enabled, communication between Accesser and Slicestor devices use this mechanism. Click **Configure** from the TLS Configuration action bar on the **Administration** page. The following options are presented.

Table 10. TLS options

Option	Description
None	Disables TLS
Integrity	Enables authentication and integrity features of TLS.
Encryption	Enables authentication, integrity, and confidentiality features.

Select **None**, **Integrity**, or **Encryption** and then click **Update** to disable or use cryptographic protocols for data communications between the Accesser devices and the Slicestor devices. It is a system-wide feature, applicable to all vaults on this system.

If confidentiality is not wanted, selecting **Integrity** has a positive impact on performance.

TLS typically provides three features: Confidentiality (others cannot read the content), Authentication (you are communicating with whom you think you are), and Integrity (others cannot tamper with communications without being noticed). The **Integrity** option provides only authentication and integrity, providing improved performance in cases where confidentiality is not needed. The **Encryption** option provides all three features. If TLS is enabled, Cloud Object Storage Accesser devices establish encrypted connections before they store and retrieving data from Cloud Object Storage Slicestor devices [1]. Cloud Object Storage Slicestor devices establish encrypted connections to other Cloud Object Storage Slicestor devices while they perform data consistency scans and rebuilding.

The SecureSlice feature, although technically not encryption, guarantees that without access to a threshold number of slices, no information can be obtained without brute forcing the random symmetric key that is used to perform the transformation. SecureSlice should remain enabled [default] for most deployments as it provides data at rest protection and has minimal performance impact.

Note: TLS encryption might affect data throughput performance. It should be disabled unless encryption and added security are needed.

These items are independent of the TLS setting.

- Communication with the Cloud Object Storage Manager is always encrypted by using HTTPS, SSL tunneling, and SNMPv3 data protection.
- User passwords are always encrypted before transmission between any IBM devices. A user password is never sent in plaintext. [2]
- The Simple Object over HTTP (SOH) interface supports access over HTTP (plaintext) or HTTPS (encrypted). The HTTPS option is always made available for clients that are configured to use the proper port.
- Software clients that are developed with the DSAF Server SecureSlice can always establish encrypted connections. Such client applications might not acknowledge the TLS setting and should be designed to allow the user to configure separately whether TLS is used.

[1]Communication between clients and Cloud Object Storage Accessers devices is governed separately. With SOH, the client determines whether encryption is used. The Cloud Object Storage Accesser device does not support IPSec.

[2]User passwords can be sent over plaintext between a client and an Accesser device over SOH when HTTP (port 80, plaintext) is used. Clients can be configured to avoid it by using HTTPS (port 443, encrypted).

Configure drive health

Configure disk health thresholds for Slicestor models. The health icon for each Slicestor Node changes as thresholds are met. Meeting the warning threshold changes the device health icon to yellow. If the number of drives meets or exceeds the error threshold, the icon changes to red. Messages are also sent to the **Event Console**. The Slicestor Node continues to function at reduced capacity until the disks are replaced.

In the Administration tab, click **Configure** on the **Drive Health Configuration** action bar. Then, click **Enable** to configure drive health thresholds for individual Slicestor Nodes.

For a multi-node Slicestor Node configuration, the disk health thresholds can be set for all nodes or customized on a per node basis. When you have two separate multi-node devices with 2 and 3 node configurations, setting the thresholds on the two node device and three node device sets those thresholds for all nodes in that configuration. In addition, it is possible to set the disk health thresholds for each node separately, similar to any other Slicestor Node.

Configure organization

In the Administration tab, click **Configure** on the **Organization Configuration** action bar to configure owner information. Provide general and contact information. Only the Name is needed. This content is used by the System Usage and Configuration Summary Report.

Configure SNMP

SNMP v2c can be enabled for each of the device types in the system. Enabling allows polling of any selected devices from all of the provided IP addresses. This feature allows numerous statistics to be collected by a third-party application. The following check boxes and fields are available:

Table 11. SNMP fields

Field	Description
Enable SNMP	Checking this box enables SNMP collection.
Manager, Accesser Devices, Slicestor Devices, File Accesser Devices	Select the device types for polling.
Community String	Configure the community string (for example, public).
IP Addresses/Subnets	Identify the IP addresses/subnets used for polling. If the IP addresses field is left blank, IP access is open, and any IP address can poll the selected devices. IP addresses can be separated by a space or comma.

Changes that are made take a few minutes to take effect on each selected device.

Configure Device Level API

Enable **Device Level API** to configure statistic and health APIs on devices for third-party applications. Select the device type: Manager, Accesser Devices, Slicestor Devices, File Accesser Devices, and SMC Devices. Two new device-level APIs are available to access health and statistics content.

Note: For more information, see the *Device API Developer Guide*.

Configure preferences

These Manager Web Interface global settings can be changed and are seen by all users.

Table 12. Global settings

Parameter	Default	Description
Time Zone	Greenwich Mean Time (GMT)	The drop-down allows a different time zone to be selected. Time zones for individual accounts can be set on the Security tab. Note: The backup schedule uses this time zone.
Device Display	Alias	This setting applies only to the UI. All files that are exported have both host names and alias fields.
Display Name		Text value is displayed in the upper right corner of the Manager Web Interface.
Storage Units	SI units	In SI units, 1,000 Bytes equal 1 KB. Can be changed to IEC units where 1,024 Bytes equal one kibibyte.
Session Timeout	30 minutes	The drop-down allows range 5 minutes from 8 hours.

Table 12. Global settings (continued)

Parameter	Default	Description
Automatic Logout	Disabled	<p>When Automatic Logout is enabled, a pop-up window appears when the session is about a minute from expiring and gives the user the option to continue the session or to log out. Pages in the Monitor tab do not reload every 5 minutes.</p> <p>When Automatic Logout is disabled, pages in the Monitor tab will reload every five minutes.</p> <p>Note: When Automatic Logout is enabled, avoid opening the Manager Web Interface in more than one browser tab. Because the focused tab uses more processing power than unfocused tabs, the countdown to session timeout can differ by a few seconds between tabs.</p>

Configure custom login banner

Through the **Administration** tab, a user can configure a login banner.

Custom login text can be entered, with a maximum of 64,000 characters and no constraints on the type of characters. The login banner displays on devices before SSH login.

Note:

Some SSH clients might not display the login banner under certain configurations, even if the server is presenting one.

If you are using OpenSSH, setting LogLevel to QUIET, FATAL, or ERROR suppresses the banner.

You must set LogLevel to an option at least as verbose as INFO (which is the default, if unspecified) for the banner to be displayed.

For other SSH clients, refer to their documentation if they are not displaying the banner.

When configured, the user-provided banner appears on the login page.

Configuring system properties

Advanced System Properties may be adjusted for system behavior.

About this task

Attention: These options allow modification to the internal operation of the system software and must be used with caution and under the guidance of technical support.

Procedure

1. Click **Configure** on the **System Properties Configuration** title bar to open the **Edit System Properties** page.
2. Edit the values that you want to change.
 - a. Change the **Automatic Shutdown Timeout** value to modify the system process to forcefully terminate after the selected timeout period.

By default, the system process waits indefinitely for transactions to close on Slicestor devices and Accesser devices when requested to stop. Select one of the following values from the **Change Timeout** drop-down menu.

Attention: When you change this value, data being written may be lost.

- 15 minutes
- 30 minutes
- 1 hour
- 2 hours
- 4 hours
- 8 hours
- 16 hours

- b. Change how long the system waits before reporting certain incidents.

Device timers trigger when certain conditions are detected, when a corresponding incident is opened, or both.

- Clear the default value to unlock the value slider.
- Change the **Processes Not Running** value (Default: 3 minutes, Range: 3 - 60 minutes). This timer determines the "hysteresis count" for generating process down incidents. This is measured in polling cycles internally which roughly translates to minutes for a user (unless the polling cycle is modified).
- Change the **Slicestor Unresponsiveness** value (Default: 30 minutes, Range: 3 - 60 minutes). This timer determines the "hysteresis count" for generating Slicestor device unresponsiveness incidents. This is measured in polling cycles internally which should roughly translate to minutes for a user (unless the polling cycle is modified).
- Change the **NTP Sync** value (Default: 120 minutes, Range: 3 - 180 minutes). This timer determines the "hysteresis count" for generating NTP sync incidents. This threshold is for syncing with NTP server(s) configured in /etc/ntp.conf (NOT the event for when the local clock is not synced with the Manager clock). This is measured in polling cycles internally which should roughly translate to minutes for a user (unless the polling cycle is modified).
- Change the **Rebuilder Agent Hung** value (Default: 1 hour, Range: 1 - 24 hours). This timer determines the amount of time the rebuilder agent waits before sending an event which opens an incident.

- c. Set the default destination rate limit for reallocation, which can be limited to minimize impact on I/O performance. (Default: 20 MB/s.

- 1) Clear the check box to disable rate limiting by default.
- 2) Change the default rate limit value.

- d. Set whether the Accesser Application enforces vault quotas. (Default: Disabled.)

- 1) Select the check box to enable the Accesser Application to retrieve information from the Manager to enforce vault quotas.

Note: This allows anonymous access to two APIs containing potentially sensitive information.

- e. Enable protection against cross-site request forgery attacks. (Default: Disabled.)

For more information, see "Enabling Cross Site Request Forgery protection" on page 79.

- f. Select the preferred IP family.

This setting prioritizes which family of IP address to use when establishing connections in the system. The preferred IP type is attempted first, and if it fails, it falls back to the alternate IP type address. This preference only applies for devices setup with dual stack addresses. For single stack devices, IPv6 or IPv4, the preference has no effect and uses the one address available on the device.

Note: IPv6 is available for ClevSO 3.10.1 and newer. Once the system has been upgraded, devices can be configured to use IPv6 stack in single or dual stack mode.

- 1) **IPv4**
- 2) **IPv6**
3. Determine which action you want to take:
 - Click **Update** to apply the changes.
 - Click **Cancel** to leave the Properties untouched.

Enabling Cross Site Request Forgery protection

The Manager can protect against Cross Site Request Forgery attacks. This protection is disabled by default.

About this task

When enabled, all POST requests from forms in the Manager Web Interface use a double-submit technique. The web browser submits a randomly generated CSRF token to the Manager as both a hidden form input and as a cookie. Then, the Manager application verifies that these two values are equal. This ensures that requests are being processed only from users interacting with the Manager Web Interface and not as the result of an attacker auto-submitting forms from a malicious third-party website without the user's knowledge or consent.

Note: CSRF Protection is only enabled for HTTP POST requests. HTTP GET requests are exempt from CSRF Protection with no CSRF token needed.

Enabling CSRF Protection might interfere with existing workflows, which use the Manager REST API to issue POST requests because existing REST API clients would not be aware of the details of the new CSRF protection feature. For systems that would be impacted by this change in behavior, several options are available:

1. Do not enable CSRF Protection. It is disabled by default. Any new or upgraded system is not impacted in any way when an administrator never enables this feature on the manager.
2. When enabling CSRF Protection, you can also enable the bypass option with settings relevant to your workflow.
 - If the CSRF Protection Bypass Header Name is set to "Authorization", any incoming requests that contain HTTP Basic Authorization credentials are exempted from CSRF Protection enforcement.
 - If the REST API clients have a known request header that can distinguish them from a browser request, configure this header as the exemption rule. Set CSRF Protection Bypass Header Name to "User-Agent" and set CSRF Protection Bypass Header Value to "curl/.*" to allow requests to be made from curl by using its default user agent header.
 - Add a custom header to REST API client requests and then add an exemption for that header. For example, include the header "X-My-Header: <any value>" in all requests and then set the CSRF Protection Bypass Header Name to "X-My-Header".
3. Modify the API client to mimic the double-submit behavior of the Manager UI. Generate any value (the exact value is unimportant) and submit it both as a "csrfToken" request parameter and as a "csrfToken" cookie value.

To enable CSRF Protection, use the following procedure.

Procedure

1. Click the **Administration** tab.
2. Click **Configure** in the **System Properties Configuration** section.
3. Check the **Enable CSRF Protection** box in the **CSRF Protection** section.

4. To allow bypass of CSRF Protection based on the presence of a specific HTTP request header, check the corresponding box.
 - a. Enter a CSRF Protection Bypass Header Name.
 - b. Enter a CSRF Protection Bypass Header Value.

Configure advanced system settings

This option allows detailed modification to the internals of the software and must be used with caution and under the guidance of IBM Customer Support.

CAUTION:

Contact IBM Support before you modify anything.

Configure advanced system properties.

1. Go to the **Administration** tab.
2. Click **Configure** on the **System Advanced Configuration** action bar.
3. Three subcategories appear on the **System Advanced Configuration** page.
 - **Detailed System Advanced Configuration**
 - **Selector-Based Advanced Configuration Rules**
 - **Existing Detailed Configuration Rules**
4. Enabling **Detailed System Advanced Configuration** presents administrators with an advanced configuration options box on the individual **Edit Device**, **Edit Access Pool**, and **Edit Storage Pool** Pages.

To enable this option, check the **Enable Detailed System Advanced Configuration** check box.

1. Click **Update**.

Note: The **Selector-Based Advanced Configuration Rules** should be entered with the guidance of IBM Customer Support. Configuration rules appear at the bottom of the page under the **Existing Detailed Configuration Rules** section.

Configuring SSH keys

Configure device SSH authentication by choosing whether to require the current localadmin password when changing the password, adding SSH keys, or removing SSH keys.

Procedure

1. Click the **Security**.
2. In the **Device SSH Authentication** section, click **Configure**. Configured keys are shown. The Device SSH Configuration page opens. The password requirement and any configured keys are shown.
3. In the **Password Requirement** section, click **Change**.
4. Choose whether the current *localadmin* password is **Required** or **Not Required** when changing the password and click **Update**.
5. In the **SSH Keys** section, click **Change**.
6. Click **Add SSH Key**.
7. Paste a new SSH public key into the **SSH Key** field in OpenSSH ssh-rsa format. You can add one or more keys at a time. Separate each key with a new line. Duplicate keys are not allowed.

For example, if you use **ssh-keygen** to generate your SSH key, the public key value is typically contained in the file at `~/.ssh/id_rsa.pub`.
8. Click **Update** to add the SSH key.
9. Click **Remove** to delete an existing key.

Set license keys

Access to some system features are enabled through a license key. No license keys are currently required.

A license key, provided by IBM Customer Support, is 18-characters long and case insensitive. Spaces or dashes are also optional.

Enter the Key and click **Update**.

Details of the current license and enabled features are summarized.

Configuring notifications

The Notification Service generates a notification whenever an object is written, overwritten, or deleted using the S3 protocol. The Notification Service utilizes an Apache Kafka Producer to publish a stream of notifications to a Kafka topic. Each notification represents a single object-modifying event. You can then use a Kafka Consumer to consume the object change records and take any necessary actions.

About this task

For more information on the Notification Service, see the Notification Service Feature Description Document.

Procedure

1. Click the **Administration** tab.
2. In the Notification Service section, click **Configure**.
3. On the Notification Service Configurations page, click the name of an existing configuration to view details.
4. To create a new configuration, click **Add Configuration**.
5. In the General section, enter the following information:
 - **Name:** The name of this configuration.
 - **Topic:** The name of the topic on the Kafka cluster that notifications will be sent to by default.
 - **Hostnames:** A list of Kafka endpoints in host:port format. A single node cluster has a single entry, but larger clusters could have multiple nodes.
 - **Type:** The type of configuration.
6. Optional: In the Authentication section, select **Enable authentication**. Enter your Kafka username and password.
7. Optional: In the Encryption section, select **Enable TLS for Apache Kafka network connections**.
 - a. If the Kafka cluster is encrypted using a self-signed TLS certificate, paste the root CA key for your Kafka configuration in the **Certificate PEM** field. For more information, see TLS encryption.
8. Click **Save**. A banner appears on the Notification Service Configurations page indicating that the configuration was created successfully and the configuration is listed in the Notification Service Configurations table.
9. To assign vaults to the configuration, click the name of the configuration in the Notification Service Configurations table.
10. In the Assignments section, click **Change**.
11. In the Not Assigned tab, select a vault or vaults and click **Assign to Configuration** to add those vaults to this configuration. You can filter the list of available vaults by selecting or typing a vault name in the **Vault** field.

Note: Notification configurations cannot be assigned to container vaults, mirrored vaults, vault proxies, or vaults that are involved in migrating data. Once assigned to a configuration, this vault

cannot be used in a mirror, with a vault proxy, or for data migration.

Notifications are sent only for new operations that occur after the vault is assigned to the configuration. The Assigned count updates to reflect the number of vaults you assigned to the configuration.

12. Click **Update**.

Editing or disabling a notification configuration

Modify the notification configuration's general, authentication, or encryption information and vault assignments or disable the configuration. The configuration settings are preserved while the configuration is disabled.

About this task

For more information on the Notification Service, see the .

Procedure

1. Click the **Administration** tab.
2. In the Notification Service section, click **Configure**.
3. In the Notification Service Configurations table, click the name of the configuration you want to edit or disable.
4. Click **Change**.
5. To modify the configuration:
 - a. Edit the information you want to change, then click **Update**.
 - b. In the Assignments section, click **Change**.
 - c. In the Assigned tab, select a vault or vaults and click **Unassign from Configuration** to remove those vaults from this configuration. The Not Assigned count updates to reflect the number of vaults you removed from the configuration.
 - d. In the Not Assigned tab, select a vault or vaults and click **Assign to Configuration** to add those vaults to this configuration. The Assigned count updates to reflect the number of vaults you added to the configuration.
6. To disable the configuration:
 - a. Deselect **Enable configuration**. A warning dialog box appears stating that disabling a Notification Service Configuration might cause application inconsistencies.

Note: When a configuration is disabled, no notifications are sent and they are not queued up for future reference. Any operations that occur after the configuration is disabled do not generate notifications.

- b. Click **OK**.
7. Click **Update**.

Deleting a notification configuration

Deleting a Notification Service configuration disassociates it from any assigned vaults, stopping notifications for those vaults from that point on.

About this task

For more information on the Notification Service, see the .

Procedure

1. Click the **Administration** tab.
2. In the Notification Service section, click **Configure**.

3. In the Notification Service Configurations table, click the name of the configuration you want to edit.
4. Click **Delete Configuration**.
5. Enter your current password and click **Delete**. A banner appears on the Notification Service Configurations page indicating that the configuration was deleted successfully and the configuration is removed from the Notification Service Configurations table.

Chapter 6. Monitoring the system

Monitoring the system, software, and hardware components helps to verify that tasks have completed successfully. If a problem occurs during a lengthy operation, monitoring the task ensures quick detection and resolution.

Monitored components

Through the monitoring application, reporting and notification of events and performance statistics are available.

The following components are reported:

- Device
- Vault
- Site
- Cabinet
- Storage Pool

Note: When Automatic Logout is not enabled, these statistics are updated every 5 minutes, so a short delay in graphing will happen after start. For more information, see “Configure preferences” on page 76.

Device health summary

The overall health for each device is displayed in the **Devices** tab of the navigation tree. By selecting a particular device, a summary of its health is displayed with some additional details while in the **Monitor** page.

Devices that are up and functional have an icon in front of the device that appears green with an arrow that is pointing up. This icon indicates that the Manager is able to communicate with the device and all critical processes are running.

Devices in a critical state have a red icon with an "x" in it. This icon appears if the Manager cannot communicate with the device and one or more critical processes are not running. In some cases, a gray icon appears which means that this device is not polled. The letter to the left of the icon indicates the device type.




















Within the device section of the left navigation tree, for a multi-node Slicestor device configuration, individual nodes are considered logical Slicestor devices and are treated like any other non multi-node Slicestor device. It means that one entry exists in the tree per node.

Device Health

The device health includes an overall summary, more details about the symbols, and associated interpretations.

In addition to the device health summary, the following details are provided: device ping time, SNMP response, checks for critical system processes that should be running on the device, and whether the system time is synchronized between the Manager and the device. [Time synchronization does not apply to Network Nodes.] Of these items, the SNMP response and critical system process checks directly contribute to the device health. If either is red, then the device health appears red.

The small letter next to the device symbol indicates type: [A = Accesser Device; M = Manager; N = Network Node; S = Slicestor Device; F = File Accesser Device; C = SMC Device]

Summary	
M 	The Manager is able to communicate with the device and all critical processes are running. The number of failed drives is less than the warning threshold.
S 	The number of failed drives is equal to or greater than the warning notification level.
G 	The Manager either cannot communicate with the device and one or more critical processes are not running and the number of failed drives is equal to or greater than the error notification level.
S 	The Manager has not polled this new device.
Details	
	Device must be reimaged (after data is evacuated).
	Device Ping Time (time)
	Device is not reachable by ping.
	Device has not been pinged or polled.
	Device is Responding to SNMP.
	Device is not responding to SNMP.
	All critical system processes are running.
	Not all critical system processes are running properly (see Event Console).
	System time is synchronized.
	System time is not synchronized [more than 10 seconds out of sync].
	All drives are online.
	Device's drives are in a "Not Good" state.
	Device is either S  or S  depending on the drive health configuration setting.
	Source and Destination devices for Data Evacuation; no impact on device health.

Note: The other devices use the Manager as a time reference, and use NTP (Network Time Protocol) for time synchronization.









Monitor device health

The following components appear as part of health.

- Device ping time.
- Device reporting status.
- Critical system processes running.
- System time is synced (devices other than the Manager).
- Disk Online (Slicestor device specific).

The device icons for each item are as follows.

Table 13. Icons displayed in the Device Health Summary

Icon	Detail
	Device ping time.
	Device is not reachable by ping.
	Device is not pinged or polled.
	All critical system processes are running.
	Not all critical system processes are running properly. (Check the Event Console).
	System time is synchronized.
	System time is not synchronized (more than 10 seconds out of sync).
	All drives are online.

Monitor storage pool health

Storage Pool health status reports on storage pool, vault, and device health. Summaries for sites and storage pools are shown as red, yellow, or green:

Table 14. Storage Pool health status

Status	Detail
Green	All vaults are healthy in the storage pool.
Yellow	If any vault in the storage pool reports as degraded, then the storage pool reports as degraded.
Red	If any vault in the storage pool reports as failed, then the storage pool reports as failed.

Note: In some circumstances, Slicestor devices in a Concentrated Dispersal storage pool can appear "red" because of disk errors, but the storage pool health appears "green" since the health of all vaults on that pool are all "green."

Monitor vault health

Vault health status reports on both vault health and device health for devices that are associated with the vault. Summaries for sites, pools, and vaults are shown as red, yellow, or green.

Vault Health

Data Availability															
									Threshold		Write Threshold		Alert Level		Width
R	R	R	R	R	R	R	R	R	R	R	Y	Y	Y	G	G
1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Data Unavailable									Data Readable		Data Writeable				
Alert Generated														No Alert	

	The number of available Slicestor® Devices ¹ is above the alert level.
	The number of available Slicestor® Devices ¹ is between the alert level and the write threshold.
	The number of available Slicestor® Devices ¹ is below the write threshold.

¹ For multiple Storage Pools, the threshold is calculated against the worse case pool.

In addition to the threshold indicators mentioned, separate icons indicate the status of the deployed Accesser® Device.

	All deployed Accesser® Devices ² are available.
	Some deployed Accesser® Devices ² are available.
	No deployed Accesser® Devices ² are available.

² Note that these status indicators do not take into account embedded Accesser service.


Monitor site health

Site Health is based on device availability.

Table 15. Monitor site health status

Status	Detail
	Site has all devices up.
	Site has at least one device up and at least one device down.
	Site has all devices down.

Table 15. Monitor site health status (continued)



Status	Detail
	Site has no devices.

Devices at this site section list all the devices or a grouped view of devices in the site.

Note: To be more scalable with increasing number of devices, the devices are grouped based on device health and device type (in the same order as in left navigation bar). When grouped, they are routed to the deviceSummary page with the necessary filters applied. The device threshold to switch to a grouped view is more than eight devices at that site.

Accesser devices that are configured for high availability (HA) have a unique set of icons to indicate the health of the pair.

Table 16. Accesser high-availability device pair health status

Icon	Detail
	The device is paired with another device as a HA failover pair. The Manager is able to communicate with the devices and all critical processes are running.
	The pair either cannot communicate with the device and one or more critical processes are not running.

Within the sites section of the left navigation tree, for a multi-node Slicestor device configuration, nodes are grouped since they are inside the same physical box. The name is displayed as **Multi-Node Chassis - {x}**, where {x} is the chassis ID for the physical box that contains the nodes.

Monitor drive health

To aid in monitoring and troubleshooting a front view exists for all Slicestor devices on the **Monitor** page. Click a drive bay to provide additional information on drive health.

On select supported hardware, the bay status information and available bay action are displayed.

Note: For a multi-node Slicestor configuration, the front view displays all the disks that belong in the same chassis. The node that is selected is highlighted, but the others are shaded out. The disk status is visible on all nodes, and disk lifecycle management actions (see “Drive lifecycle states”) are available as well.

Drive lifecycle states

Manager monitoring reports states for data drives, healthy or not, per Slicestor device. Alerts appear in the **Event Console** when user intervention is needed.

Table 17. Drive Lifecycle States

Drive State	Status LED Color	Disk Condition
ONLINE	Green	Healthy, usable.
MIGRATING	Yellow	Failed manually after quarantine. Triggers data move to other disks.

Table 17. Drive Lifecycle States (continued)

Drive State	Status LED Color	Disk Condition
DIAGNOSTIC	Yellow	Temporarily inaccessible. Can be used for running diagnostics, tests, and so on.
FOREIGN	Yellow	Identified with a different Slicestor device and cannot be used in this Slicestor device.
OFFLINE	Gray	Unavailable to use, either physically removed or undetected by the system due to inconsistent hardware state.
UNUSABLE	Red	Kernel cannot initialize for use.
FAILED	Red	Disk is failed and needs to be replaced. It shows in the Failed FRU Report.
UNKNOWN	Gray	State undetermined. Call IBM Customer Support for more assistance.
INIT	Gray	Formatting or reformatted. Can also appear briefly during upgrade or when a Slicestor device or process is restarted.

Drives can change through various states over the course of their useful life.

Note:

Disk Lifecycle operations should be performed if the dsnet-core process is running. Check the **Monitor** page for the Slicestor device. If the following message appears:

dsnet-core is expected to be running

but it is not with a red "x", wait for 15 minutes.

If this message does not go away after 15 minutes, contact IBM Customer Support.

Drive Impact on Device and Vault Health

As part of disk management, Slicestor disk health thresholds can be configured at the warning and error levels. The configuration of these thresholds can be performed for each model globally via the **Administration** page or at the individual Slicestor device level in the Manager Web Interface. Defaults are defined for each model. The device and vault health are a function of these thresholds.

The following disk states contribute to the warning and error thresholds:

- DIAGNOSTIC
- MIGRATING
- OFFLINE
- FOREIGN
- UNUSABLE

If the total number of non-healthy drives is greater than or equal to the warning threshold but less than the error threshold, the device health is set to yellow. If the total is greater than or equal to the error threshold, the device health is set to red.

Note: The vault health is affected if the error threshold is achieved.

Thus, the warning threshold does not impact the vault health.

With the aid of IBM Support, caching (1-wide) vaults can be created. The vaults are intended to be used along with a high reliability vault configuration within the same storage pool (although they can be used in a separate storage pool) for improving read performance. If the device health for a caching vault becomes red (due to disk failures or other issues), an incident appears stating the vault is unusable due to the number of accessible Slicestor devices that are falling below read/write threshold. Although data on the unhealthy drives is lost, it can still be obtained from the high reliability vault. Contact IBM Customer Support for more information, including reliability and performance tradeoffs.

Drive lifecycle state descriptions and troubleshooting

ONLINE

The normal state that most disks are in for most of their life. This state indicates that the disk is functioning normally and that all data operations are supported on the disk.

INIT

The disk state shows Initializing when the disk is formatting or reformatting. When complete, the disk changes state to ONLINE unless a disk health issue is detected.

OFFLINE

If a disk is physically removed or undetected by the system due to inconsistent hardware state, its state changes to OFFLINE. In an OFFLINE state, read and write operations to the disk are not supported. Missing slices that were written to the system while the disk was OFFLINE are rebuilt on the disk.

Recommended user remediation options

- Reinsert the OFFLINE disk.
- Replace the OFFLINE disk with a new disk to trigger rebuild of slices on the new disk.
- Leave the disk in an OFFLINE state. Use the dispose command on the Slicestor console or on the Manager to ensure data reliability for the data slices on the OFFLINE disk.

Note: Other disks within the Slicestor device become responsible for slices on an OFFLINE disk. A replacement disk cannot be used immediately. The Slicestor device starts using the new disk when the used space on other disks exceeds a system-defined threshold.

DIAGNOSTIC

A disk can change from an ONLINE state to a DIAGNOSTIC state if the application detects problems that are related to disk performance or significant number of I/O errors. It can also change into this state if it fails a SMART test or SMART attributes fall below threshold. The purpose of the DIAGNOSTIC state is to allow the operator to investigate whether the disk needs to be replaced or if it is healthy enough to resume operation within the system.

In a DIAGNOSTIC state, read and write operations to the disk are not supported. Existing slices on the disk are preserved indefinitely.

DIAGNOSTIC disks are considered "quarantined" and are occasionally referred to as so throughout the documentation.

CAUTION:

Leaving a disk as DIAGNOSTIC for an extended duration affects data reliability and is discouraged.

Recommended user remediation options

See Table 18 for recommended user remediation options.

Note:

If multiple disks in the same Slicestor change to DIAGNOSTIC state at the same time, it can indicate a backplane, disk controller, or other component failure. Contact IBM Customer Support if it happens.

When a DIAGNOSTIC disk is pulled and reinserted, a new event appears.

Disk in drive bay X with S/N Y is diagnostic state, due to Z.

To find more details about the reason, go to **Advanced Search** in the **Event Console** and specify the disk serial number. Identify the corresponding quarantined disk event, which has the reason.

Diagnostic states and recovery actions

The event on the UI indicates why the disk was placed into a DIAGNOSTIC (or in earlier releases, quarantined) state. Multiple reasons can exist for placing a disk into one of these states, but the software presents only the first reason it detects.

The version of ClevOS that a device uses determines the event description that is displayed. Refer to the following tables for the event descriptions specific to the version of ClevOS in use on the impacted device.

Note: Contact IBM Customer Support if the following situations occur:

- The drive does not change to an ONLINE (or in earlier releases, a good) state after you take the recommended actions in the following tables.
- The failure rate of data disks in a storage pool or set is higher than 4% AFR (Annual Failure Rate). See AFR Rate Calculation for details.
- Multiple data disks in a storage pool, set, or Slicestor® device fail in a coordinated fashion, which might indicate a higher-order failure.
- The Event Console contains an event description not included in the following tables.

Include failed drive serial numbers and event description text (if applicable) with all RMA requests.

Table 18. ClevOS 3.10.1+ DIAGNOSTIC state recovery actions

Event Description	Recommended Action
User initiated action.	Resume the disk through the IBM COS Manager UI or by using the CLI command: # storage resume <disk-id>
Possible storage metadata issues detected.	1. Resume the disk through the IBM COS Manager UI or by using the CLI command: # storage resume <disk-id> 2. If the disk is not in an ONLINE state after it is resumed, then follow these steps: a. If the storage pool or set that the disk is part of is less than 95% used: 1) Fail the disk through the IBM COS Manager UI or by using the following CLI command: # storage fail <disk-id> 2) When the failing disk migration finishes, replace the disk. b. If the storage pool or set that the disk is part of is greater than or equal to 95% used, replace the disk with a new disk.
I/O timeouts exceeded threshold.	

Table 18. ClevOS 3.10.1+ DIAGNOSTIC state recovery actions (continued)

Event Description	Recommended Action
I/O errors exceeded threshold.	If the storage pool or set that the disk is part of is less than 95% used: 1. Fail the disk through the IBM COS Manager UI or by using the following CLI command: # storage fail <disk-id> 2. When the failing disk migration finishes, replace the disk. If the storage pool or set that the disk is part of is greater than or equal to 95% used, replace the disk with a new disk.
Drive attribute exceeded threshold.	
Possible file system issues detected.	
Possible software issue detected.	Contact IBM Customer Support.

Table 19. Pre-ClevOS 3.10.1 quarantine state recovery actions

Code	Event Description	Recommended Action
1	A SMART failure.	If the storage pool or set that the quarantined disk is part of is less than 95% used: 1. Fail the disk through the IBM COS Manager UI or by using the following CLI command: # storage fail <disk-id> 2. When the failing disk migration finishes, replace the disk. If the storage pool or set that the disk is part of is greater than or equal to 95% used, replace the disk with a new disk.
2	A SMART command failure.	
4	Excessive I/O errors on disk.	
3	User initiated operation.	Resume the disk through the IBM COS Manager UI or by using the CLI command: # storage resume <disk-id>
5	Excessive timeouts on disk.	1. Resume the disk through the IBM COS Manager UI or by using the CLI command: # storage resume <disk-id> 2. If the disk is not in a good state after it is resumed, then follow these steps: a. If the storage pool or set that the disk is part of is less than 95% used: 1) Fail the disk through the IBM COS Manager UI or by using the following CLI command: # storage fail <disk-id> 2) When the failing disk migration finishes, replace the disk. b. If the storage pool or set that the disk is part of is greater than or equal to 95% used, replace the disk with a new disk.
6-9, 11-13	An invalid internal structure on the data drive.	
All other codes		Contact IBM Customer Support.

AFR Rate Calculation

IBM Cloud Object Storage System is designed with a reliability calculation that assumes a 4% disk annual failure rate (AFR).

Monitor the disk failure rate to spot trends of increasing AFR over a month-to-month basis. The best method to track this trend is to quantify the rate of monthly disk failures that a particular storage pool or set can sustain, while not exceeding the 4% AFR. You can accomplish it by using the formula:

$$\# \text{ of drives failures per month} = (4\% \text{ AFR}) * (\# \text{ of disks in the Set}) / (12 \text{ Months})$$

Example - If 1200 disks exist in a set and the wanted AFR is 4%, then the monthly disk failure rate should not exceed $[(0.04) * (1200) / 12] = 4$ disks per month.

Alternatively, another approach is to calculate the AFR based on the number of disks that failed relative to the total population of disks in the storage pool or set yearly. Calculating AFR over other durations of time can be accomplished by considering the time duration relative to a year.

Example 1 - If 36 disks failed out of a population of 1200 disks in a Set within 12 months, the AFR would be $[(36/1200*100)*(12/12)] = 3\%$ AFR.

Example 2 - If 36 disks failed out of a population of 1200 disks in a Set within 6 months, the AFR would be $[(36/1200*100) * (12/6)] = 6\%$ AFR.

Contact IBM Customer Support if the monthly drive replacements exceed the expected rate or the monthly AFR is trending upwards toward 4% AFR.

Failure reason codes

The disk failure code is intended to give an indication for why the disk failed.

Table 20. Failure reason codes

Code	Reason Text	Details
0	I/O check failed.	Cannot communicate with drive to determine identification characteristics.
1	Unable to build platform metadata.	Cannot build a partition table or sign the drive.
2	Unable to create file system.	Cannot format the remainder of the disk with a file system.
3	Data migration completed.	A disk that is failed by the user and they attempted failure migration.
4	No data migration attempted.	A disk that is failed by the user and requested not to attempt failure migration.
5	Drive has preexisting data.	

MIGRATING

A disk is reported as MIGRATING after a DIAGNOSTIC disk is manually failed. After a disk is marked as MIGRATING, data migration begins. Data migration moves data from the MIGRATING disk to neighboring drives. This process reduces the amount of data that needs to be rebuilt when the disk is replaced. The disk is permanently removed after all data is migrated.

Note: During upgrade, disk migration aborts and does not resume when the upgrade completes.

Recommended user remediation options

Migrating data can take some time. It depends upon the amount of data to be moved. In general migration is preferred as fewer system resources are involved to recover slices when compared to rebuilding the entire disk. IBM suggests monitoring the migration process closely and aborting the operation if performance degrades. The dsnet-core process must be restarted to abort migration.

Replace the disk after data is migrated and the disk is permanently removed.

To abort the migration process, do the following steps.

1. Go to the **View from the Front**.
2. Click **Stop Migration** for the MIGRATING disk.
3. When the Slicestor device comes back online, the disk is reported as FAILED and the disk can be replaced.

Note: Contact IBM Customer Support for more guidance if you run into this issue.

UNUSABLE

A disk can change to an UNUSABLE state if the OS fails to initialize the disk on startup or when the disk is inserted into the Slicestor device.

If a disk is in the UNUSABLE state, you can issue a password-protected Reset operation through the Manager Web Interface that starts a disk reformat. The functionality is equivalent to the **nut storage** reset command.

A **Reset** action button is available for UNUSABLE disks in the physical disk layout. This action reformats the disk. The Operator role cannot perform this action. A password is needed to proceed with the reset operation.

In an UNUSABLE state, any slices on the disk are inaccessible. When a disk changes from a DIAGNOSTIC state, the process of rebuilding missing slices is in progress on other disks.

Recommended user remediation options

- Replace the UNUSABLE disk with a good disk.
- Attempt to Reset the disk under the following scenarios.
 1. Insertion of a failed drive
 2. Insufficient initial drive state (for example, refurbished drive or incorrect factory initialization)

Note: Since other disks within the Slicestor device become responsible for slices on a pulled disk, a replacement disk cannot be used immediately. The Slicestor device starts using the new disk when the used space on other disks exceeds a system-defined threshold.

FOREIGN

A disk can be reported as FOREIGN if it is removed from one Slicestor device and inadvertently placed in a different Slicestor device. I/O operations are not allowed on a FOREIGN disk.

If a disk is in the FOREIGN state, you can issue a password-protected Reset operation through the Manager Web Interface that starts a disk reformat. The functionality is equivalent to the **nut storage** command.

A **Reset** action button is available for FOREIGN disks in the physical disk layout. This action reformats the disk. The Operator role cannot perform this action. A password is needed to proceed with the reset operation.

Reset FOREIGN disk

In the **Event Console** the messages appear in the following sequence.

- A disk was detected as FOREIGN.
- An event is sent to the console that the disk was reset by the user who initiated the action.
- After the disk is reset, it is reported as a new disk.

It is recommended to reset the disk as soon as possible to minimize the amount of rebuilding and protect the overall data reliability.

UNKNOWN

UNKNOWN state is reported when the disk returns an unexpected message.

Note: The event can clear on its own. It is recommended to call IBM Customer Support to report the problem.

When a disk is moved from UNKNOWN to ONLINE, the message in the **Event Console** is cleared. The view from the front status color is also changed from gray to green.

RAID states

The Manager separates RAID member states from the state of the entire array.

The RAID array health reports as OPTIMAL when all the disks are ONLINE.

The entire array is reported as DEGRADED when any RAID member is reported as OFFLINE.

A RAID member reports REBUILD when it is replaced. The incident for the entire array is only cleared when all RAID members report as ONLINE and when the array reports as OPTIMAL.

Event console

Incidents are separate from other events. This area focuses attention on incidents - stateful events that require operator intervention. The event content consists of two sections: **Open Incidents**, and **Event Console** (event logs). The **Event Console** displays events based on context that is defined by filters that are selected in **Advanced Search** (default time frame is the last two weeks).

Open incidents

The **Open Incidents** section highlights events that require intervention. After the issue is resolved, the incident will close; most will close automatically.

In some cases, manual intervention might be needed by IBM Customer Support to close an incident.

The incidents are grouped under a corresponding Manager device, Storage Pool, or Access Pool. Each incident group can be opened and closed by using the open/close icon. The incidents are also categorized into four types:

- System - Incidents that are related to the system as a whole. Examples include backup failures and Slicestor device to Slicestor device communication issues.
- Hardware - Incidents that are related to a specific device's hardware. Examples include high file system usage, high CPU temperature, and power supply fail.
- Software - Incidents that are related to a specific device's software. Examples include processes not running and the device in an inconsistent state.
- Disk - Incidents that are related to Disk Lifecycle Management. Examples include when a disk is DIAGNOSTIC and when a disk is FAILED.

Each category has a corresponding filter at the top of the section. Clicking the check box shows/hides any incidents that match the respective category.

In the following scenario, a disk is quarantined, pulled, permanently removed, disposed, Slicestor device is powered down, disk that is replaced, Slicestor device that is powered on, an incident that indicates:

Disk in drive bay 1 with S/N 9QJ1SRR5 is a previously removed disk will remain in the Open Incident view of the {mgr}.

Note: If all disk issues are addressed and a mis-match persists between disk states that are presented on the Manager Web Interface disk diagram and the **Open Incidents** view, contact IBM Customer Support. Manual closing of disk incidents might be needed.

Events

When certain conditions are detected, events are generated and displayed in the **Event Console**, which is accessible from both the Home and **Monitor** pages.

The console has a number of features to aid in troubleshooting and incident management. The displayed events are context-specific; events can be displayed for the overall system and its components. Event severity is displayed as follows.

1. **Critical**
2. **Error**
3. **Warning**
4. **Info**

Manager events include status information from all nodes in the system. Events are ordered by most recent.

Note: When you upgrade a device to ClevOS 3.10.1 or later, existing disk lifecycle management events become part of the incident state cycle. New events that are opened after upgrade are added to the old state cycle and indicate that the disk was previously quarantined.

Audits also appear in the **Event Console**. Audits are special events, which have no severity level, and are only visible by users with the role of Super Admin, System Administrator, or Security Officer.

For audit messages, an empty string, denoted by ' ', indicates that the specified value is unset.

Event search

Event Console

The Event Console is similar to an event log and displays the most recent 50 events (most recent first), based on the context defined by the Current Filters. The Current Filters are selected via the **Advanced Search** (Pressing Advanced Search opens/closes the display). After the filters are selected, press **Search** to view the corresponding events. The default context is a single filter that indicates a time range that represents the last week. Selecting the “x” on a filter removes the filter and automatically initiates a new search, establishing a new context. Click **Show Audits** to include 30 days of audit information with the events. Audits can be filtered via **Advanced Search**.

Click **Show More** to display more events in increments of 50. Use the **Remove/Add** scrollbar to hide/add back the inner scroll bar in the display.

Select **Export** to create a .csv (comma-separated values) file for use with spreadsheet applications. The export file, limited to the most recent 50000 events, contains all events, regardless of the severity filter setting. Events can also be forwarded to an email account based on severity, frequency, and timing. Use the **Preference** menu to configure this function to the wanted operation.

Note: A New Event Count (#) prefix is added to the HTML title tag in the browser when new events exist in the **Event Console**. After the **New Events** link is selected, the count disappears until a new event occurs.

In **Advanced Search**, the Message text box accepts standard text or regular expressions (a special pattern that specifies a set of strings - see http://en.wikipedia.org/wiki/Regular_expression for an overview). When you use standard text (Regular expression box cleared), an “AND” search is done of the individual terms that are provided in the text box. In particular, all terms in any order are returned when the regular expression box is not checked.

Alternatively, a search based on regular expressions can be initiated. It is accomplished by entering a regular expression in the **Message** text box, selecting the **Regular expression** box, and clicking **Search**. Several examples of regular expressions are provided.

Table 21. Regular expression examples






"OR" functions.	usage space (matches events with usage or space).	Note - The vertical bar separates alternatives.
Matching preceding character zero or one time.	file? (matches events that contain fil or file, such as file system).	Note - In addition to "?", the "*" and "+" characters can be used to match a set of strings.
Bracket construct.	slices[te] (matches events with terms such as slicestor, sliceserver).	Note - The items in the bracket are interpreted as "t" or "e".
One term followed by another.	reporting status (matches events with "reporting status").	Note - This pattern represents a request to match based on a specific ordering, one term after another. It is a constrained "AND" search.
Combining parentheses with a preceding element zero or one time.	r(eb)?oot (matches events with root and reboot).	

Numerous other constructs can be used as part of regular expressions, which are not described. An error occurs if an invalid regular expression is provided.

Event Console Details

The display includes event: **Status** [Severity], **Summary** [Description], and **Time** [Occurrence], including amount of time relative to the current time. When selected, a detailed view of an event is shown within this box. An event might occur multiple times. The count appears in a rectangular box (17) next to the event. Older events that were migrated from an earlier Manager version are denoted with an asterisk (*).

Times are shown in GMT (Greenwich Mean Time) by default. Use the **Preference** menu to change the display to local time.

Severity Key -  = Critical,  = Error,  = Warning,  = Information,  = Cleared

Selecting a cleared event displays all related events. Selecting of any of the related events in the **Event Console** displays the same collection of events as the clear event. In addition, the duration from the time of the event occurrence until the time that it is cleared is displayed.

Clicking any event in the **Event Console** displays additional information on the event.

Example - Click a diagnostic disk event to show the suspend reason: List Disk Suspend Codes.

Note: Virtual device disk events reference the block device name name, instead of bay number and drive serial number.

Note: Only a Security Officer account (or a Super User account) can access the **Audit Search** utility.

Attention: Any time a device is added or a vault, site, cabinet, or an administration configuration is changed, the manager must be backed up by using the **Backup and Restore** utility. Permanent data loss can occur if the manager database becomes corrupted. Periodic backups must also be performed to preserve historical statistics and log information. See **Backup** settings on the **Administration** menu.

Regular expression examples

A small sample of the possible constructs for Regular Expressions.

Table 22. Regular Expression Examples

Example of	Code	Description
OR functionality	<code>usage space</code>	Matches events with usage or space . The vertical bar separates alternatives.
Matching preceding character zero or one time	<code>file?</code>	Matches events that contain fil or file , such as filesystem . In addition to ? , the * and + characters can be used to match a set of strings.
Bracket construct	<code>slices[te]</code>	Matches events with terms such as, slicestor , slicesterver , etc. The items in the bracket are interpreted as t or e .
One term followed by another	<code>reporting status</code>	Matches events with reporting status . This pattern represents a request to match based on a specific ordering, one term after another. It is a constrained AND search.
Combining parentheses with a preceding element 0 or 1 time	<code>r(eb)?oot</code>	Matches events with root and reboot

An error is displayed if an invalid regular expression is provided.

Search covers all items, both events and incidents. The severity filter allows for multiple severity types to be identified. Multiple devices, cabinets, sites, and storage pools can be selected, depending on context.

All filters appear on the dashboard and top **Monitor** page, but the **Monitor Storage Pool** page contains Slicestor devices.

As filters are chosen, they are displayed in the **Current Filter** section. It is recommended to specify a time range to reduce the search scope.

Audit search

The audit search functionality – available to users assigned the Super User or Security Officer role – can be found under the Audit tab of the **Advanced Search** portion of the **Event Console**. For a Security Officer's use, the **Security** page includes an **Event Console** that initially filters out all events except Audits.

The Audit utility allows a user to perform a search of audit records based on type of object, user, action code (type of operation performed), source (REST API or UI), and date of the record. Click **Apply Filter** after the categories are selected. Each audit record contains the change date, user name, action, source, and a modification summary.

Click **Export** to create a .csv (comma-separated values) file for use with Excel or other spreadsheet applications. The export file contains the audit records corresponding to the filter setting.

Export

Export creates a .csv file of the current **Event Console** contents.

The file contains:

- Events
- Reported Time (GMT)
- Alias
- Hostname
- Type
- IP
- Site
- Severity
- Stream Type
- Status
- Summary
- Audits
- Change Date (GMT)
- Account (Username)
- Action
- Source
- Modification

Performance graphs

Performance graphs are available for devices and vaults. The graphs allow the user to view the system over six default time ranges and also allow for custom time ranges to be entered.

To access the Device Graphs.

1. Click **Monitor** from the action bar.
2. Click the **Devices** tab in the left navigation tree.
3. Select the wanted **Devices** tab. It generates the performance graphs.

To view graphs for an individual vault.

1. Click **Monitor** on the **Summary** section of the wanted vault. To get to the page, use the left navigation tree.
2. Performance graphs can be expanded or hidden; hidden is the default.
 - Click the plus icon to expand the graph.
 - Click the minus icon on the graph's title to hide the graph.

Standard computer performance parameters are calculated by the device or collected via SNMP, Im-sensors, and so on, from its system board and system-specific information for the Accesser, and Slicestor devices. Axes for all graphs are scaled automatically depending upon the time selected. By default, six-hour graphs are shown. The time range can go from 5 minutes to yearly. The time frame can be selected by the dropdown under the Performance action bar or by changing the **To** and **From** dates to an acceptable range.

Several basic capabilities exist to support further inspection and troubleshooting:

- Select the icon in the upper right corner of the graph to expand it.
- Click **Export** to get data provided in .csv format. This data can be analyzed further by tools.
- Click **Share** to get a link to the current page and time range.
- Click the device or vault links to go directly to their pages. It is only available in some graphs, such as **Scanned Sources** and the **Outgoing/Incoming Rebuild**.
- Click **Hide All** or **Show All** to hide or show all lines.

- Click the line in the legend to hide or show individual lines.
- Click buttons that appear on hover to pan left or right through the data.
- Click and drag on the graph to select a new time range to zoom into based on the range that is represented in the highlighted section.

Vertical Crosshairs are synchronized across all graphs and can be set/unset by clicking any graph.

A zoom in/out widget is available in the upper right corner of each graph to allow for quicker investigation into surrounding data points or diving into a cluster to get a better resolution.

When the graph time range is changed in anyway, the date range picker and all graphs are synchronized to the appropriate range, resulting in the same time window being displayed. Click **Go** to refresh the data in the graphs regarding the current time reference.

Note: In some circumstances, particularly when a device is down, statistics are not reported correctly. It results in gaps or erroneous values in the performance graph for that device. When the device is online, statistics are reported correctly.

For a graph, gaps in the last week, month, and year views can be of different sizes due to granularity. To accurately assess a gap, anything less than a 6-hour view should be used.

Table 23. Available Graphs by component

Graph Type	Storage Pool	Vault	Accesser Device	Slicestor Device	Manager
Storage Pool Capacity and Usage	X				
Raw Space Used		X			
Total Number of Slices per Device		X			
Aggregate Client-Accesser Throughput		X			
Aggregate Accesser-Slicestor Throughput		X			
Client-Accesser Throughput		X			
Accesser-Slicestor Throughput		X			
Scanning Rate (sources/sec)				X	
Rebuild Slices Sent				X	
Rebuild Deletes Sent				X	
Rebuild Bytes Sent				X	
Rebuild Slices Received				X	
Rebuild Deletes Received				X	
Rebuild Bytes Received				X	
Disk Usage			X	X	X
Device Load			X	X	X
CPU Usage			X	X	X
Network Usage			X	X	X
Accesser Requests			X		
Message Acknowledge Time			X	X	
CPU Temp			X	X	X

Table 23. Available Graphs by component (continued)

Graph Type	Storage Pool	Vault	Accesser Device	Slicestor Device	Manager
Fan Speed			X	X	X
Hard Drive Temp			X	X	X

Table 24. Performance Graph Summary

Metric	Units	Description
Storage Pool Capacity and Usage	bytes	Displays the overall raw capacity and storage pool usage over time.
Raw Space Used	bytes	Raw space usage. Note: When the manager is down, gaps appear in this graph.
Total Number of Slices per Device	slices / device or slices / pillar	The storage amount in a device that is measured in slices.
Aggregate Client to Accesser Device Throughput	bytes / second	Aggregate rate at which data is traveling (reads and writes) between the external client and the Accesser device. This graph is not defined if an Accesser device is not deployed for this vault.
Aggregate Accesser to Slicestor Device Throughput	bytes / second	Aggregate rate at which data is traveling (reads and writes) between the Accesser device and the Slicestor devices. This graph is not defined if an Accesser device is not deployed for this vault.
Client to Accesser Device Throughput	bytes / second	The rate at which data is traveling (reads and writes) between the external client and the Accesser device. This graph is not defined if an Accesser device is not deployed for this vault.
Accesser to Slicestor Device Throughput	bytes / second	The rate at which data is traveling (reads and writes) between the Accesser device and the Slicestor devices. This graph is not defined if an Accesser device is not deployed for this vault.
Scanning Rate	sources / second	The rate at which the scanning agent of one or more devices are scanning the storage for missing writes/deletes.
Rebuild Slices Sent	slices / second	The rate at which the rebuilding agent of one or more Slicestor nodes have reconstructed missing writes and have sent to a destination Slicestor node. Measured in slices/sec.
Rebuild Deletes Sent	slices / second	The rate at which the rebuilding agent of one or more Slicestor nodes have discovered missing deletes and have sent to a destination Slicestor node.
Rebuild Bytes Sent	bytes / second	The rate at which the rebuilding agent of one or more Slicestor nodes have reconstructed missing writes and have sent to a destination Slicestor node.
Rebuild Slices Received	slices / second	The rate at which missing writes have been recovered.
Rebuild Deletes Received	slices / second	The rate at which missing deletes have been recovered.
Rebuild Bytes Received	bytes / second	The rate at which missing writes have been recovered.

Table 24. Performance Graph Summary (continued)

Metric	Units	Description
Disk Usage	bytes / second	<p>Provides a consolidated view of disk read/write speeds for all drives in device.</p> <p>Read Show</p> <p>Display all read lines for all drives in the graph.</p> <p>Read Hide</p> <p>Hide all read lines for all drives in the graph.</p> <p>Write Show</p> <p>Display all write lines for all drives in the graph.</p> <p>Write Hide</p> <p>Hide all write lines for all drives in the graph.</p> <p>Aggregate Read</p> <p>Sum all toggled-on read lines and display a single summed line.</p> <p>Aggregate Write</p> <p>Sum all toggled-on write lines and display a single summed line.</p> <p>The legend items can be sorted in three different ways:</p> <ul style="list-style-type: none"> • bay number (increasing) • read speed (decreasing) • write speed (decreasing). <p>Click the headers in the legend to trigger the sorting. The legend is sorted by bay number when the graph is initially loaded.</p> <p>VM devices display SCSI name instead of bay numbers and are sorted in lexicographically increasing order.</p> <p>When aggregate lines are displayed instead of individual lines, the aggregate lines adjust as individual lines are toggled by using the color box in the legend.</p> <p>For export, only toggled-on lines are included in the .csv output file.</p>
Device Load	Processes	<p>Average number of processes that are either in a runnable or uninterruptible state. A process in a runnable state is either using the CPU or waiting to use the CPU. A process in uninterruptible state is waiting for some I/O access, for example, waiting for disk. Load averages are not normalized for the number of CPUs in a system, so a load average of one means that a single CPU system is loaded all the time while on a 4 CPU system it means it was idle 75% of the time. For a lightly loaded device, the vertical axis might show m for milli-percent.</p>

Table 24. Performance Graph Summary (continued)


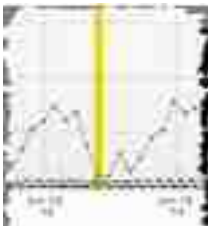




Metric	Units	Description
CPU Usage	%	A measure of how much time the CPUs spend on user applications and O/S (Operating System) functions. Even when the CPU usage is 0% the CPUs are still performing basic system tasks. The graph shows cumulative CPU utilization by the device. Dual and quad core CPUs might report two, four, or even eight CPUs, depending on the specific device. Therefore, loads of greater than 100% are typical. CPU Wait, also shown, represents the percentage of time that the CPU is waiting on I/O, which includes disk, network, and memory delays. If your device's CPU uses hyperthreading of CPUs, it could take the percentage to greater than the total number of CPUs x 100%.
Network Usage	bytes / second	Inbound and outbound network traffic.
Accesser Requests	requests / second	HTTP requests to the Accesser device.
Message Acknowledge Time	ms	Measured from the time an Accesser device is ready to send a data packet until the Slicestor device has written to the operating system and responded back. Long message acknowledge times can indicate slow Slicestor devices. If an Accesser device does not have a connection to Slicestor devices, no data is displayed in the graph during this period.
CPU Temperature		A critical alarm is sent to the Monitor Event Console if the CPU (Central Processing Unit, that is, the server microprocessor) temperature exceeds 90°C. Check the ambient temperature of the device immediately.
Fan Speed	rpm or %	A critical alarm is sent to the Monitor Event Console if the speed exceeds the threshold; the threshold is specific to the device and is set automatically by the system. Excessive fan speed generally indicates an ambient cooling problem or poor air flow around the device. Each fan is shown as a separate line on the graph.
Hard Drive Temp	°C	Depending on the number of hard drives, there can be more than one graph shown. The graph indicates what drives it is displaying in its title and legend. A critical alarm is sent to the Monitor Event Console if the temperature of any disk exceeds 60°C. Check the ambient temperature of the device immediately. Each disk is shown as a separate line on the graph.

Note: In some graphs, such as the Storage Pool Capacity and Usage and Raw Space Used graphs, a temporary drop is observed, particularly during upgrade. After device upgrades are complete, the value returns to normal.

If no vaults are created on a storage pool, the scanning and rebuilder graphs do not appear on the **Monitor Device** page. In addition, a **Toggle Aggregate** button appears on these graphs, which aggregates all lines into a single line, representing the overall performance.

Performance graph features

Each Performance graph has several key features.

Component	Image	Description
Zoom widget		Click the plus and minus buttons in the upper right corner of the graphs to zoom in or out of the current time range.
Crosshair		<ul style="list-style-type: none"> • Move the mouse on any graph to update the synced crosshairs. • Click the plot one time to set the crosshair at that location across all of the graphs. • Click a second time to unlock the crosshair. If the crosshair is locked in a time range, it stays at that location until the graph is clicked again to prevent losing one's place.
Expand Graph Button		Click the icon that is displayed below to enlarge a graph to full screen.
Plot selection		Click and drag over a time range in the graph to zoom into that time range. A light orange highlight is added until the mouse is released, indicating what the new range is. Plot selection range is limited to greater than 5 minutes.
Share functions		Click the Share link to share a graph. A link is generated that can be distributed. When the link is visited the time range visible upon creation of the link is displayed.
Export link		The current time range that is viewed on the graph can be exported as a .csv file and displays two columns for each data series. One contains the date. The other contains the value.
Pan functions		The graph pans left or right approximately 3/4 of the current view to allow for graph traversal in the same time range.

Graph controls

Storage Pool page (scanning, outgoing and incoming rebuild, data reallocation graphs):

A graph control mechanism on the Storage Pool page can be used with Scanning/Rebuilding charts.

Graph control is only visible when the Storage Pool is configured in the following ways:

- For unmerged Storage Pools, there exist one or more Vaults with IDA width that is less than the Storage Pool width.
- For any merged Storage Pools regardless of number of unique IDA widths of deployed Vaults.

Graph control allows the Scanning/Rebuilding statistics to be grouped and rolled-up by logical entities, such as Storage Pool Sets and Stripes, for monitoring its activity at different levels of granularity.

The least granular view is the Scanning/Rebuilding statistics rolled up to the Storage Pool level. In this view, all Scanning/Rebuilding activity is summed across all devices that are associated with the Storage Pool for a comprehensive view of Scanning/Rebuilding performance.

There is a middle view called Stripes. Stripes are grouping of devices within a Storage Pool. Depending on the unique IDA width of deployed vaults in the Storage Pool, there can be one or more numbers of Stripes. Stripes are categorized by unique IDA width of deployed Vaults in a Storage Pool. There is a selection that lists out all the unique IDA widths. When a width is clicked, it refreshes all the open Scanning/Rebuilding graphs to show statistics for all Stripes that exist in the define IDA width.

The most granular view is seeing statistics for individual devices. They can be viewed after a Storage Pool Set is selected (if there is more than one), a unique IDA width (if there is more than one), and a Stripe.

The graph control has the following selections:

- Selection by Storage Pool Set
- Selection by Unique IDA Width
- Selection by Stripe

Monitor Accesser page: message acknowledge time (MAT) graphs:

The MAT graph is enabled with a statistical analysis tool for more efficient troubleshooting and diagnosis.

This tool is available when an Accesser appliance is deployed to multiple Storage Pools, expanded/merged Storage Pools, or to multiple vaults with differing unique widths. Three statistics are available - max, 95th percentile, and average. When selected, it applies the statistical analysis on the devices that is consistent with the entity that is selected. For example, if you select average and a Storage Pool, it averages all the data points of all the devices that belong in that Storage Pool. This analysis is intended to be used to find outliers and anomalies by looking at a summarized view and be able to drill down to find problematic sources. The statistic analysis is disabled when looking at data on a per-device basis.

Monitor Accesser page: Accesser request graphs:

The Accesser request graph displays incoming HTTP requests to the Accesser device.

Above the graph are controls to switch the graph mode and to filter the graph. The graph has two modes: request and response. In request mode, graph lines are displayed for each request type, and the filter control allows the selection of a particular response code. In response mode, graph lines are displayed for each response code, and the filter control allows the selection of a particular request type.

Any unsupported HTTP request types are grouped under a single label: INVALID.

Annotated graphs

A small **Event Console** annotated graph that contains annotation data points appears under each event and audit data graph.

Note: This graph shows data consistent with the currently displayed **Event Console** entries. It is possible to pan the annotated graph to time periods before and after this time frame, but the data will not be refreshed. To show more data, first click the **Show More** link below the list in the **Event Console**. The graph updates automatically when the **Event Console** data refreshes.

Annotation data points are color-coded based on the severity/type of each event. When annotations are clicked, the corresponding entry is clicked/highlighted in the **Event Console** so it can be scrolled into view to correlate events with graph inconsistencies.

Annotated graphs have the same interactive functions as the regular performance graphs (Crosshair, plot selection, data point labels on-hover).

The **Event Console** and the annotated graphs have the same exact content so they are always in sync. Every time the **Event Console** is updated, the annotated graphs reflect the new values. If no **Event Console** data exists, the annotated graphs are hidden.

Disabling events on a device

About this task

Note: To disable events on a device, Event Suppression must first be enabled per “Suppressing device events” on page 68.

Disabling events on a device can be useful when, for example, a bad/misbehaving device is generating numerous events and email alerts.

To enable/disable the ability to suppress events on a device:

Procedure

1. When on the **Monitor** page, under the **Device Summary** section, click the device where you want to disable events.
2. Click **Disable Events** in the upper right corner of the page.

Note: If **Disable Events** is not present, then you need to enable **Event Suppression** per “Suppressing device events” on page 68.

3. Select one of the three options.
 - Hours
 - End-date
 - Indefinite
4. Pick the wanted option and complete secondary fields (if any) that become enabled.
5. Click **Update**.

Device summary

Filtering devices by using tabs

On the **Monitor** and **Configure Device Summary** pages, three tabs at the top of the page provide quick filters to limit the list of devices displayed.

- The left tab is for **All Devices**, which include both SliceStor devices and Accesser devices (Manager devices are not included on the device summary pages).
- The center tab is for **SliceStor Devices**.
- The right tab is for **Accesser Devices**.

Click one of these three tabs to filter the list of devices to include only the ones that are specified on the tab.

Default tab filters

If the user clicks the **Devices** node in the left navigation tree, the **All Devices** tab is selected and no filters are applied.

If the user clicks the SliceStor devices or Accesser devices node in the left navigation tree, the page is rendered with the respective device tab already selected and the devices pre-filtered.

In the quick filters for SliceStor devices and Accesser devices, a pie chart and health bars are also available. Click the segments of the pie chart or the health bar in a tab to filter the list of results. The individual slices of the pie chart change color when the user hovers over them. The individual health bars are highlighted with a colored border when the user hovers over them.

Filtering devices by using extra filters

The **Monitor** and **Configure Device Summary** pages also have a gray box with a collection of filters to limit the list of devices.

- Health by using check boxes
- Sites, Cabinets, Models, Storage Pools, Storage Pool Sets and Access Pools by using drop-down menus

Select one or more of any of the options in the drop-down menus or check boxes to filter the list of devices.

The user can enter extra free text to further limit the results in the **Search Results** field below the gray box. The **Search Results** field accepts one or more terms that are separated with a space. It does not support AND, OR, or None logical operators or keywords.

Apply filters to device list results

All filters are combined when filtering the list ((logical)AND). Within each filter, if more than one option is selected, any of those selected options can be included when filtering the list ((logical)OR).

To limit the devices that are displayed to the ones with the following attributes:

- Warning or Unhealthy states
- Storage Pools ABC or DEF
- The terms Seattle and S3 in their names

Perform the following steps:

1. Select Warning and Unhealthy in the **Health** check boxes.
2. Select ABC and DEF in the **Storage Pool** drop-down menu.
3. Type Seattle S3 in the **Search Results** field.

The requested filtered result is displayed.

Chapter 7. Maintenance

Overview

This information covers maintenance activities, such as upgrade, tools to support troubleshooting (particularly log collection), and reports. All reports can be exported to a .csv file, sent via email on demand, or sent automatically based on a configured schedule.

Note: Any custom configuration modifications that are made to the Accesser device that were not made by using the Manager Web Interface will need be redone after you upgrade.

Upgrade

Upgrading is performed from the **Maintenance** tab. The upgrade process supports $n-2$ upgrades, which means the starting point of a direct upgrade is a release version whose second digit differs by less than or equal to two. Contact IBM Customer Support for further information.

Upgrading is based on the concept of a single upgrade queue, containing the devices to be upgraded. A device selected for upgrade is placed in a queue and upgraded, while ensuring the health of the vaults across the system are not compromised. Before a device upgrade, the Manager application tests to see whether the health of any of the vaults that are associated with the device are negatively impacted if the device goes down. Upgrade can be performed on a single device, a handful of devices, or an entire access/storage pool. Checks will be performed to ensure that a device comes back online correctly after upgrade.

Note: When upgrading to a new version of ClevOS, some incidents are automatically closed and reopened if they still exist after upgrade.

Note: The Protection Level cannot be set to **Retention** until all devices that will host the retention vault have been upgraded to ClevOS 3.12.0 or newer.

In a typical device upgrade, the state transitions as follows:

Table 25. Typical device upgrade state transitions

State transitions	Detail
Requiring	The device software version differs from the Manager device.
Enqueuing	The command to upgrade this device is sent, but a status update has not been performed. Status updates occur every 10 seconds.
Precheck	The device is performing a pre-upgrade consistency check to ensure that the system is capable of being upgraded.
Precheck Failed	When a precheck fails, this status appears and remains until the device is upgraded again. Reloading the page removes this status.
Pending	The device is added to the upgrade queue.
Initiated	The device has begun the upgrade process.
Downloading	The upgrade package is copied to the device.
Stopping	The command to stop the device process has been issued.

Table 25. Typical device upgrade state transitions (continued)

State transitions	Detail
Capturing Integrity	The device is capturing integrity state that will be used after upgrading the device to ensure the upgrade completed successfully.
Rebooting	The device is being rebooted with a new software version.
Verifying Integrity	The device is verifying that it upgraded successfully.
Success	All data integrity checks passed and the correct version of the software is confirmed.
Failure	The device did not pass the data integrity checks and the correct version of the software is not confirmed.
Current	The device has the most recent software version.

Prechecks are performed before adding devices into the upgrade queue. During an upgrade, several scenarios can occasionally occur that require operator intervention. As part of an upgrade, various options are presented to the user if unexpected conditions arise during a device upgrade. In particular, the following scenarios might arise:

- Prechecks can fail.
- Device upgrade takes longer than expected.
- Device processes take longer than expected to stop.
- General issues arise during a device upgrade.

If problematic devices exist within a pool, a message appears. An exclamation point (icon) appears on the left side of the page. It indicates that user intervention is needed on that pool.

If the device does not upgrade successfully, an icon appears with an alert message on the same page; manual intervention is needed. The device can be removed from the upgrade queue and the upgrade retried. Remove applies to a single device in the Pending or Failure state. A **Remove** button exists to the right of the device (right portion of the upgrade page), which when selected removes the device from the upgrade queue. After a device is removed, the status will be Requiring. Upgrade is attempted only if vault health will not be compromised. Problems encountered when software updates fail are identified through messages.

Upgrade detects circumstances in which device processes cannot be stopped in a timely manner. When detected, an icon appears, and a **Force Kill** button is displayed. If selected, upgrade forces termination of the process.

Upgrade also detects scenarios where restart, download, or integrity verification take longer than expected. If this occurs, the device can be removed from the upgrade queue by using the **Force Removal** button. Integrity failures result in an incident appearing in the **Open Incidents** view.

Note:

If IE is being used to upgrade, set the following option:

1. Go to **Internet Options**. It is available from Tools (IE8) or a circular settings icon (IE9).
2. In the **General** tab, select **Settings** | **Browsing history**.
3. In the pop-up window for the **Check for newer versions of stored pages** section, select the **Every time I visit the webpage** option.

The high-level procedure consists of the following steps:

1. Transfer the upgrade compressed file to the machine running the browser. (See alternative procedure in the note below.)
2. Browse and upload one upgrade compressed file from the **Upgrade** page. (See alternative procedure in the note below.)
3. Initiate the Manager upgrade from the **Upgrade** page.

Note: Perform a backup of the Manager database before continuing. At the end of the upgrade, log back in to the Manager.

4. Initiate the upgrade of the remaining devices from the **Upgrade** page.

In some circumstances, a desirable alternative to Steps 1 and 2 is to transfer the image file directly to the Manager device from an alternative source machine. It is beneficial when the operator is working over a low-bandwidth connection (for example, wifi or trans-continental) for which the upload time of the upgrade image is prohibitive. In such cases, the operator can log in to a jump host elsewhere on the network that already contains the upgrade image and send it directly to the Manager device. The following curl command accomplishes this:

```
$ curl --insecure -u <ADMIN_USERNAME> -F upgradeFile=@<UPGRADE_IMAGE> -F action=installUpgradeFile https://<MANAGER_IP>/manager/upgrade.adm
```

Table 26. Description of Options

Option	Description
--insecure	Manager device certificate does not need to verify as the Upgrade file is not secret information.
-u <ADMIN_USERNAME>	Administrative user name
-F upgradeFile=@<UPGRADE_IMAGE>	Upgrade image to send to the manager.
-F action=installUpgradeFile	Manager device extracts the provided file to disk.
(https://<MANAGER_IP>/manager/upgrade.adm)	URL of the upgrade page.

When run, the command prompts for the password that is associated with the specified user.

To determine the current version of the system, go to the **Maintenance** tab in the Manager Web Interface.

Before clicking **Upgrade**, the upgrade compressed file that is provided with the release needs to be copied to the machine running the browser.

Click **Upgrade** from the **Maintenance** tab to show the **Upgrade System Software** page.

The **Upgrade System Software** page shows three steps:

1. Upload One Manager Upgrade File

Browse for the upgrade file and click **Upload**.

Note: Only one upgrade file can be uploaded to the Manager at a time. If another file is uploaded during an upgrade, an error message appears until the page is reloaded.

Note: The behavior of the upload progress is different based on the browser:

- Firefox and Internet Explorer display an upload progress bar in the Manager Web Interface.
- Chrome and Safari do not show a progress bar. The browser window displays progress.

2. Upgrade Manager

Click **Start Upgrade**.

Attention: Do not restart the manager while upgrade is in progress. Manager upgrade takes approximately 15 minutes. If the upgrade takes longer, contact IBM Customer Support.

After the Manager upgrade completes, a dialog box appears. Click **Log back into the manager** to return to the log in page.

Note: When you upgrade the Manager to ClevOS 3.10.1 or later, it might not be possible to log in immediately since an extra 20 - 30 minutes might be needed for the Manager application to become available. On systems with large databases, the time might be longer. Contact Customer Support if it takes longer than 30 minutes to successfully log in to the Manager.

3. Upgrade Devices

Accesser devices are grouped by Access Pool. Slicestor devices are grouped by Storage Pool, and devices that are not associated with any particular pool are placed in their own group "Devices not in a storage pool or access pool." At the left of the page, for a storage pool in which all devices are to be upgraded, the following options are available: "Upgrade Entire Storage Pool," "Upgrade Entire Set," or "Upgrade" a single device. Similarly, for an access pool, the following options are available: "Upgrade Entire Access Pool" or "Upgrade" a single device. For "Devices not in a storage pool or access pool," the following options are available: "Upgrade All Devices Not in a Pool" or "Upgrade" a single device. Upgrade is possible when the status is Requiring, Pre-check Failed, or Failure.

Note: Vault access (unavailable) events will be generated after an Accesser device upgrade completes if the core process is down, particularly in single Accesser device configurations.

The full upgrade procedure is as follows:

1. Browse and upload the upgrade file as part of Step 1 on the **Upgrade** page.

Note: Only one upgrade file can be uploaded to the Manager at a time. If another file is uploaded during an upgrade, an error message appears until the page is reloaded.

2. If a backup is configured for this Manager, go to Step 8.
3. Click **Start Upgrade** in Step 2 under the **Upgrade Manager** section.
4. If the Upgrade has new End User License Agreements (EULA), then **Accept License Agreement to Continue** is made available for the user under **Step 3 - Upgrade Devices**, which on click, displays an IBM standard and non-IBM End User License Agreements (EULA) in a pop-up window similar to the End User License Agreement figure shown. The user needs to complete the Print Name (License Acceptor) field and check the appropriate box and click **Accept IBM & non-IBM Licenses** to accept the EULA, otherwise, device upgrades are not allowed.
5. If no backup configuration is created for this Manager, you are prompted to configure the backup, or upgrade without performing a backup.

Attention: IBM strongly recommends that the Manager is backed up before upgrade to ensure preservation of Manager data.

Click **Configure backup**.

6. On the **Backup Configuration** screen, specify encryption and FTP values as prompted and click **Configure** in the **Backup Configuration** action bar.

It is also possible to perform a manual backup from the **Administration** tab, but when you return to the upgrade step you are prompted again to configure the backup or upgrade without performing a backup. Then, you can select to upgrade without performing backup.

7. Refer to "Set backup configuration parameters" on page 62 for how to configure the backup parameters. On completion of the Backup Configuration, return to the **Upgrade** page by clicking the **Maintenance** tab and then clicking **Upgrade**.
8. Click **Start Upgrade** in Step 2 of the **Upgrade** page (Upgrade Manager).
9. A prompt offers the option to back up the Manager or cancel the operation. Click **Backup Manager Now**.
10. On completion of the backup, you are prompted to proceed with the upgrade or cancel the operation. Click **Proceed with Upgrade**.

11. The Manager will now be upgraded. Status updates are presented in the dialog box. Once the upgrade is verified, the button allowing you to log back in to the Manager Web Interface is enabled. Click **Log back into the Manager Web Interface**.

Note: When you upgrade the Manager to ClevOS 3.10.1 or later, it might not be possible to log in immediately since an extra 20 - 30 minutes might be needed for the Manager application to become available. On systems with large databases, the time might be longer. Contact Customer Support if it takes longer than 30 minutes to successfully log in to the Manager.

12. The Manager home page displays a banner at the top of the page that indicates there are devices in the system running a different version of software and provides a link.
Click this link to proceed to the **Upgrade** page.
13. The Manager upgrade portion of the UI shows the upgraded version of the Manager and report that no upgrade is available.
14. **Upgrade Devices** in Step 3.

Table 27. Supported Operations for Upgrade

Operation	Action
Remove	This applies only to a single device in the Pending or Failure state. Click the "Remove" button to the right of the device (right portion of the upgrade page) to remove the device from the upgrade queue.
Remove Entire Set	The "Remove Entire Set" button is enabled if there are 1 or more devices in the storage pool set that are in the Pending or Failure state.
Remove Entire Storage Pool	The "Remove Entire Storage Pool" button is enabled if there are 1 or more devices in the storage pool that are in the Pending or Failure state.

Constraints to upgrades

Some constraints exist regarding device upgrades.

- Only one Accesser device for the same vault can be upgraded at a time.
- During an upgrade, a threshold number of Slicestor devices must be available to ensure data availability and to protect data reliability.

The most conservative upgrade approach (and default) is to upgrade a single device per storage pool at a time. For faster upgrades, multiple devices can be upgraded simultaneously by adjusting the Upgrade Throttle parameter from 1 [default] to n . The number of available Slicestor devices equals to $[n = \text{Width} - (\text{Alert Level} + 1)]$.

If no alert levels are set for any of the vaults in the system, an upgrade throttle of one is used. Intuitively, one corresponds to a "conservative" upgrade, and "n" corresponds to an "aggressive" upgrade. It is possible that the upgrade throttle that is selected cannot be achievable due to the vault configurations across the storage pools. In all circumstances, vault health is never compromised.

- When upgrading SMC devices, you must have three devices in the SMC pool to upgrade those devices.
- When upgrading devices in a storage pool that uses Concentrated Dispersal vaults, the upgrade process waits 72 hours between initiating each device upgrade.

This is done to ensure data integrity. Once the 72-hour window has elapsed, another device can begin upgrade. This will also apply to storage pools operating in a mirrored setup where if a device on one side of the mirror has upgraded, all other devices on both sides of the mirror must wait 72 hours.

In the **System Properties** configuration options, the **Automatic Shutdown Timeout** parameter can be used to specify a maximum time devices wait for outstanding transactions to complete when attempting

to cleanly shut down. This parameter is not specific to the upgrade process and applies to all stops of the `dsnet-core` service but it is of particular importance in the context of system upgrades.

By default, this parameter is set to `None`, which means that devices wait indefinitely for transactions to close before shutting down cleanly. It prevents the possibility of inducing write errors as a result of the upgrade process and can lead to unbounded upgrade times if long-standing I/O operations are being performed on the system. The **Automatic Shutdown Timeout** can be used to place an upper bound on the amount of time the `dsnet-core` process waits for the transaction to close before forcefully ends outstanding transactions and completing the shutdown process.

The following options are available:

- `None`
- `15 min`
- `30 min`
- `1 hour`
- `2 hours`
- `4 hours`
- `8 hours`
- `16 hours`

Selecting a value other than `None` prompts a warning message that indicates active writes are dropped and in-flight data that is not yet committed to the system might be lost.

When a device begins the graceful shutdown of the `dsnet-core` process in preparation for upgrade, it waits up to the selected timeout for the shutdown to complete. If shutdown completes before the timeout, the device upgrade proceeds normally. If not, after timeout expiration, automatic termination will be performed on the `dsnet-core` process, and any active writes will be canceled.

On the right side of the Manager Web Interface under the upgrade queue, for the device being upgraded, a countdown indicates the amount of time that is left before automatic termination. This operation can be changed or stopped by setting the Automatic Shutdown Timeout to a different value or `None`. If a new value is specified (something other than `None`), any previous time that is spent in the stopping state is accounted for.

Migrating devices to IPv6

Migrate your devices to IPv6 to guard against the eventual depletion of IPv4 addresses and ensure that no outage in continuous connectivity occurs.

About this task

To use IPv6 in the system, all devices must be upgraded to ClevOS 3.10.1 or later. Devices that are brought in to an IPv6-enabled system that are not capable of IPv6 can function only if Slicestor[®] Devices within the same Vault do not have IPv6 addresses configured. These devices must be set to the same version as the system and might not provide all functions until this requirement is fulfilled.

Note: Native File Interface and geo-dispersed efficiency upgrades do not yet support IPv6.

When all devices are upgraded to an IPv6-enabled load, the system can run several network implementations:

- Single-stack IPv4 (The Manager appliance and all devices are configured with IPv4 addresses only).
- Dual-stack IPv4/IPv6 (The Manager appliance and all devices are configured with both IPv4 and IPv6 addresses).

Tip: You can set IPv6 as the preferred IP protocol for object operations within the system on the **Configure > Configure System Properties** page. For more information, see “Configuring system properties” on page 77.

- Single-stack IPv6 (The Manager appliance and all devices are configured with IPv6 addresses). IPv4 addresses are removed from the Manager and all devices.
- Mixed IPv4/IPv6 (The Manager appliance is configured with both IPv4 and IPv6 addresses). Some storage pools have devices that are configured with IPv6 addresses, while other storage pools have devices that are configured with IPv4 addresses only.

Use the following procedure to upgrade your system to a single-stack IPv6 implementation.

Procedure

1. Upgrade the Manager device to ClevOs 3.10.1 or later. For more information about upgrading the Manager, see “Upgrade” on page 111.
2. Log in to the Manager device and update the IP address, network mask, and network gateway to use IPv6 addresses. For more information on the **nut** commands, see the “Channel utility” section of the *Appliance Configuration Guide*.
3. Upgrade each device to ClevOs 3.10.1 or later.

Note: The system does not use IPv6 addresses until all devices are upgraded.

4. Log in to each device and update the IP address, network mask, network gateway, and Manager IP to use IPv6 addresses. For more information on the **nut** commands, see the “Channel utility” and “Manager utility” sections of the *Appliance Configuration Guide*.

Note: Ensure that you fully upgraded and updated each device before you move on to the next step.

5. Log in to each device and remove any IPv4 addresses by specifying the IPv4 IP, network mask, and network gateway addresses with no value.

IBM Cloud Object Storage Insight™

From the **Maintenance** tab, a user can read current IBM Cloud Object Storage Insight™ information and determine how the Manager device executes sessions with the IBM Cloud Object Storage Insight™ server. It is also possible to manually connect to the IBM Cloud Object Storage Insight™ server to send system data.

Leaving this feature enabled assists IBM support staff in resolving cases in a timely fashion. It can also help find potential issues with drives, configuration, or software.

Currently, only a small amount of information is sent to IBM support. It includes some disk-related metrics, metadata files, and the output of the following REST API calls:

- View System
- Disk Drive and Device Report
- Firmware Report
- Version History Report

Authentication details are redacted and are not sent to or visible to IBM employees:

- Passwords (accounts, backup, LDAP, and so on)
- Vault secret access keys
- Account secret access keys
- Private keys that are used for certificate generation.

When the anonymization settings are enabled, the best level of effort is made to redact or anonymized potentially sensitive data. Most configuration elements that indicate whether certain features are enabled or disabled are never redacted. For other types of information, the following tables illustrate how it's processed before it is sent to IBM.

Table 28. Content never anonymized or redacted in IBM Cloud Object Storage Insight™

Entity	Fields
Devices	Advanced configuration, software, and firmware versions for devices and drives, hardware information (model, serial number, and so on) for devices and drives, health-related information for devices and drives
Sites	Health
Cabinets	Health, how the slots are allocated (no device aliases or host names are present).
Vaults	How access permissions are used (Note: Account names are not present), IDAs, health, vault proxy types.
Mirrors	Synchronous operations
Storage Pools	Advanced configuration, storage engine in use, protocol type if the Accesser service is embedded.
Access Pools	Advanced configuration, http ports in use, how vaults and mirrors are deployed (does not include any names, descriptions, and so on).
Tags	How tags are allocated across vaults, devices, and so on.
Accounts	Roles and authentication type selected.
Groups	Authentication type that is selected.
Administration Settings	External agent deployments (no device aliases, host names, or site names are present), how the following settings are enabled: http authentication, back up alert forwarding, log collection (metadata information, such as credentials, are redacted).

Table 29. Examples of anonymizable or redactable fields in IBM Cloud Object Storage Insight™

Entity	Anonymizable Fields	Redacted Fields
Devices	IP addresses, host names	aliases, descriptions, certificate data
Sites	names	abbreviations, all contact information
Cabinets	names	descriptions
Vaults	names, access control	descriptions
Mirrors	names	descriptions
Storage Pools	names	descriptions
Access Pools	names	descriptions, subject alternative names
Tags		names, descriptions
Accounts	usernames	names, email addresses, any data related to configured authentication mechanism
Groups	names	aliases, any data related to configured authentication mechanism
Organizations	names, domains	descriptions, contact information

Table 29. Examples of anonymizable or redactable fields in IBM Cloud Object Storage Insight™ (continued)

Entity	Anonymizable Fields	Redacted Fields
Administration Settings	configured external certificate authorities,	Name, login banner, email alert rules, SSH keys configuration, SMTP server information and user name, back up server information and user name, LDAP bind+search model configuration settings, SMPV2 community string and allowed IP addresses, email addresses that receive automatic reports, Alert Forwarding IP addresses/hostnames and community string, Keystone configuration settings, External Agent names, and file names.

Configuring IBM Cloud Object Storage Insight™

About this task

Note: The hostname of the IBM Cloud Object Storage Insight server is `ph1.cleversafe.com`.

Procedure

- Click **Configure** on the **Cloud Object Storage Insight** action bar.
The last successfully performed IBM Cloud Object Storage Insight™ session is displayed.
 - View under the **Automatic Settings** bar.
 - Enabled/Disabled status
 - Automatic schedule
 - Next automatic session time
 - View under the **Proxy Settings** bar.
 - Enabled/Disabled status
 - The configured proxy host name and port
 - Whether all requests that are routed through the proxy require basic HTTP authentication and the configured credentials.
 - View under the **Anonymization Settings** bar.
 - Whether data is anonymized before it is sent to IBM.
- Click **Configure** on the **IBM Cloud Object Storage Insight Settings** action bar to configure these settings.
- Enable or disable automatic operation. Automatic operation is enabled by default. When automatic operation is enabled, the service runs periodically in the background based on a schedule set by the IBM server.

Note: IBM determines the IBM Cloud Object Storage Insight™ automatic schedule. It cannot be modified.

- Enable or disable proxy settings.
 - Enable or disable **Use an HTTP Proxy**. Enter the **Proxy Hostname** or endpoint and the **Proxy Port**.
 - If authentication is also required, enable **Use Authentication**. Enter the **Proxy Username** and associated **Proxy Password**.
- Enable or disable **Data Anonymization**. When enabled, potentially sensitive information is anonymized or redacted before it is sent to IBM Customer Support.

Manually starting IBM Cloud Object Storage Insight™ sessions

Procedure

1. Click **Configure** on the **IBM Object Cloud Storage Insight** action bar.
2. Click **Execute Session**.

CAUTION:

Execute Session cannot be pressed again for 1 hour after the session completes.

Viewing an anonymized object

Procedure

1. Click **Configure** on the **IBM Cloud Object Storage Insight** action bar.
2. Click **View** on the **View Anonymized IBM Cloud Object Storage Insight Data** action bar.
3. Enter the token that is provided by IBM support under the **View Anonymized IBM Cloud Object Storage Insight Data** action bar.

Objects that match the token are displayed. Click a result to view to the appropriate **Monitor** page.

Note: The full token does not need to be entered. The Manager device returns a full list of objects that match the partial token provided.

Logs

From the **Maintenance** tab, you can collect internal logs across several different devices and send them to an SFTP or HTTP server, or choose not to send collected logs. Log collection can be performed manually on demand or automatically based on a schedule. Log collection status can be viewed, which displays "real-time" status.

Status that is presented on the Manager Web Interface includes the following items.

- Pending (default state until log collection starts).
- Success (log collection completed successfully).
- Error states (for example, device log collection failures, device communication issues, and so on).

If you choose not to send collected logs to an SFTP or HTTP server using the Manager Web Interface or Manager REST API, but have Privacy Controls enabled, then dump-logs initiated from any device are still redacted. For more information on Privacy Controls, see "Redacting client IP addresses" on page 123.

Note: Automatic log collection currently does not allow configuration of time and selection of devices. If these capabilities are needed, manual log collection should be used.

Case numbers, added to logs in ClevOS 3.4.0, specified in the Manager Web Interface, apply to devices that are running ClevOS 3.4.x or later. They serve as the directory into which the logs are placed on the destination server.

CAUTION:

Only one log collection activity should be performed at a time. Make sure that log collection completes before you start another one. The status of the most recent log collection activity can be viewed by going to the Maintenance tab on the Manager, selecting **Configure** in the Logs section, and selecting **View** in the Log Collection Status section of the Logs page.

Collecting logs manually

Selected logs can be manually collected and uploaded to an HTTP Server or an SFTP server.

Procedure

1. Click the **Maintenance** tab.
2. Click **Configure** on the **Logs** action bar.
3. Click **Configure** on the **Log Collection Configuration** action bar.
4. Provide the following information, depending upon the backup server type.

Backup Server	Actions
SFTP	<p>Provide valid SFTP server connection details:</p> <ul style="list-style-type: none"> • host name • path • user name • password <p>All of these fields must be entered.</p> <p>The manager validates the SFTP credentials before the selected devices start collecting logs.</p> <p>Note: It is not possible to use ~ for the users directory. It must be written out.</p>
HTTP	<p>Provide valid HTTP server connection details:</p> <p>Required</p> <ul style="list-style-type: none"> • Host name (for example, 192.168.14.63 or fe00:d7::1 or logs.company.com) Note: Do not include http:// or https://. • Path (for example, /s3/vault) <p>Optional</p> <ul style="list-style-type: none"> • Check the SSL/TLS check box to use HTTPS. • Check the Use Custom Port check box (then type in the port number) to set a different port than the defaults. <ul style="list-style-type: none"> – The default HTTP port is 80. – The default HTTPS port is 443. • Check the Use authentication check box if necessary. <ul style="list-style-type: none"> – If it is checked, type values into the Username and Password fields. • Check Use an HTTP Proxy check box. <ul style="list-style-type: none"> – If checked, select either the Automatically configure the Manager as an HTTP Proxy or Configure an HTTP proxy option: <ul style="list-style-type: none"> - Automatically configure the Manager as an HTTP Proxy <ul style="list-style-type: none"> • When this option is selected, the Manager device configures itself to act as an HTTP proxy, open to each of the devices. Logs sent to the destination are proxied through the Manager device. - If Configure an HTTP proxy is selected, enter values in the following fields: <ul style="list-style-type: none"> • Proxy Hostname • Proxy Port • Check the Use authentication check box if necessary. <ul style="list-style-type: none"> – If it is checked, type values into the Proxy Username and Proxy Password fields.

5. Click **Update**.
6. Click **Collect** from the **Collect Logs Manually** action bar and provide the information as specified.

Criteria	Actions
Filter Criteria	<p>Indicate the time period (From/To). By default, it is one week. Smaller time periods limit the log file size that is transferred from each device to the SFTP or HTTP server. Categories specify the type of log files to collect and further limit the log file size. Contact IBM Customer Support for log file details.</p> <p>Note: Select the wanted logs to minimize excess network traffic. Log files can be large. Selecting a smaller time range limits the file size and speedup the operation. The Cloud Object Storage Manager queues the requests and process five devices at a time.</p>
Devices	<p>Select devices. For any one request, logs can be collected on the Manager device, all Accesser devices, and all SliceStor devices on any one storage pool. Only one storage pool can be selected at a time. The site is shown next to the device name.</p>

7. Enter a **Case Number** (optional) to be used as a directory for the collected log files (spaces are trimmed).
8. Click **Collect Logs** to start log collection or **Cancel** to exit. When log collection is in process on the manager, an information box shows up at the top of the page, which allows a log collection to be canceled. Click the **can be terminated** link to start a cancel log collection request. It prevents any pending log collection runs from running, sends a cancel request to the log collection process on each device that is in progress, and closes the log collection execution thread in the manager. When the log collection process on each device receives the cancel request, it makes a best-effort attempt to prevent the log collection from sending the logs to the destination.
9. Click **View** to view log collection status.

Collecting logs automatically

System logs can be collected automatically and uploaded to an HTTP Server or an SFTP server.

Procedure

1. Click the **Maintenance** tab.
2. Click **Configure** from the **Logs** action bar.
3. Click **Configure** on the **Log Collection Configuration** action bar.
4. Provide the following information, depending upon the backup server type.

Backup Server	Actions
SFTP	<p>Provide valid SFTP server connection details:</p> <ul style="list-style-type: none"> • Host name • Path • User name • Password <p>All of these fields must be entered.</p> <p>The manager validates the SFTP credentials before the selected devices start collecting logs.</p>

Backup Server	Actions
HTTP	<p>Provide valid HTTP server connection details:</p> <p>Required</p> <ul style="list-style-type: none"> Host name (for example, 192.168.14.63 or fe00:d7::1 or logs.company.com). Note: Do not include http:// or https://. Path (for example, /s3/vault) <p>Optional</p> <ul style="list-style-type: none"> Check the SSL/TLS check box to use HTTPS. Check the Use Custom Port check box (then type in the port number) to set a different port than the defaults. <ul style="list-style-type: none"> The default HTTP port is 80. The default HTTPS port is 443. Check the Use authentication check box if necessary. <ul style="list-style-type: none"> If it is checked, type values into the Username and Password fields. Check HTTP Proxy check box. <ul style="list-style-type: none"> If checked, enter values in the following fields: <ul style="list-style-type: none"> Proxy Hostname Proxy Port Check the Use authentication check box if necessary. <ul style="list-style-type: none"> If it is checked, type values into the Proxy Username and Proxy Password fields.

5. Enable **Automatic Log Collection** and select the collection period.

Note: Select the wanted logs to minimize excess network traffic. Log files can be large. Selecting a smaller time range limits the file size and speeds up the operation.

Redacting client IP addresses

A system administrator can enable options to perform log redaction (remove client IP addresses from dump-log output or management vaults), which can help meet regulatory restrictions on collecting personally identifiable information.

Before you begin

Log redaction can be configured from two locations in the Manager Web Interface, which offer different behavior.

Removing client IP addresses during log collection

Log collection redaction applies to access logs and HTTP logs (including rotated and zipped logs) as well as netstat output. The original contents of the logs is not redacted, but the dump-log content is redacted before it is sent to an SFTP or HTTP server.

Procedure

1. Click the **Maintenance** tab.
2. Click **Configure** on the **Logs** action bar.
3. Click **Configure** on the **Log Collection Configuration** action bar.
4. Enable **Remove client IP addresses during log collection** in the **Privacy Controls** section.

5. Click **Update**.

Removing client IP addresses from management vaults

Management vault redaction only applies to rotated and zipped access log files that were or are placed as objects in a device's management vault. Existing objects are downloaded from the management vault, unzipped, redacted, zipped, and placed back in the management vault.

Procedure

1. Click the **Configure** tab.
2. Click **Configure** in the **Configure Management Vault** action bar.
3. In the Management Vault Options section, enable **Backup HTTP access logs** and **Remove client IP addresses**. A hyperlink to view the redaction progress displays.
4. Click **Update**.

Collecting log status

Status of last log collection. Only devices that were included in the last run appears.

About this task

Collection status updates periodically until no devices are in progress or pending. Devices are in a 'Pending' state until a log collection is initiated; at that point, each device is in an 'In Progress' state until it either succeeds or fails. A failed log collection is marked red with an error message dependent on the type of failure. There is one 'Success' state, marked green.

Procedure

1. Click the **Maintenance** tab.
2. Click **Configure** in the **Logs** action bar.
3. Click **View** on the **Log Collection Status** taskbar to view the **Log Collection Status** page.
4. Click **Export** in the **Log Collection Status** taskbar to export this data as a CSV file.

Note: Log collection status does not retain any historical information. Only the last known collection is retained.

Note: Failed states often contain troubleshooting information (text) regarding the cause of the failure.

Configuring device logs

You can set the Access Log Retention Period and Access Log Rotation Period.

About this task

Note: A system administrator can enable options to perform log redaction (remove client IP addresses from management vaults or dump-log output), which may help with meeting regulatory restrictions on collecting personally identifiable information. The Cloud Object Storage Manager has two locations where log redaction can be configured. The two locations configure different behavior. The first is on the Configure Management Vault page under the Management Vault Options section (Remove client IP addresses). For this case, Management Vault redaction applies to rotated and zipped access log files that exist or will be placed as objects in a device's management vault. Existing objects will be downloaded from the management vault, unzipped, redacted, zipped, and placed back in the management vault. The second place where log redaction can be configured is the Log Collection Configuration page under the Privacy Controls section (Remove client IP addresses during log collection). In this case, Log Collection redaction applies to access logs and HTTP logs (including rotated and zipped logs) as well as netstat output. The original contents of the logs will not be redacted; the dump-log content will be redacted before it is sent to an SFTP or HTTP server.

Procedure

1. Click the **Maintenance** tab.
2. Click **Configure** on the **Logs** action bar.
3. Click **Configure** on the **Device Log Configuration** action bar.
4. Select the desired **Retention Period** from the drop-down menu. The Retention Period is the maximum amount of time that access logs are persisted on Accesser devices and Slicestor devices.
5. Select the desired **Rotation Period** from the drop-down menu. The Rotation Period applies to both Vault Mode and Container Mode, controlling how often access logs are rotated. If "Not Set", the rotation only occurs once the access log reaches 500 MB. Otherwise, the logs rotate based on the selected time interval. The Management vault upload of the access log occurs within one (1) hour of the log rotation.
6. Click **Update**.

Troubleshooting console

The troubleshooting console allows a collection of commands to be run on one or more devices to help support troubleshooting.

1. Click the **Maintenance** tab.
2. Click **View** from the **Troubleshooting Console** action bar.
3. Select one or more devices from the **Devices** area.
 - The **Device Group** tab has several drop-down filters that you can use to filter the list of available devices. The resulting list shows the device type, health, and name. Select the check box next to each device you want to run commands on. Select the **Select All** check box to select all of the devices in the filtered list.
 - The **Devices** tab contains a list of all available devices that are organized by Device Type. Click in the text box to show the list. Select each device you want to run commands on. You can also use CTRL+click to select multiple devices. Click the "X" next to the device name to remove it from the list.
4. Commands are grouped into six categories. Select the check box next to each command in the **Select Commands** area that you want to run on the devices that you selected in the **Devices** area. Select the check box next to a category name to select all commands in that category.

The available categories and their respective commands include (commands are valid for all releases going forward from the one indicated):

Table 30. Available categories and commands.

Category	Command	Supported in Release			
		3.12.0	3.11.0	3.10.1.x	3.8.3.x
General	<code>cat /proc/cpuinfo</code>	Yes	Yes	Yes	Yes
	<code>ipmitool mc info</code>	Yes	Yes	Yes	Yes
	<code>cat /proc/meminfo</code>	Yes	Yes	Yes	Yes
	<code>ipmitool sel list</code>	Yes	Yes	Yes	Yes
	<code>uptime</code>	Yes	Yes	Yes	Yes
	<code>nut health statistic</code>	Yes	Yes	Yes	Yes
	<code>free</code>	Yes	Yes	Yes	Yes
	<code>nut health state</code>	Yes	Yes	Yes	Yes
	<code>cat /proc/cmdline</code>	Yes	Yes	Yes	Yes
	<code>appliance</code>	Yes	Yes	Yes	Yes
Process	<code>ps ax</code>	Yes	Yes	Yes	Yes
	<code>pstree -ap</code>	Yes	Yes	Yes	Yes
	<code>top -b -n1</code>	Yes	Yes	Yes	Yes
Storage	<code>mount</code>	Yes	Yes	Yes	Yes
	<code>nut storage list</code>	Yes	Yes	Yes	Yes
	<code>cat /proc/scsi/scsi</code>	Yes	Yes	Yes	Yes
	<code>storagectl list</code>	Yes	Yes	Yes	No
	<code>df -ah</code>	Yes	Yes	Yes	No
	<code>sg_map -x</code>	Yes	Yes	Yes	Yes
	<code>storagectl list all</code>	Yes	Yes	Yes	Yes
	<code>df -i</code>	Yes	Yes	Yes	Yes
	<code>nut enclosure bay list</code>	Yes	Yes	Yes	Yes
	<code>storagectl info</code>	Yes	Yes	Yes	No
	<code>sdf -h</code>	Yes	Yes	Yes	Yes
	<code>service dlm status --extend</code>	Yes	Yes	Yes	No
	<code>storagectl history</code>	Yes	Yes	Yes	No
	<code>sdf -i</code>	Yes	Yes	Yes	Yes
	<code>df -h</code>	No	No	No	Yes

Table 30. Available categories and commands. (continued)

Category	Command	Supported in Release			
		3.12.0	3.11.0	3.10.1.x	3.8.3.x
Network	<code>ifconfig</code>	Yes	Yes	Yes	Yes
	<code>ping -c3 -W5 -s9000 manager.dsnet</code>	Yes	Yes	Yes	No
	<code>ipmitool lan print</code>	Yes	Yes	Yes	Yes
	<code>netstat -anp</code>	No	No	No	Yes
	<code>ifconfig -a</code>	Yes	Yes	Yes	Yes
	<code>ping6 -c3 -W5 manager.dsnet</code>	Yes	Yes	Yes	No
	<code>arp -n</code>	Yes	Yes	Yes	Yes
	<code>netstat -rn</code>	No	No	No	Yes
	<code>netstat -anpW</code>	Yes	Yes	Yes	No
	<code>ping6 -c3 -W5 -s9000 manager.dsnet</code>	Yes	Yes	Yes	No
	<code>ip -6 neigh show</code>	Yes	Yes	Yes	No
	<code>netstat -rn -46</code>	Yes	Yes	Yes	No
	<code>nut system dns</code>	Yes	Yes	Yes	Yes
	<code>lsof -n -i :5000</code>	Yes	Yes	Yes	No
	<code>ping -c3 -W5 manager.dsnet</code>	Yes	Yes	Yes	Yes
	<code>host cleversafe.com</code>	Yes	Yes	Yes	Yes
	<code>ethtool_summary</code>	Yes	Yes	Yes	No
Clock	<code>data</code>	Yes	Yes	Yes	Yes
	<code>hwclock</code>	Yes	Yes	Yes	Yes
	<code>ntpq -pn ::1</code>	Yes	Yes	Yes	No
	<code>nut system ntpservers</code>	Yes	Yes	Yes	Yes
	<code>ntpq -pn</code>	No	No	No	Yes

Table 30. Available categories and commands. (continued)

Category	Command	Supported in Release			
		3.12.0	3.11.0	3.10.1.x	3.8.3.x
Other	cat /proc/slabinfo	Yes	Yes	Yes	Yes
	dmesg	Yes	Yes	Yes	Yes
	cat /proc/vmstat	Yes	Yes	Yes	Yes
	dsnet-storage-alloc	Yes	Yes	Yes	Yes
	lspci	Yes	Yes	Yes	Yes
	lsmod	Yes	Yes	Yes	Yes
	dmidecode	Yes	Yes	Yes	Yes
Other (Gateway device specific)	nodetool status	Yes	Yes	Yes	Yes
	nodetool cfstats	Yes	Yes	Yes	Yes
	nodetool info	Yes	Yes	Yes	Yes
	nodetool compactionstats	Yes	Yes	Yes	Yes
	nodetool gcstats	Yes	Yes	Yes	Yes
	nodetool tpstats	Yes	Yes	Yes	Yes
	crm status	Yes	Yes	Yes	Yes
	crm configure show	Yes	Yes	Yes	Yes
	crm resource show	Yes	Yes	Yes	Yes

5. Click **Run Selected Commands** to run the commands.

Changing the device local password

A device-wide local admin password can be changed if wanted.

Procedure

1. Go to the **Maintenance** tab.
2. Click **View** in the **Change Device Local Password** action bar.
3. Select the devices that require password modification on the right side.
4. Enter the current password of the selected devices in the **Current Password** field. Devices that do not match this password will not have their password modified.
5. Enter the new password for the selected devices.
6. Reenter the new password for the selected devices.
7. Click **Save**.

Automatic report emailing

Configure automated system reports.

Under the **Reports** section of the **Maintenance** tab, click **View** to access Automatic Report Emailing. The following reports can be selected for automatic emailing:

- Disk Drive and Device Report
- Storage Pool Capacity and Disk Report
- System Usage and Configuration Summary Report
- Storage Pool Usage Report
- Vault Usage Report
- Vault Summary Report
- Device Summary Report
- Failed FRU Report
- Event Report
- Firmware Report
- Redaction Status Report (only visible when **Configure management vault > Remove client IP addresses** is enabled)

Email addresses of individuals to receive the reports can be specified. Automatic report emailing can be scheduled for daily or weekly intervals. The times are in GMT.

Enter the following information and click **Update** or **Cancel**.

Reports

Select the reports, which are in HTML format with a .csv attachment.

Email Addresses

Enter the email addresses, which are separated by a comma or return.

Schedule

Enable **Automatic** reports, and select the report frequency.

Reporting

Reports can be emailed automatically or generated manually on demand.

Under the **Reports** section of the **Maintenance** tab, click **View** to access the reports.

Automatic Email

Configure the automatic report feature. The Administration SMTP Configuration must be configured to use this feature.

Disk Drive and Devices

The **Disk Drive and Device** utility can be used to view disk parameters (model, serial number, software version) for each disk in the system. [All parameters are not available for virtual devices.] Click **Generate** to begin.

Storage Pool Capacity and Disk

The **Storage Pool Capacity and Disk** utility can be used to view storage pool capacity. Click **Generate** to begin.

System Usage and Configuration Summary

The **System Usage And Configuration Summary** can be used to generate a summary of organizations, storage usage, sites, storage pools, and license key information for this system. Click **Generate** to begin.

Vault Usage

Generate a report that shows vault usage over time. Click **Generate** to begin.

Vault Summary

Generate a report of all Vault configurations in the system. Click **Generate** to begin.

Device Summary

Generate a report of all Device configurations in the system. Click **Generate** to begin.

Failed FRUs

Generate a report of failed Field Replaceable Units (FRUs). Click **Generate** to begin.

Events

Generate a report of the Events. Click **Generate** to begin.

Firmware

Generate a report of device Firmware. Click **Generate** to begin.

“Redaction status report” on page 136

If Cloud Mode is enabled, generate a report of redaction status. Click **Generate** to begin.

Generating disk drive and device reports

Create a summary of disk information for all drives in the system.

About this task

The filter functions can be used to sort the data by relevant parameters, including Drive Status, Device Alias, Site, Cabinet, Device Model, Device Software Version, SCSI Name, Drive Model, and Firmware. More columns appear if a multi-node appliance exists in the system. The columns are Chassis ID and Node Location and Node Count that describes the node configuration of a multi-node chassis. For appliances that is not multi-node, the columns are marked 'N/A'. This information is view-only and cannot be changed from this screen.

The contents can be filtered, based on items such as disk status and model, and exported to a .csv file. It can be used for troubleshooting disk and device health issues. The following information is available:

- Drive Status
- Suspend Reason
- SCSI Name
- RAID Sequence
- Bay
- Drive Serial number
- Drive Model

- Drive Firmware
- Alias
- IP Addresses
- Device Serial number
- Model
- Version
- Cabinet (if present)
- Slot number (if cabinet present)
- Site
- Pool
- Chassis ID
- Node Location
- Node Count
- Drive Capacity

Note: This report displays **Bay** identifiers with a new format: <chassisId>:<enclosureId>:<slotId>.

Procedure

1. Click **View** under the **Reports** section in the **Maintenance** tab to generate this report.
2. Click **Generate** from the **Disk Drive and Device Report** action bar displays a list of all drives for all devices within the system.
3. Click **Send** to email or **Export** to create a .csv (comma-separated values) file for use with spreadsheet applications.

Note: A reduced set of drive attributes are available for virtual devices. Bay number, drive serial number, drive model, and drive firmware appear as N/A.

Generating storage pool capacity and disk report

Generate a report of system Storage Pool capacity and drive health.

About this task

The report displays a list of all storage pools with counts of different drive states and all the unhealthy drive information in a different table within the system.

The report contents can be searched by using the **Search** bar provided on top of each table. Storage pools can be sorted based on the number of drives at a particular state. The entire report can be exported to a .csv file.

Note: This report displays **Bay** identifiers with a new format: <chassisId>:<enclosureId>:<slotId>.

Procedure

1. Click **View** under the **Reports** section in the **Maintenance** tab to generate this report.
2. Click **Generate** from the **Storage Pool Capacity and Disk Report** action bar.
3. Select **Storage Pool Capacity and Drive Health** or **Unhealthy Drives by Storage Pool**.
4. Click **Send** to email or **Export** to create a .csv (comma-separated values) report for use with spreadsheet applications.

Report contents

In the first table of the report, the contents can be filtered based on items such as storage pool name and width. It can show general storage pool health.

The first table includes:

- Storage Pool ID
- Storage Pool Name
- Storage Pool Width
- Storage Pool Total Size/ Capacity
- Storage Pool Usage
- Storage Pool Remaining free space
- Percentage of Remaining Free space
- Number of ONLINE drives
- Number of MIGRATING drives
- Number of FAILED drives
- Number of DIAGNOSTIC drives
- Number of FOREIGN drives
- Number of OFFLINE drives
- Number of UNUSABLE drives
- Number of INIT drives
- Number of UNKNOWN drives
- Total Number of unhealthy drives
- Total Number of drives

In the second table of the report, the contents can be filtered based on items such as storage pool name, host name, device model, drive status, suspend reason, drive model, and firmware.

It can be used for troubleshooting all the disk and device health issues in a storage pool.

The second table includes:

- Storage Pool Group Name
- Host Name
- Device Model
- Drive Status
- Suspend Reason
- Bay number
- Drive Capacity
- Drive Serial Number
- SCSI Name
- Model
- Firmware

Generating a vault usage report

The **Vault Usage Report** generates a vault usage summary based on a date range (monthly or user-specified) for each vault.

About this task

Daily averages of the storage that is used by each vault are calculated. When a date range is selected, the sum of daily average usage in the range is displayed in PB/TB/GB/MB/kB days. When the content is

exported, the number is converted to Byte-days that can be used for billing purposes. The report includes all the vaults (standard, service, container, and management) on the system irrespective of the vault purpose.

Storage that is used by each vault is shown as the sum of daily average usage for the days that are selected in tabular format. Each vault usage is displayed in "PB/TB/GB/MB/kB days". When the content is exported, the usage is converted to Byte-days.

The **Search** field can be used to dynamically search through the table for items that contain the search input. The **Filter** drop-down on the **Storage Pool** column can be used for filtering vaults that are associated with the selected **Storage Pool**. Both of these operations are mutually exclusive.

The table can be sorted on columns **Name**, **Usage**, and **Provisioning Code** by clicking the associated column header.

Note: The **From** date is inclusive and the **To** date is exclusive. Usage is collected every day at midnight GMT. The **Usage Reporting Start** and the **Usage Reporting End** columns present data based on either the user selected time zone or the default time zone (GMT).

Procedure

1. Select either the **Month** from the drop-down or provide a **Date Range** by using the **From** and **To** fields for the report.
2. Click **Send** to email the report to an email address in .csv format.
3. Click **Export** to download the report in .csv format.
4. Click **Cancel** to get back to the previous page.

Vault summary report

The Vault Summary Report can export the following data for all vaults to a .csv file:

- Vault ID
- Name
- Description
- Creation Date
- Storage Pool
- Protection
- Vault Configuration
 - Threshold
 - Read Threshold
 - Write Threshold
 - Alert Level
- SecureSlice
- Segment Size
- Vault Usage
 - Raw Used
 - Usable Used
 - Usable Capacity
- Soft/Hard Quotas
- Versioning
- Notification Service
- Notification Service Configuration

- Notification Service Topic
- Container Storage Location
 - Storage Location ID
 - Provisioning Code
 - Region
 - Storage Class

Click **Send** to email or **Export** to create a .csv (comma-separated values) report for use with spreadsheet applications. The report can be filtered on some parameters. The default is ALL vaults.

Reporting device summary

Generate a report of all Device configurations in the system.

About this task

The **Device Summary Report** can export the following data for all devices to a .csv file:

- Alias
- Host name
- IP address
- Site
- Type
- Software Version
- Model
- Serial Number
- MAC0(eth0)
- MAC1(eth1)
- Extra MAC Addresses
- Pool
- Device ID

Procedure

1. Click **Send** to email the report.
2. Click **Export** to create a .csv (comma-separated values) report for use with spreadsheet applications. The report can be filtered by using parameters. The default is ALL devices.

Failed FRU report

The **Failed FRU Report** identifies all failed drives, power supplies, and fans. The report can be filtered on some parameters. Default is ALL failures.

Note: The report shows any drive that is in the 'Failed' state, which means the drive should be replaced.

The failing of a drive covers these scenarios:

- A previously quarantined drive that is now in the **FAILED** state.
- Reset for drives in the **UNUSABLE** state.
- If the reset is successful, the drive is put in the **ONLINE** state.
- If the reset fails, the drive is put in the **FAILED** state.
- If system detects a previously failed drive, it puts the drive in the **FAILED** state.

All the failed drives, failed fans, and failed PSUs appear for replacement in the FRU Report.

The **Failed FRU Report** can be exported to a .csv file with the following format:

- Failed Drives
- Device
- Model
- Firmware
- Serial
- Capacity
- Bay
- Sequence
- Device Model
- Version
- Device Serial #
- Site
- Failed Fans
- Device
- Model
- Fan Name
- Failed PSUs
- Device
- Model
- PSU Name

Note: This report displays **Bay** identifiers with a new format: <chassisId>:<enclosureId>:<slotId>.

Click **Send** to email or **Export** to create a .csv (comma-separated values) report for use with spreadsheet applications.

Reporting event information

Generate a report of system events.

About this task

An **Event Report** can send all system events via email or .csv file. The **Advanced Search** can be used to filter out events in the console before you send/export its contents. The link is in the upper right of the console.

Filtered Events are applicable when the report is emailed or exported. For example, if advanced search is being performed and only incidents are showing in the UI page and then export or email of the report is performed, the items in the export/email match what is seen on the UI. Automated emails for the **Event Report** contain all log messages, all incidents, and contain event information from the past 7 days.

Note: The exported or emailed content is limited to a maximum of 50,000 items.

Procedure

1. Click the **Maintenance** tab.
2. Click **View** on the **Reports** action bar.
3. Configure the report with **Advanced Search** filters.
4. Click **Send** to send the report by email.
5. Click **Export** to generate a .csv file for use in spreadsheet applications.

Method	Action
Send via email	Click Send in the action bar.
Export CSV	Click Export in the action bar.

Firmware report

A firmware report is a firmware summary of each device. It includes device BIOS/BMC, network, disk controller, and disk drive firmware.

Method	Action
Send via email	Click Send in the action bar.
Export CSV	Click Export in the action bar.

Redaction status report

This report outlines the redaction status for logs persisted to a management vault prior to access log redaction being enabled for each device in the system. The content can be exported to CSV format and sent by email.

Note: This report is only available if **Configure management vault > Remove client IP addresses** is enabled.

The report provides the following information:

- The **Access Pool** that the device is deployed to, if applicable.
- The **Storage Pool** that the device is deployed to, if applicable.
- The **Hostname** of the device.
- The **IP** of the device.
- The **Site** of the device.
- The redaction **Status** of the device. Status values are "Initializing," "Unknown," "Upgrade Required," "In Progress," and "Complete." The "Initializing" state occurs when log redaction is starting or it recently rebooted. The "Unknown" state occurs only when there is an issue determining the current redaction state. The "Upgrade Required" state occurs when the device first needs to be upgraded before redaction can start.
- If the device has been removed from the system, then **Is Deleted** displays "Yes".
- The number of **Logs Redacted** currently.
- The **Total Logs** that need or needed to be redacted.

Method	Action
Send by email	Click Send in the action bar.
Export CSV	Click Export in the action bar.

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Accesser[®], Cleversafe[®], ClevOS[™], Dispersed Storage[®], dsNet[®], IBM Cloud Object Storage Accesser[®], IBM Cloud Object Storage Dedicated[™], IBM Cloud Object Storage Insight[™], IBM Cloud Object Storage Manager[™], IBM Cloud Object Storage Slicestor[®], IBM Cloud Object Storage Standard[™], IBM Cloud Object Storage System[™], IBM Cloud Object Storage Vault[™], SecureSlice[™], and Slicestor[®] are trademarks or registered trademarks of Cleversafe, an IBM Company and/or International Business Machines Corp.

Other product and service names might be trademarks of IBM or other companies.

Homologation statement

This product may not be certified in your country for connection by any means whatsoever to interfaces of public telecommunications networks. Further certification may be required by law prior to making any such connection. Contact an IBM representative or reseller for any questions.



Printed in USA