# IBM Storage Virtualize for SAN Volume Controller and FlashSystem Family

*Concept Guide - Version 862*

**Note**

Before using this information and the product it supports, read the information in Chapter 1, “Notices,” on page 1.

This edition applies to version 8, release 6, modification x, and to all subsequent modifications until otherwise indicated in new editions.

# Contents

# Figures

# Tables

# Chapter 1. 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.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux® and the Linux logo is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other product and service names might be trademarks of IBM or other companies.

# Getting started

## System

A system consists of one or more I/O groups. Each I/O group contains two nodes. In enclosure-based systems, a control enclosure is an I/O group.

See the following pages to learn more about the system configuration:

### Nodes

A node is a single processing unit within a system. For redundancy, nodes are deployed in pairs.

Each pair of nodes is known as an I/O group. Each node can be in only one I/O group.

In enclosure-based systems, a node canister is equivalent to a node. The two node canisters in a control enclosure make an I/O group. Within a control enclosure, nodes communicate with each other via the enclosure hardware. See Enclosure for more details.

Nodes use Fibre-Channel or Ethernet ports to communicate with hosts, external storage, other nodes (outside the enclosure), or other systems. Each node has a service IP address. The system also has one or two management IP addresses.

If Ethernet ports are used for I/O, nodes are configured with additional IP addresses. See Ethernet configuration details for more details.

#### Configuration node

At any one time, a single node in the system provides a focal point for configuration commands. If the configuration node fails, another node in the system takes over its responsibilities.

The configuration node also handles the following functions:

- Accepts user logins to the management GUI and CLI via the IP address that is bound to Ethernet port 1 and optionally to the second management IP address that is bound to Ethernet port 2.
- Sends system configuration changes to the other nodes in the system.
- Sends call home and system notifications on behalf of the system.

If the configuration node fails or is taken offline, the system chooses a new configuration node. This action is called configuration node failover. The new configuration node takes over the management IP addresses. Thus, you can access the system through the same IP addresses although the original configuration node failed. During the failover, there is a short period when you cannot use the command-line tools or management GUI.

For correct failover operation, all nodes of the system must be connected to the same subnet and the management IP configuration must be valid on the subnet.

#### Ethernet Link Failures

If the Ethernet link to the system fails because of an event that is unrelated to the system, the system does not attempt to fail over the configuration node to restore management IP access. For example, the Ethernet link can fail if a cable is disconnected or an Ethernet router fails. To protect against this type of failure, the system provides the option for two Ethernet ports that each have a management IP address. If you cannot connect through one IP address, attempt to access the system through the alternative IP address.

### Identify the configuration node using the GUI

To identify the configuration node, follow these steps:

1. In a web browser, navigate to the Service Assistance Tool (SAT) of any node in the system, by entering the following address:

   ```
   https://<service_ip>/service
   ```

   Where `service_ip` is the service IP address of the chosen node.
2. Log into the SAT by authenticating using the superuser credentials.
3. Once logged in, the dashboard displays a list of all the nodes in the system. The configuration node is labelled as **CONFIG**.
4. Click the **LED** toggle to turn on the identification LED on the node canister.
5. Check for the node canister with the identification LED lit and perform the required service action.
6. When finished, click the LED toggle to turn off the identification LED on the node canister.

### Identifying the configuration node using the command-line interface

To identify the configuration node, follow these steps:

1. In a terminal window, use Secure Shell (SSH) software to connect to the service IP address of any node in the system and authenticate using the superuser credentials:

   ```
   ssh superuser@service_ip
   ```

   Where `service_ip` is the service IP address of the chosen node.
2. Once logged in, run **`sainfo lsservicenodes`** command to display a list of all the nodes that can be serviced by using the service assistant CLI. From the list of nodes, check the `panel_name` field to determine the list of all nodes in the same system as the node that is running the command.
3. For each of the nodes above, run the following command to display the service status of the node:

   ```
   sainfo lsservicestatus
   ```

   The configuration node is the one which reports a `config_node` value of Yes.
4. Turn on the identification LED for the configuration node canister by running the following command:

   ```
   satask chnodeled -on panel_name
   ```

   Where `panel_name` is the panel name of the configuration node.
5. Check for the node canister with the identification LED lit and perform the required service action.
6. When finished, turn off the identification LED for the configuration node canister by running the following command:

   ```
   satask chnodeled -off panel_name
   ```

   Where `panel_name` is the panel name of the configuration node.

## Node status

The *node status* identifies the status of the node within the IBM Storage Virtualize system.

Table 1 on page 7 describes the operational status of a node or node canister. To display the node status on the GUI, select **Monitoring** > **System Hardware**. You can also see the node status by executing the **`lsnode`** or **`lsnodecanister`** command or using the service assistant GUI.

| *Table 1. Node status* | |
|---|---|
| **Status** | **Description** |
| Active | The node is online and can process I/O from hosts. |
| Adding | The node or node canister was added to the clustered system but is not yet synchronized with the system status. The node or node canister status changes to Online after synchronization is complete.<br>**Note:** A node can stay in the Adding status for a long time. Wait at least 30 minutes before you take further action. If after 30 minutes the node status is still Adding, you can delete the node and add it again. If the node that was added is at a lower code level than the rest of the system, the node is upgraded to the system code level, which can take up to 20 minutes. During this time, the node is shown as Adding. |
| Candidate | The node is not defined in an I/O group; however, it can be added as one of the node pairs in an I/O group. |
| Deleting | The node or node canister is in the process of being deleted from the system. |
| Flushing | The system is removing data from the cache for the node. |
| Online | The node canister is operational and is assigned to a cluster. |
| Online spare | The node has been defined to be a hot-spare node and it is being used as a surrogate for another node in the I/O group. |
| Offline | The node canister is not operational. Run the fix procedures to determine the problem. |
| Pending | The status of the node or node canister is in transition. In a few seconds, the status of the node or node canister will change. |
| Service | The node or node canister has been placed in this state to allow it to be serviced or it has a fatal node error, which means that it is not safe for this node to be part of the system. |
| Spare | The node has been defined to be a hot-spare node; however, it is not being used by the system. The node has a valid cluster ID and a valid node ID. |
| Starting | The node has been restarted. It is attempting to join the system cluster however, it cannot yet perform I/O. |

## Routing considerations

When you configure any of the protocols for event notifications, consider these routing decisions.

### Routing considerations for event notification and Network Time Protocol

The system supports the following protocols that make outbound connections:

- Email
- Simple Network Mail Protocol (SNMP)
- Syslog
- Network Time Protocol (NTP).

These protocols operate only on a port that is configured with a management IP address. When it is making outbound connections, the system uses the following routing decisions:

- If the destination IP address is in the same subnet as one of the management IP addresses, the system sends the packet immediately through the Ethernet port with the matching subnet as the destination IP.
- If the destination IP address is not in the same subnet as either of the management IP addresses, the system sends the packet to the default gateway for Ethernet port 1.

- If the destination IP address is not in the same subnet as either of the management IP addresses and Ethernet port 1 is not connected to the Ethernet network, the system sends the packet to the default gateway for Ethernet port 2.

When you configure any of these protocols for event notifications, use these routing decisions to ensure that error notification works correctly if the network fails.

## I/O group

When a write operation is performed to a volume, the node that processes the I/O duplicates the data onto the partner node that is in the I/O group. After the data is protected on the partner node, the write operation to the host application is completed. The data is physically written to the disk later.

*Volumes* are logical disks that are presented to the system by nodes. Volumes are also associated with the I/O group.

When an application server processes I/O to a volume, it can access the volume with either of the nodes in the I/O group. When you create a volume, you can specify a preferred node. Many of the multipathing driver implementations that the system supports use this information to direct I/O to the preferred node. The other node in the I/O group is used only if the preferred node is not accessible.

If you do not specify a preferred node for a volume, the system selects the node in the I/O group that has the fewest volumes to be the preferred node. After the preferred node is chosen, it can be changed only when the volume is moved to a different I/O group.

**Note:** The management GUI provides a wizard that moves volumes between I/O groups without disrupting host I/O operations.

To view the current preferred node for a volume, select **Volumes** > **All Volumes** in the management GUI. Right-click the volume and select **Properties**.

To access information about I/O groups in the management GUI, select **Monitoring** > **System**. In the **System - Overview**, you can view the configured I/O groups on the system. The **System - Overview** page displays all the hardware that is assigned to the I/O groups that are configured on the system. Use the directional arrow to expand details on the I/O groups or its related hardware.

Read I/O is processed by referencing the cache in the node that receives the I/O. If the data is not found, it is read from the disk into the cache. The read cache can provide better performance if the same node is chosen to service I/O for a particular volume.

I/O traffic for a particular volume is, at any one time, managed exclusively by the nodes in a single I/O group.

Figure 1 on page 9 shows a write operation from a host (1) that is targeted for volume A. This write is targeted at the preferred node, Node 1 (2). The write operation is cached and a copy of the data is made in the partner node, the cache for Node 2 (3). The host views the write as complete. Later, the data is written, or de-staged, to storage (4).

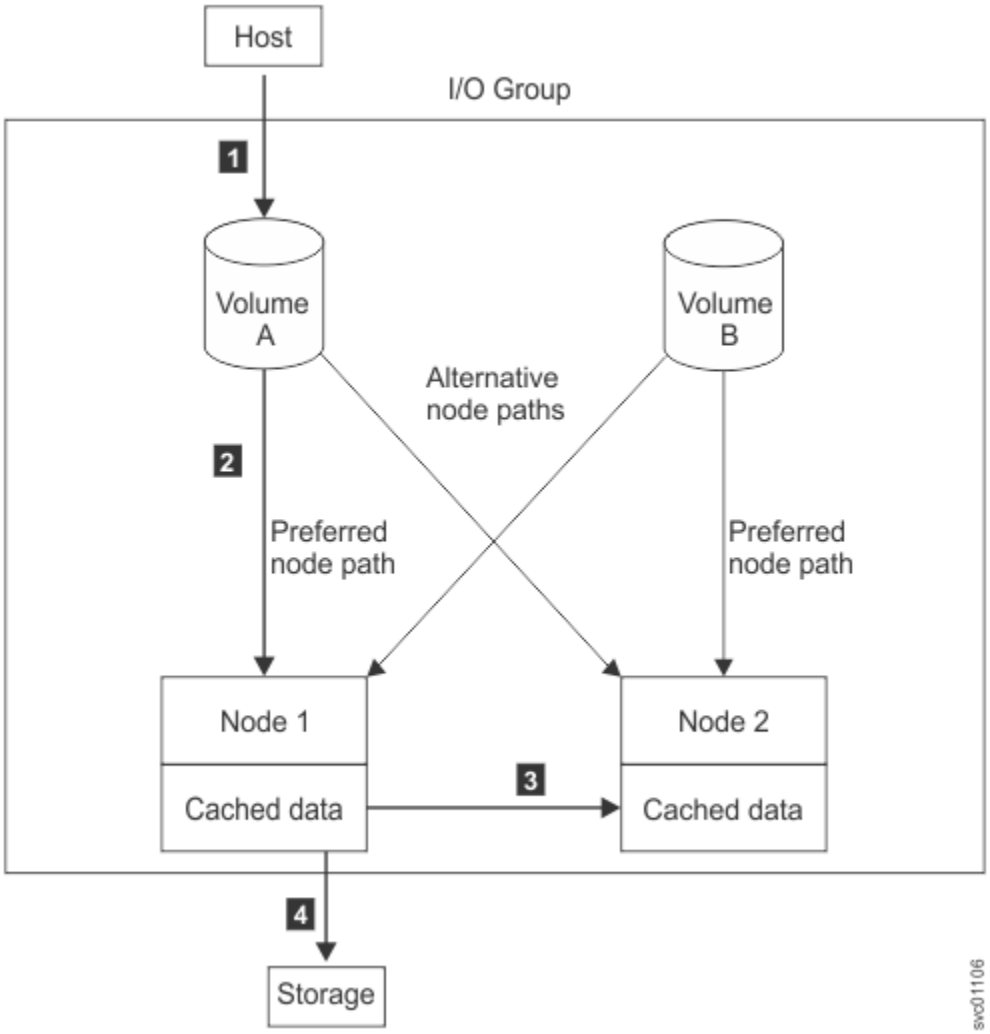


*Figure 1. Example of a write operation in an I/O group*

When a node fails within an I/O group, the other node in the I/O group assumes the I/O responsibilities of the failed node. Data loss during a node failure is prevented by mirroring the I/O read and write data cache between the two nodes in an I/O group.

If only one node is assigned to an I/O group or if a node fails in an I/O group, the cache is flushed to the disk and then goes into write-through mode. Therefore, any writes for the volumes that are assigned to this I/O group are not cached; they are sent directly to the storage device. If both nodes in an I/O group go offline, the volumes that are assigned to the I/O group cannot be accessed.

When a volume is created, the I/O group to provide access to the volume must be specified. However, volumes can be created and added to I/O groups that contain offline nodes. I/O access is not possible until at least one of the nodes in the I/O group is online.

## Partnerships

Partnerships are used to connect systems together to enable migration, data replication, and high-availability solutions.

A system can have partnerships with up to three remote systems. The connectivity for each partnership can be either Fibre Channel or IP. Systems also become indirectly associated with each other through partnerships. If two systems each have a partnership with a third system, those two systems are indirectly associated. A maximum of four systems can be directly or indirectly associated with each other.

A partnership configuration requires actions on both systems involved. This ensures that there is authority to access each system and share data between them.

You can create a partnership in the following ways:

- By using Fibre Channel connectivity
- By using IP connectivity.

## Policy-based Replication

To use partnerships for policy-based replication (asynchronous or high-availability), IP connectivity is required between management IP addresses of partnered systems.

Replication management requires access to the REST API on the remote system. Ensure that firewalls between the systems allow inbound traffic to port 7443 on the system management IP address.

The management traffic uses authentication certificates to prevent unauthorized access and ensure secure communications between the systems. Therefore, ensure that valid authentication certificates are installed on both the systems. For more information, see Managing certificates for secure communications.

## Background copy management

Certain types of replication differentiate between foreground host writes and background synchronisation traffic. The background copy rate is specified as a percentage of the partnership link bandwidth that is available to background synchronisation activities. Policy-based replication treats all traffic as background work so the background copy rate should be set to 100% if the system is using only policy-based replication.

HyperSwap, Metro Mirror, and Global Mirror (if supported by your system) use the background copy rate to control synchronisation traffic. Multi-cycling Global Mirror (Global Mirror with Change Volumes) uses only background copy therefore to achieve the best possible recovery point the background copy rate should be set to 100%. If you are using Metro Mirror, HyperSwap or non-cycling Global Mirror on your system, a lower value should be used to ensure that there is sufficient bandwidth to replicate host writes.

## Replication between IBM Storage Virtualize systems

Systems that run IBM Storage Virtualize software are in one of two layers: the replication layer or the storage layer.

- A SAN Volume Controller system is always in the replication layer.
- A FlashSystem is in the storage layer by default, but the system can be configured to be in the replication layer instead.

To create a partnership between systems, both systems must be in the same layer. For more information, including how to change the layer, see System layers.

## Partnership states

The state of the partnership helps determine whether the partnership operates as expected. A partnership can have the following states:

**Configured**
Both the local and remote systems have a partnership that is defined and are running as expected.

**Partial Local**

For the partnership to be fully configured, you must create a partnership from the remote system to the local system.

**Local Stopped**
Indicates that the partnership is defined on both the systems, but the partnership is stopped on the local system.

**Remote Stopped**
Indicates that the partnership is the defined on both the systems, but the partnership is stopped on remote system.

**Partial Local Stopped**
Indicates that only the local system has the partnership that is defined and the partnership is stopped on the local system.

**Local Excluded**
Indicates that both the local system and the remote system have the partnership that is defined, but the local system is excluding the link to the remote system. This state usually occurs when the link between the two systems is compromised by too many errors or slow response times of the partnership.

**Remote Excluded**
Indicates that both the local system and the remote system are defined in a partnership, but the remote system is excluding the link to the local system. This state usually occurs when the link between the two systems is compromised by too many errors or slow response times of the partnership.

**Exceeded**
Indicates that the partnership is unavailable because the network of systems exceeds the number of systems that are allowed in partnerships. To resolve this error, reduce the number of systems that are in partnerships in this network.

**Not Present**
Indicates that the remote system is not visible. This state can be caused by a problem with the connectivity between the local and remote system or if the remote system is unavailable.

For more information, see Creating IP partnership.

# Partnerships using Fibre Channel connectivity

Partnerships can be established over Fibre Channel infrastructure.

## Environment

See the following pages to understand the Fibre Channel network and port configuration needed to support Fibre Channel partnerships.

- Fibre Channel Zoning
- Planning for more than four fabric ports per node canister
- Long-distance requirements for partnerships

## Creating a partnership using Fibre Channel connectivity by using the management GUI

1. Connect to the GUI on either system. Select **Copy Services** > **Partnerships and Remote Copy** and select **Create Partnership**.
2. To create a partnership for policy-based replication (asynchronous or high-availability), ensure that the **Use Policy-Based Replication** checkbox is selected.
3. To create a partnership for Global Mirror or Metro Mirror, nondisruptive system migration, or 3-site partnerships, deselect the **Use Policy-Based Replication** checkbox. The GUI also provides step guidance for configuring 3-site partnerships using 3-Site Orchestrator.
4. If the certificate retrieved from the second system is signed by an authority that is not yet recognized on this system, then from the **Validate certificate** select **Upload File** to upload the root certificate of the certificate authority that signed the partner system's certificate.

### Creating a partnership using Fibre Channel connectivity by using the command-line interface

To create a Fibre Channel partnership between the two systems, use the following command on both systems:

```
mkfcpartnership -linkbandwidthmbits <link_bandwidth_in_mbps> -backgroundcopyrate <percentage>
<remote_system_id | remote_system_name>
```

where **mkfcpartnership** command defines a new partnership created over a Fibre Channel connection. For more information, see **mkfcpartnership command**.

## Zoning for partnerships

Partnerships configured using Fibre-Channel connectivity require fabric zoning.

Redundant Fabrics shows an example of a configuration that uses dual redundant fabrics that can be configured for Fibre Channel connections. Part of each fabric is at the local system and the remote system. There is no direct connection between the two fabrics.

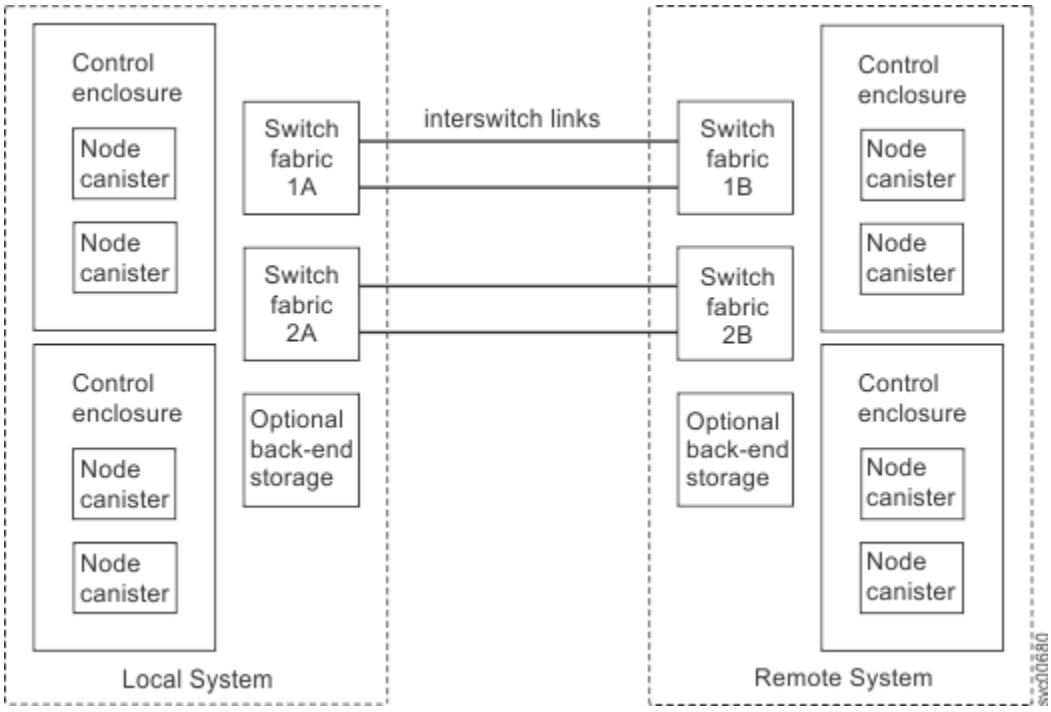


*Figure 2. Redundant fabrics*

### ISL vs FCIP configurations

Where the Fibre Channel fabric spans two physical sites, the inter-site connectivity can be achieved either by using ISLs to connect the sites, or by using Fibre Channel over IP (FCIP) routers to connect the sites by using an IP link.

Where FCIP is used, the configuration must provide guaranteed bandwidth for private (local node to node) traffic either by using separate dedicated links or by implementing QoS.

### Public / Private ports

Best practice is to ensure that cluster ports performing node to node traffic are segregated from those performing inter-system replication, or host I/O. This ensures that if there is an issue with a host, or with inter-system replication, any Fibre Channel credit loss does not impact local system's node to node communication. 2-4 ports per node have visibility of nodes in required I/O groups on the partnered system. Usually this is 1 port per redundant fabric.

### Inter Site Latency (Asynchronous)

Both asynchronous policy-based replication and Global Mirror support up to 250 ms RTT with appropriate zoning.

### Policy-based replication (Asynchronous or HA) zoning

Zone two ports from each node or canister in the first cluster I/O group with two ports from each canister in the second cluster I/O group. If dual-redundant fabrics are available, zone one port from each node across each fabric to provide the greatest fault tolerance. No other Fibre Channel ports on any node should have remote zones.

### Metro Mirror and Global Mirror (where RTT is less than 80 ms) zoning

For Metro Mirror and Global Mirror configurations where the RTT between systems is less than 80 ms, zone two Fibre Channel ports on each node in the local system to two Fibre Channel ports on each node in the remote system. If dual-redundant fabrics are available, zone one port from each node across each fabric to provide the greatest fault tolerance. No other Fibre Channel ports on any node should have remote zones.

**Optional:** Reducing the number of nodes that are zoned together can reduce the complexity of the intersystem zoning and might reduce the cost of the routing hardware that is required for large installations. Reducing the number of nodes also means that I/O must make extra hops between the nodes in the system, which increases the load on the intermediate nodes and can increase the performance impact, especially for Metro Mirror configurations.

### Global Mirror (where RTT is greater than 80 ms) zoning

If the RTT between systems is greater than 80 ms, stricter configuration requirements apply:

- Use SAN zoning and port masking to ensure that two Fibre Channel ports on each node that is used for replication are dedicated for replication traffic.
- Apply SAN zoning to provide separate intersystem zones for each local-to-remote I/O group pair that is used for replication. See the information about long-distance links for Metro Mirror and Global Mirror partnerships for further details.

**Optional:** As an alternative, choose a subset of nodes in the local system to be zoned to the nodes in the remote system. Minimally, you must ensure that one whole I/O group in the local system has connectivity to one whole I/O group in the remote system. I/O between the nodes in each system is then routed to find a path that is permitted by the configured zoning.

### Host Zoning

For policy-based High Availability, hosts need to be zoned so they can see the host ports from both systems.

## Partnerships using IP Connectivity

Partnerships can be established over Ethernet links that use the IPv4 and IPv6 addresses associated with Ethernet ports. These IP partnerships can be connections through Ethernet switches, or direct connections between local and partner systems. Partnerships can be created using an IPv4 address, IPv6 address, or the domain name of the remote system.

### IP partnership requirements

When you create and manage IP partnerships, consider the following requirements and constraints:

- A system can be part of up to three IP partnerships and only one of the remote systems can be at Storage Virtualize software level below 8.4.2.0.

- You can create IP partnerships between systems that are in the same layer. In other words, systems in the IP partnership must be at the storage layer or both systems must be at the replication layer.
- A secured IP partnership uses certificate-based authentication to create a secure communication channel between the partnered systems. You must ensure that the following requirements are met so that replicated data is encrypted and communication between partnered systems is secure. For more information, see "Planning secured IP partnerships" on page 17.
- You cannot use link-local addressing.
- If you use IPv4 addressing, the management IP addresses on both systems must be IPv4-compliant, and these addresses must have connectivity with each other.
- If you use IPv6 addressing, the management IP addresses on both systems must be IPv6-compliant, and these addresses must have connectivity with each other.
- You must configure all links between the production and recovery site with either IPv4 or IPv6 addresses.
- A system can have simultaneous partnerships over Fibre Channel and IP but with separate systems.
- Data compression is supported for IPv4 or IPv6 partnerships. To enable data compression, both systems in an IP partnership must be running a software level that supports IP partnership compression.
- To fully enable compression in an IP partnership, each system must enable compression. When compression is enabled on the local system, by using the **mkippartnership** or **chpartnership** commands, data sent to the remote system is compressed. To send compressed data to the local system, the remote system in the IP partnership must also enable compression.
- The use of WAN optimization devices such as Riverbed are not supported in native Ethernet IP partnership configurations.
- Bandwidth limiting on IP partnerships between both sites is supported.
- IP partnerships are not supported on 1 GB Ethernet ports on IBM Storage FlashSystem 9500 and SAN Volume Controller SV3 nodes.
- Using an Ethernet switch to convert a 25 GB to a 1 GB IP partnership, or a 10 GB to a 1 GB partnership is not supported.
- The IP infrastructure on both partnership sites must match.

## Port configuration requirements

- Specific ports are used by systems for IP partnerships (long-distance using TCP and short-distance using RDMA) communications. See the table TCP and UDP ports that are supported in section *TCP and UDP ports* for more information.
- The maximum supported round-trip time between systems is 80 milliseconds (ms).
- For IP partnerships, the recommended method of copying is by using Global Mirror with change volumes because of the performance benefits. Also, Global Mirror and Metro Mirror might be more susceptible to the loss of synchronization.
- *Portsets* are groupings of logical addresses that are associated with the specific traffic types. The system supports both Fibre Channel and IP portsets for host attachment, IP portsets for backend storage connectivity, and IP replication traffic. The system supports a maximum of 72 portsets.
- Each IP partnership can be mapped to two portsets, one for each link between systems. For a partnership with a single link, a single portset can be defined with the **-link1** attribute in the **mkippartnership** command for partnerships with a single intersite link. For a partnership with dual links, a second portset must be defined with the **-link2** attribute to specify the second portset for a dual link configuration. You can also use the management GUI to specify portset for each intersite link.
- You can configure ports from at most two I/O groups from each system for an IP partnership.
- If your system supports Ethernet ports of several speeds , **only one** of those speeds can be used to configure all of the remote copy links between the local and remote site.

**Note:** You can use 25 Gbps Ethernet ports to establish an IP partnership between systems. However, 25 Gbps Ethernet ports do not provide an additional performance advantage over IP partnerships that are established using 10 Gbps Ethernet ports.

- iSCSI hosts can access volumes over IP ports that are participating in an IP partnership; however, this access might result in an impact on performance.
- VLAN tagging of the IP addresses configured for remote copy is supported.
- IP partnerships do not support Network Address Translation (NAT) traversal. Most NAT implementations change the IP address and ports of the IP packets when the packets traverse the NAT or a network firewall. IP replication does not work with transport protocols that change the IP header source and destination IP addresses and their associated port numbers.
- If you configure two intersite WAN links, you must assign each WAN link to separate portset, one for each link.
- If you have one intersite link, you must configure one remote-copy port group for that link.
- No more than two intersite links are supported.
- If you have one remote-copy port group, then configure one port from each node in one I/O group in that remote-copy port group.
- For systems with more than one I/O group, ports from a second I/O group can be added to the remote-copy port group.
- If you have two remote-copy port groups and one I/O group, then on each system, configure one port from one node in the first remote-copy port group. Then, configure a port from the other node in the second remote-copy port group.
- For systems with more than one I/O group, ports from a second I/O group can be added to each of the two remote-copy port groups.
- Only one port in a node can be configured in an IP partnership.
- If you connect systems by directly attaching them without switches, you must have only two direct-attach links. Both direct-attach links must be on the same I/O group. You should use two port groups, where a port group contains only the two ports that are directly linked.

If you choose to use compression and IP partnership in the same system, certain hardware updates or configuration choices might help increase IP partnership performance:

- If you choose to create an IP partnership on a system that has compressed volumes, and you have multiple I/O groups, configure ports for the IP partnership in I/O groups that do not contain compressed volumes.
- Use a different port for IP partnership traffic, which is not used for iSCSI host, iSCSI virtualization, iSCSI backend, iSER host attachment, and RDMA clustering. Also, use a different VLAN ID for iSCSI host I/O and IP partnership traffic.

## Requirements for short-distance partnerships using RDMA

When you create and manage short-distance partnerships using RDMA, consider the following requirements and constraints:

- You can create short-distance partnership using RDMA between systems that are in the replication layer. In other words, systems in the short-distance partnership using RDMA must be at the storage layer or both systems must be at the replication layer.
- The system supports a maximum of 6 highspeedreplication portsets.
- You cannot use link-local addressing.
- Configure all links between the production and recovery site with either IPv4 or IPv6 addresses.
- You can configure ports at most with two I/O groups from each system for a short-distance partnership using RDMA.

**Note:** You can use 10 or 25 Gbps Ethernet ports with RDMA (iWARP) capabilities to establish a short-distance partnership using RDMA between systems. For support for RoCE Ethernet ports, contact IBM support.

- A system can have simultaneous partnerships over native-IP and short-distance partnership using RDMA but with separate remote systems.
- Systems that are configured in active short-distance partnership using RDMA must not be zoned with each other for Fibre Channel.
- VLAN tagging of the IP addresses configured for highspeedreplication portsets is supported.
- Short-distance partnership using RDMA do not support Network Address Translation (NAT) traversal. Most NAT implementations change the IP address and ports of the IP packets when the packets traverse the NAT or a network firewall. IP replication does not work with transport protocols that change the IP header source and destination IP addresses and their associated port numbers.
- If you have one intersite link, you must assign one highspeedreplication portset for that link that is, all ports that belong to the portset must be connected by using that link.
- If you configure two inter-site links, you can configure two highspeedreplication portset and assign one portset per link.
- No more than two intersite links are supported.
- If you have one highspeedreplication portset, then configure maximum of two ports from each node in one I/O group in that highspeedreplication portset.
- If you have two highspeedreplication portsets and one I/O group, then on each system, configure maximum of two ports from one node in the first highspeedreplication portset. Then, configure maximum of two ports from the other node in the second highspeedreplication portset.
- For systems with more than one I/O group, ports from a second I/O group can be added to each of the two highspeedreplication portsets.
- If you connect systems by directly attaching them without switches, you must have only two direct-attach links. Both direct-attach links must be on the same I/O group. Use two port groups, where a port group contains only the two ports that are directly linked.

**Related tasks**

“Planning secured IP partnerships” on page 17
A secured IP partnership uses certificate-based authentication to create a secure communication channel between the partnered systems. You must ensure that the following requirements are met so that replicated data is encrypted and communication between partnered systems is secure.To enable a secured IP partnerships, must install an encryption license on both of the partnered systems. See Activating encryption license for more information, or contact support for assistance.

## Configuring VLAN for IP partnerships

To configure VLAN when you use IP (Internet Protocol) partnerships, consider the following requirements and procedures.

### Before you begin

- VLAN tagging is supported for IP partnership traffic between systems.
- VLAN provides network traffic separation at the layer 2 level for Ethernet transport.
- VLAN tagging by default is disabled for any IP address of a node port. You can use the management GUI or the command-line interface (CLI) to optionally set the VLAN ID for port IPs on systems in the IP partnership.
- When a VLAN ID is configured for the port IP addresses that are mapped to the portsets, appropriate VLAN settings on the Ethernet network must also be properly configured to prevent connectivity issues.
- Setting VLAN tags for a port is disruptive. Therefore, VLAN tagging requires that you stop the partnership first before you configure VLAN tags. Then, restart again when the configuration is complete.

### About this task

Follow this procedure to configure VLAN tags for existing IP partnership setups:

### Procedure

1. Stop the partnership between the local and remote system.
2. Configure VLAN on node ports in the portsets on the local system.
3. Configure all intervening switches with appropriate VLAN tags.
4. Configure VLAN on node ports in the portsets on the remote system.
5. Check to see whether connectivity between the local and remote sites are restored.
6. Restart the partnership.

## Planning secured IP partnerships

A secured IP partnership uses certificate-based authentication to create a secure communication channel between the partnered systems. You must ensure that the following requirements are met so that replicated data is encrypted and communication between partnered systems is secure.To enable a secured IP partnerships, must install an encryption license on both of the partnered systems. See Activating encryption license for more information, or contact support for assistance.

### Before you begin

You must plan what type of system certificate the partnered systems will use. The system certificate is used for all functions on the system that use certificate authentication. A system certificate can be an internally signed certificate or an externally signed certificate. An internally signed certificate is a certificate that has been issued by the system's root certificate authority. If the system is currently using a self-signed certificate, the self-signed certificate can be used until it expires. After it expires, you must use either an internally signed certificate or an externally signed certificate. Externally signed certificates are issued and signed by a third-party CA. A signing request must be generated on the system and presented to the third-party CA, which signs the request and returns a signed certificate that can be installed. For more information about managing system certificates, see Managing certificates for secure communications.

**Mutual authentication**

A secured IP partnership between the partnered systems requires the systems to be mutually authenticated. Mutual authentication requires that the system certificate or certificate authority certificate for each local system is installed in a truststore on the partner system. The local system presents its certificate and chain of signing CA certificates over the network when the secured partnership is established. The partner system authenticates the certificate presented by the local system by using the certificates and authorities that are installed in its truststore.

**Setting up mutual authentication**

To set up mutual authentication, you must install certificates on the partnered systems based on the type of certificate you choose. The system supports both internally-signed certificates and externally signed certificates.

**Internally signed certificates**

- Create an internally signed certificate with either the management GUI or use the **svctask chsystemcert -mksystemsigned** command. The certificate is installed on the system automatically. To configure an internally-signed certificate, see Updating an internally signed certificate.
- Export the system's root certificate to the partnered system's truststore.

**Externally signed certificates**

- – Generate a certificate signing request (CSR) with the management GUI or use the **svctask chsystemcert -mkrequest** command.

- Send the CSR that is generated to your external CA.
- When the third-party CA returns the signed certificate, install the certificate and the intermediate CA certificates on the system.
- Install the third-party CA's root certificate in the partnered system's truststore. To configure an externally-signed certificate, see Requesting and installing an externally signed certificate.

The certificates are added to a system-wide truststore on the partnered system and managed with a set of CLI commands. For more information, see **`Truststore management commands`**.

When using an externally signed certificate, the certificate can be signed by the third-party root CA or by an intermediate CA which itself is signed by the third-party root CA. If an intermediate CA is used then there is a certificate chain, and the complete chain of trust must be established in order for one system to authenticate the other. The complete chain of trust can be established in the following ways:

- Only the signed system certificate is installed on the local system. The local system presents only the endpoint certificate over the network and the intermediate CA certificates and root CA certificate must be installed in the truststore on the partner system.
- The signed system certificate along with the intermediate CA certificates (the root CA certificate can optionally be included but is not required) are installed on the local system. The local system presents the endpoint certificate and the intermediate CA certificates over the network, and only the root CA certificate must be installed in the truststore on the partner system.

It is recommended that the full certificate chain is installed on the local system. If the full chain is not installed then errors in other functions that use certificate authentication such as the management GUI can occur.

The network planning requirements for Secured IP partnership are same as the existing IP partnership requirements. For information about IBM Storage Virtualize for Public Cloud networking considerations, see Planning networking for Microsoft Azure.

Best practices for secured IP partnerships:

- In a secured IP partnership, the partnered systems mutually authenticate with each other using certificates. These certificates have a validity period for which they are valid, and expire after this time. The certificate for each system can be configured separately, and may have a different validity period to the partner system's certificate. Therefore, it is possible for one system's certificate to expire while the partner system's certificate is still valid. In this scenario, the certificate authentication for the IP partnership will fail. It is the best practice to:
  - Ensure the certificates of each system in the partnership are configured with similar validity periods to avoid disruption in the IP partnership.
  - Synchronize the date and time of each system. If using NTP, it is recommended to use the same NTP server for all the remote copy peer systems.
- User gets better encryption efficiency and better replication throughput in secured IP partnerships with higher MTU size. However, larger MTU sizes might not be possible outside controlled IP networks. Users need to plan this carefully with network administrators.
- If MTU sizes configured at endpoints of a secured or unsecured IP partnership do not match, the partnership may remain in *not_present* state. Hence it is recommended to configure same MTU size at all endpoints such as IBM Storwize® devices and network devices that are involved for establishing an IP partnership.

## Long-distance Ethernet partnerships using TCP

Partnerships are used to connect systems together to enable migration, data replication, and high-availability solutions.

*Secured IP partnerships* secure the data as it travels through an untrusted network between production and recovery systems. Although securing IP partnerships is optional, it supports authentication of the

production and recovery systems, and verifies the confidentiality and integrity of the replicated data. Secured IP partnerships minimize the risk of hackers manipulating or intercepting data in untrusted networks. Secured IP partnerships use Internet Protocol Security (IPsec) suite of protocols that covers important security aspects such as:

- Enhanced mutual authentication
- Stronger encryption algorithms
- Encryption key management mechanisms

In secured IP partnerships, the partner systems authenticate to each other, negotiate the security parameters, exchange encryption keys, and establish secured network tunnels through which encrypted data travels. Partner systems are authenticated by certificates issued by internal root certificate authorities (CA) or trusted third-party's root CA or intermediate CA. Secured IP partnerships are created when the necessary certificates and authorities are installed on the partner systems.

To enable secured partnerships, you must purchase an encryption license and activate the license on both partnered systems. For more information, see Activating encryption license.

Portsets replace the requirement for creating groups for IP partnerships. Dedicated portsets can be created for remote copy traffic. The dedicated portsets provide group of IP addresses for IP Partnerships. Each node can have one IP address that is assigned to a portset for traffic. If the local system in the IP partnership contains four nodes, a portset can be created that defines four IP addresses, one per each node. Similarly, the remote system with four nodes, a portset on that system can also have four IP addresses to handle traffic exclusively. Before you can configure a new IP partnership, you need to define a portset and assign IP addresses to nodes.

You can configure portsets so that each IP partnership can be mapped to two portsets, one for each WAN link between systems. For network configurations that have a single link between systems in an IP partnership, a single portset can be defined in the **Portset Link 1** field on the **Create Partnership** page from GUI. For a partnership with dual links, a second portset must be mapped defined in the **Portset Link 2** field.

## Supported IP partnership configurations

The following general configurations are supported, but the number of I/O groups that are configured for each site can be different.

**Configuration 1:** In this configuration, only a single WAN intersite link is available. Therefore, only one portset is configured on each node.

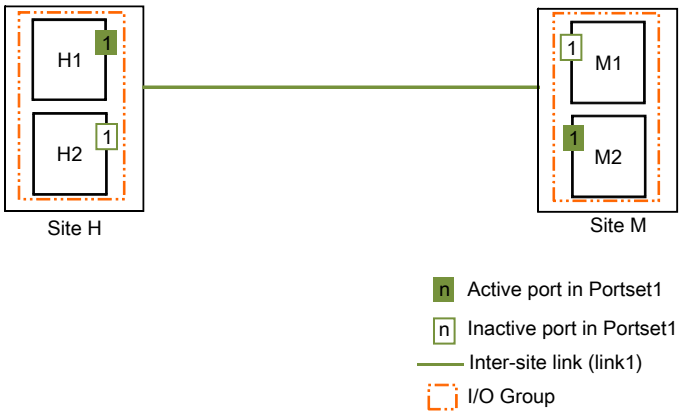


*Figure 3. One intersite link, one I/O group per system*

Only one port from either of the nodes in each system actively participates in the IP partnership. The other port acts as the failover port. If a critical failure is observed on node H1 in Site H, the IP partnership will fail over to node H2 and continue. Remote copy relationships might stop momentarily during the failover.

**Note:** The system supports two IP partnerships that use the same physical link between systems however the single link can have lower throughput than a single IP partnership.

**Configuration 2:** In this configuration, only one intersite link is available. Each system uses a single portset where each node in the system has an IP address assigned. However, out of all of the available ports, only one port from either of the nodes in each system actively participates in the IP partnership. The other ports act as failover ports.

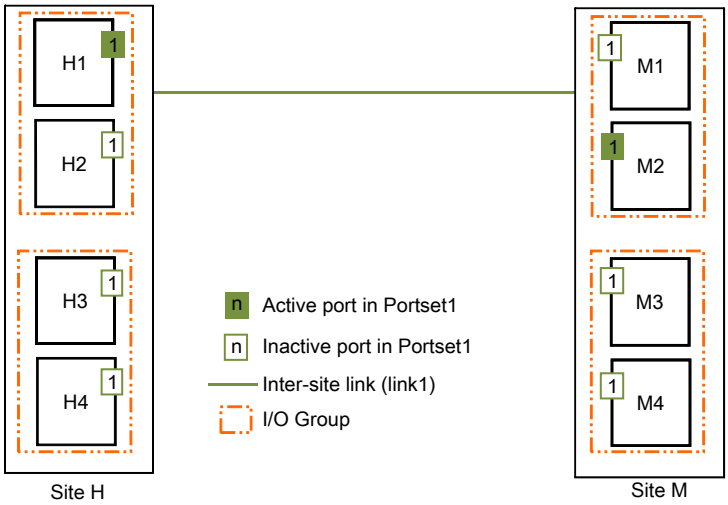


*Figure 4. One intersite link, two I/O groups per system*

If a critical failure is observed on node H1 in Site H, the IP partnership fails over to node H2, H3, or H4 and continues. Remote copy relationships might stop momentarily during the failover.

**Configuration 3:** In this configuration, eight-node systems are available. However, only two I/O groups in a system can have ports that are configured in IP partnerships. Each system uses a single portset where each node in the system has an IP address assigned. In this configuration, each system has a portset with eight IP addresses. However, out of all the available ports, only one port from either node in each system actively participates in IP partnership. The other ports act as failover ports.

**Note:** Configuration 3 also applies to systems with four I/O groups. In such systems, while only two I/O groups can have ports configured in IP partnerships, all I/O groups in a system can contain remote-copy relationships. Any replication-related operations that are generated by nodes that are not connected directly to the remote system is forwarded to connected nodes for onward transmission to the remote system.

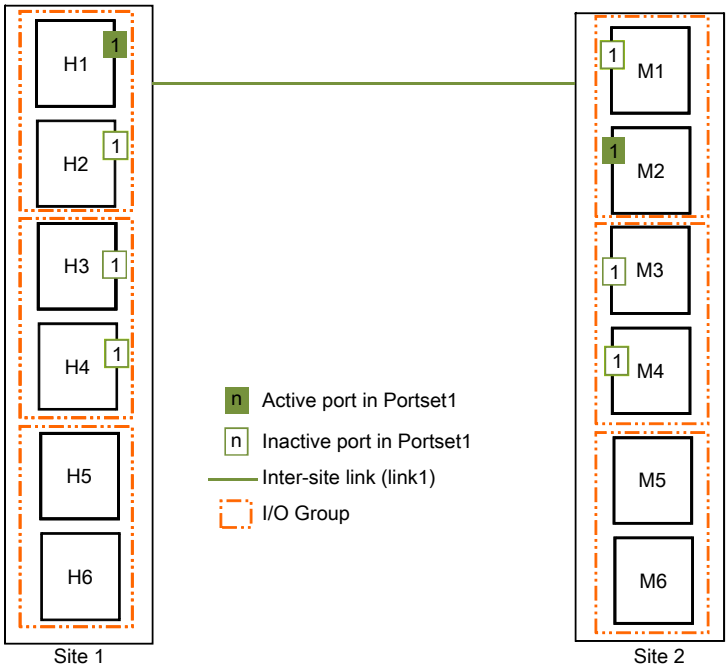


*Figure 5. One intersite link, three I/O groups per system*

If a critical failure is observed on node H1 in Site 1, the IP partnership fails over to node H2, H3, or H4 and continues. Remote copy relationships might stop momentarily during the failover.

**Configuration 4:** In this configuration, two intersite links are available; therefore, two portsets are configured. One port from each node in each system actively participates in the IP partnership. If a critical failure is observed on node H1 in Site H, the IP partnership continues over the other port on node H2.

No failure occurs in this scenario; however, the effective bandwidth is reduced to half; only one of the two links is available to facilitate IP partnership traffic. When the failure is corrected, ports will fail back and the IP partnership continues to operate over both links.

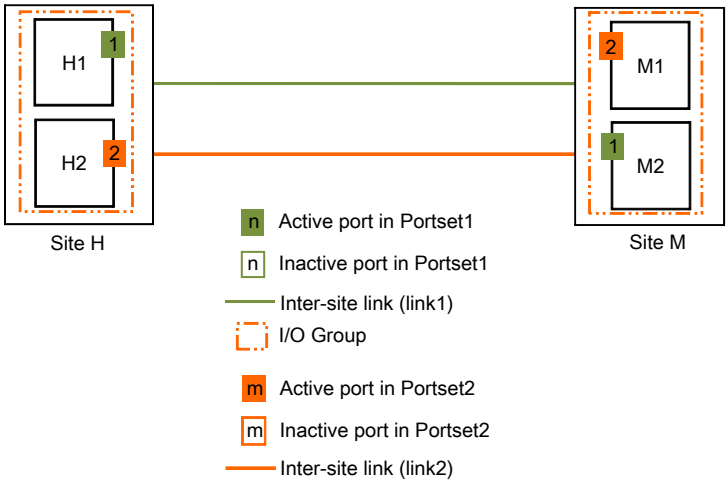


*Figure 6. Two intersite links, one I/O group per system*

**Configuration 5:**In this multi-node configuration, two intersite links are available. Each link must be assigned to a different portset. Each portset contains one IP address for each node in the system. Out of the four ports, only two ports actively facilitate the IP partnerships. This port and path selection is maintained by an internal algorithm. The other ports act as failover ports.

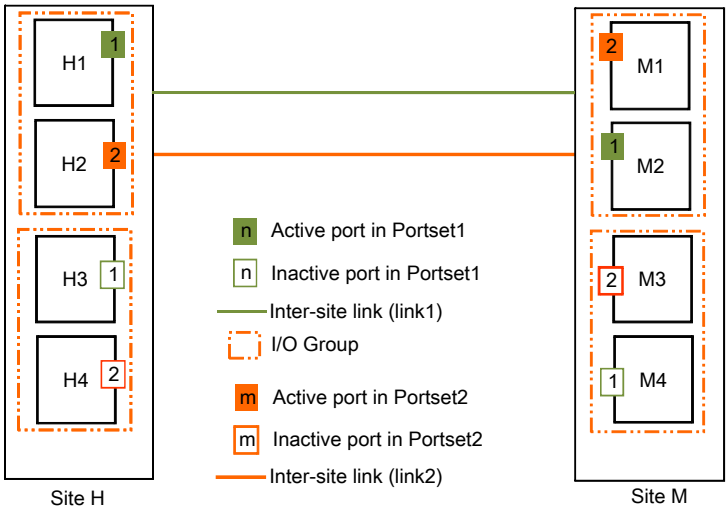


*Figure 7. Two intersite links, two I/O groups per system*

If a critical failure occurs on node H1 in Site H, the IP partnership will fail over to node H3 and continue. The link bandwidth is not affected, as the failover happens immediately and completes quickly while IP partnership traffic continues from node H2.

**Configuration 6:** In this configuration, eight-node systems are available. However, only two I/O groups in a system can have ports that are configured in IP partnerships. In this multi-node configuration, two intersite links are available. Each link must be assigned to a different portset. Each portset contains one IP address for each node in the system. Out of the four ports, only two ports actively facilitate the IP partnerships. This port and path selection is maintained by an internal algorithm. The other ports act as failover ports.

**Note:** Configuration 6 also applies to systems with four I/O groups. In such systems, while only two I/O groups can have ports that are configured in IP partnerships, all I/O groups in a system can contain remote-copy relationships. Any replication-related operations that are generated by nodes that are not connected directly to the remote system are forwarded to connected nodes for onward transmission to the remote system.

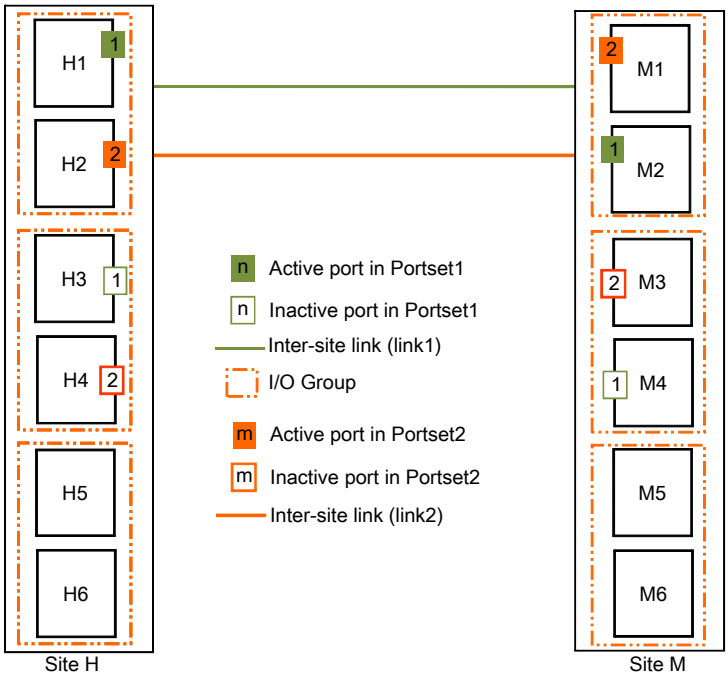


*Figure 8. Two intersite links, three I/O groups per system*

If a critical failure is observed on the node H1 in Site H, the IP partnership will fail over to node H3 and continue. The link bandwidth is not affected because the failover happens immediately and completes quickly while IP partnership traffic continues from node H2. The other I/O groups, or all the I/O groups, can be connected over a Fibre Channel partnership with another system.

### *Intersite link planning*

Specific intersite link requirements must be met when you are planning to use IP partnership for remote copy.

If you use IP partnership, you must meet the following requirements:

- IP partnerships can also support data compression. Both systems in the IP partnership must be running a software level that supports IP partnership data compression. Compression must also be functional on each system. For more information, see "Long-distance Ethernet partnerships using TCP" on page 18.
- The amount of intersite heartbeat traffic is 1 megabit per second (Mbps) per link.
- The minimum bandwidth requirement for the intersite link is 10 Mbps. However, this requirement scales up with the amount of host I/O that you choose to do. Figure 9 on page 23 describes the scaling of host I/O.

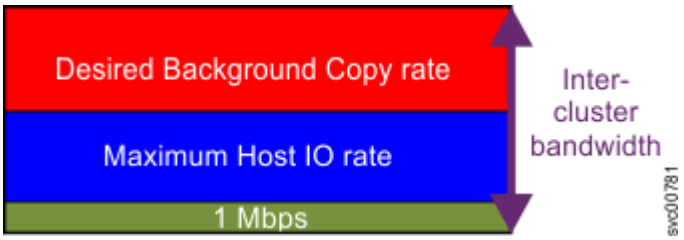


*Figure 9. Scaling of host I/O*

The equations that can describe the approximate minimum bandwidth that is required between two systems with < 5 ms round-trip time and errorless link follow.

For systems that use Global Mirror or Metro Mirror:

```
Minimum intersite link bandwidth in Mbps > Maximum(Minimum Link Bandwidth of
10Mbps , Required Background Copy in Mbps + Maximum Host I/O in Mbps + 1 Mbps
heartbeat traffic)
```

For systems that use only Global Mirror with change volumes:

```
Minimum intersite link bandwidth in Mbps > Maximum(Minimum Link Bandwidth of
10Mbps , Required Background Copy in Mbps + 1 Mbps heartbeat traffic)
```

Increasing latency and errors results in a higher requirement for minimum bandwidth.

Figure 10 on page 24 shows how the IBM Storage Virtualize solution maintains near line-speed performance by masking the latency of the line. Even as the line latency increases, the performance of the technology enables the line to continue to exceed that of a plain link. The throughput limit is the uncompressed data rate, which is measured as the remote node send data rates. The round-trip time (RTT) is measured as the IP link RTT (pinging between the two nodes on the same ports that are being used for replication). The throughput limit is based on the distance and the uncompressed data rate.

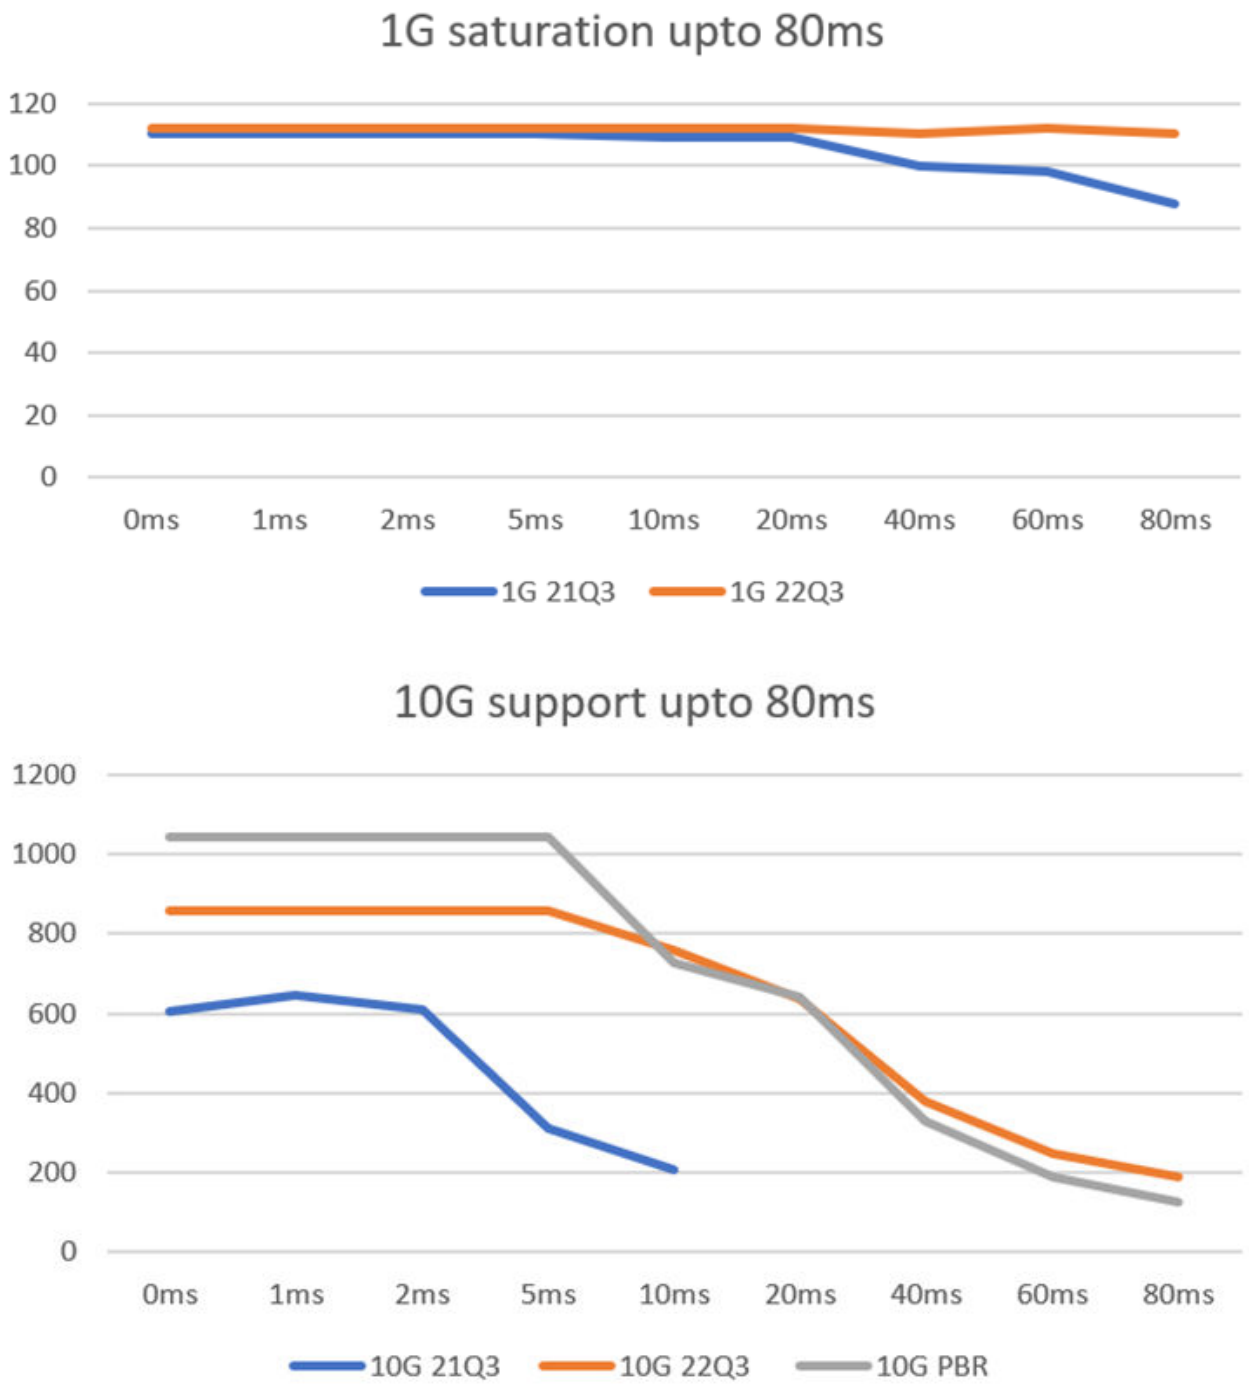


*Figure 11. Product transfer rate comparison for systems built with IBM Storage Virtualize – Latest data for 10G link (Grey line)*

*Figure 10. Product transfer rate comparison for systems built with IBM Storage Virtualize – Latest data for 1G link (orange line)*

### *IP partnership performance best practices*

A number of factors affect the performance of an IP partnership. Some of these factors are latency, link speed, number of intersite links, host I/O, MDisk latency, and hardware.

The following actions might help improve this performance. See "IP partnership requirements" on page 13 for additional considerations.

To avoid losing synchronization during node failover, use policy-based replication or Global Mirror with change volumes. These modes of replication provide consistency protection and restarts replication automatically.

## Short-distance partnerships using RDMA

Data replication is important for the integrity, availability, and disaster recovery of critical business information. High-performance computing environments and ultra-fast network infrastructures need replication solutions that match technological advancements. *Short-distance partnerships using RDMA* allows data replication with high bandwidth over a short distance. Therefore, enterprises can use their existing Ethernet infrastructure to fully adopt an end-to-end Ethernet-based solution.

Short-distance partnership using RDMA offers several advantages:

**Increased Throughput**
Short-distance partnership using RDMA is suited for throughput-hungry applications, offering compatibility with link bandwidths that surpass 10 Gbps. Experience data replication at unprecedented speeds.

**Optimized for Short Distances**
For enterprises that require disaster recovery (DR) solutions within a 100-kilometer range, short-distance partnership using RDMA provides dependable links, minimizing frame loss and maximizing data integrity.

**Low Round-Trip Time**
Short-distance partnership using RDMA is ideal in scenarios up to 1 ms Round-Trip Time (RTT). Replicate data seamlessly across racks, within data centers, or throughout city-based campuses, while the highest standards of performance are maintained.

**Comprehensive Ethernet-Based Solution**
Suitable for enterprises who want Ethernet-only environment such as Ethernet data centers deployments for Business Continuity and Resiliency.

**Note:** SVC nodes would still require FC clustering to form a cluster.

**Total Cost of Ownership (TCO)**
Short-distance partnership using RDMA is cheaper than traditional FC partnerships because expensive FCIP routers are not required for this solution.

Some of the major use cases that this solution enables include:

- *3-Site replication.* All three sites can operate only in the Ethernet environment.
- *Setting up DR Sites in sync and async environments.* This solution enables enterprises to set up multiple DR sites, where replication partnerships over IP networks maintain data volume copies, both in a synchronous and asynchronous environment.

### *Planning for short-distance partnerships using RDMA*

Planning for short-distance partnership using RDMA involves defining the partnerships between two systems and associating portsets.

**Note:** Short-distance partnership using RDMA is only supported on RDMA ports. In addition, only iWARP ports are allowed by default. To configure short-distance partnership by using RoCE adapters, contact IBM Support.

For short-distance partnership using RDMA partnerships, you must specify highspeedreplication portsets. The following are the requirements for highspeedreplication portset:

- Node IP addresses in Service GUI must not be preconfigured in the system.
- Both link1 and link2 portsets must be of type *highspeedreplication* portsets.
- All ports in a *highspeedreplication* portset must be of an identical type; that is either both must be iWARP or RoCE.
- All the IP addresses within a *highspeedreplication* type portset must be of the identical type that is, either IPv4 or IPv6.

The following are the system limits for short-distance partnerships using RDMA :

| *Table 2. System limits* | |
|---|---|
| **Limit** | **Description** |
| Max IP addresses per portset per node | 2 |
| Max *highspeedreplication* portset per system | 6 |
| Max portset per link | 1 |
| Max links per partnership | 2 |

## *Intersite link planning*

Specific intersite link requirements must be met when you are planning to use IP partnership for remote copy.

If you use IP partnership, you must meet the following requirements:

- IP partnerships can also support data compression. Both systems in the IP partnership must be running a software level that supports IP partnership data compression. Compression must also be functional on each system. For more information, see “Long-distance Ethernet partnerships using TCP” on page 18.
- The amount of intersite heartbeat traffic is 1 megabit per second (Mbps) per link.
- The minimum bandwidth requirement for the intersite link is 10 Mbps. However, this requirement scales up with the amount of host I/O that you choose to do. Figure 12 on page 26 describes the scaling of host I/O.

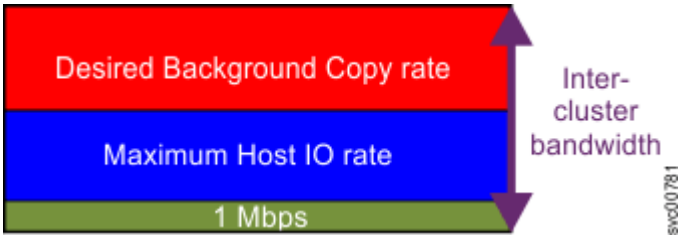


*Figure 12. Scaling of host I/O*

The equations that can describe the approximate minimum bandwidth that is required between two systems with < 5 ms round-trip time and errorless link follow.

For systems that use Global Mirror or Metro Mirror:

```
Minimum intersite link bandwidth in Mbps > Maximum(Minimum Link Bandwidth of
10Mbps , Required Background Copy in Mbps + Maximum Host I/O in Mbps + 1 Mbps
heartbeat traffic)
```

For systems that use only Global Mirror with change volumes:

```
Minimum intersite link bandwidth in Mbps > Maximum(Minimum Link Bandwidth of
10Mbps , Required Background Copy in Mbps + 1 Mbps heartbeat traffic)
```

Increasing latency and errors results in a higher requirement for minimum bandwidth.

Figure 13 on page 27 shows how the IBM Storage Virtualize solution maintains near line-speed performance by masking the latency of the line. Even as the line latency increases, the performance of the technology enables the line to continue to exceed that of a plain link. The throughput limit is the uncompressed data rate, which is measured as the remote node send data rates. The round-trip time (RTT) is measured as the IP link RTT (pinging between the two nodes on the same ports that are being used for replication). The throughput limit is based on the distance and the uncompressed data rate.

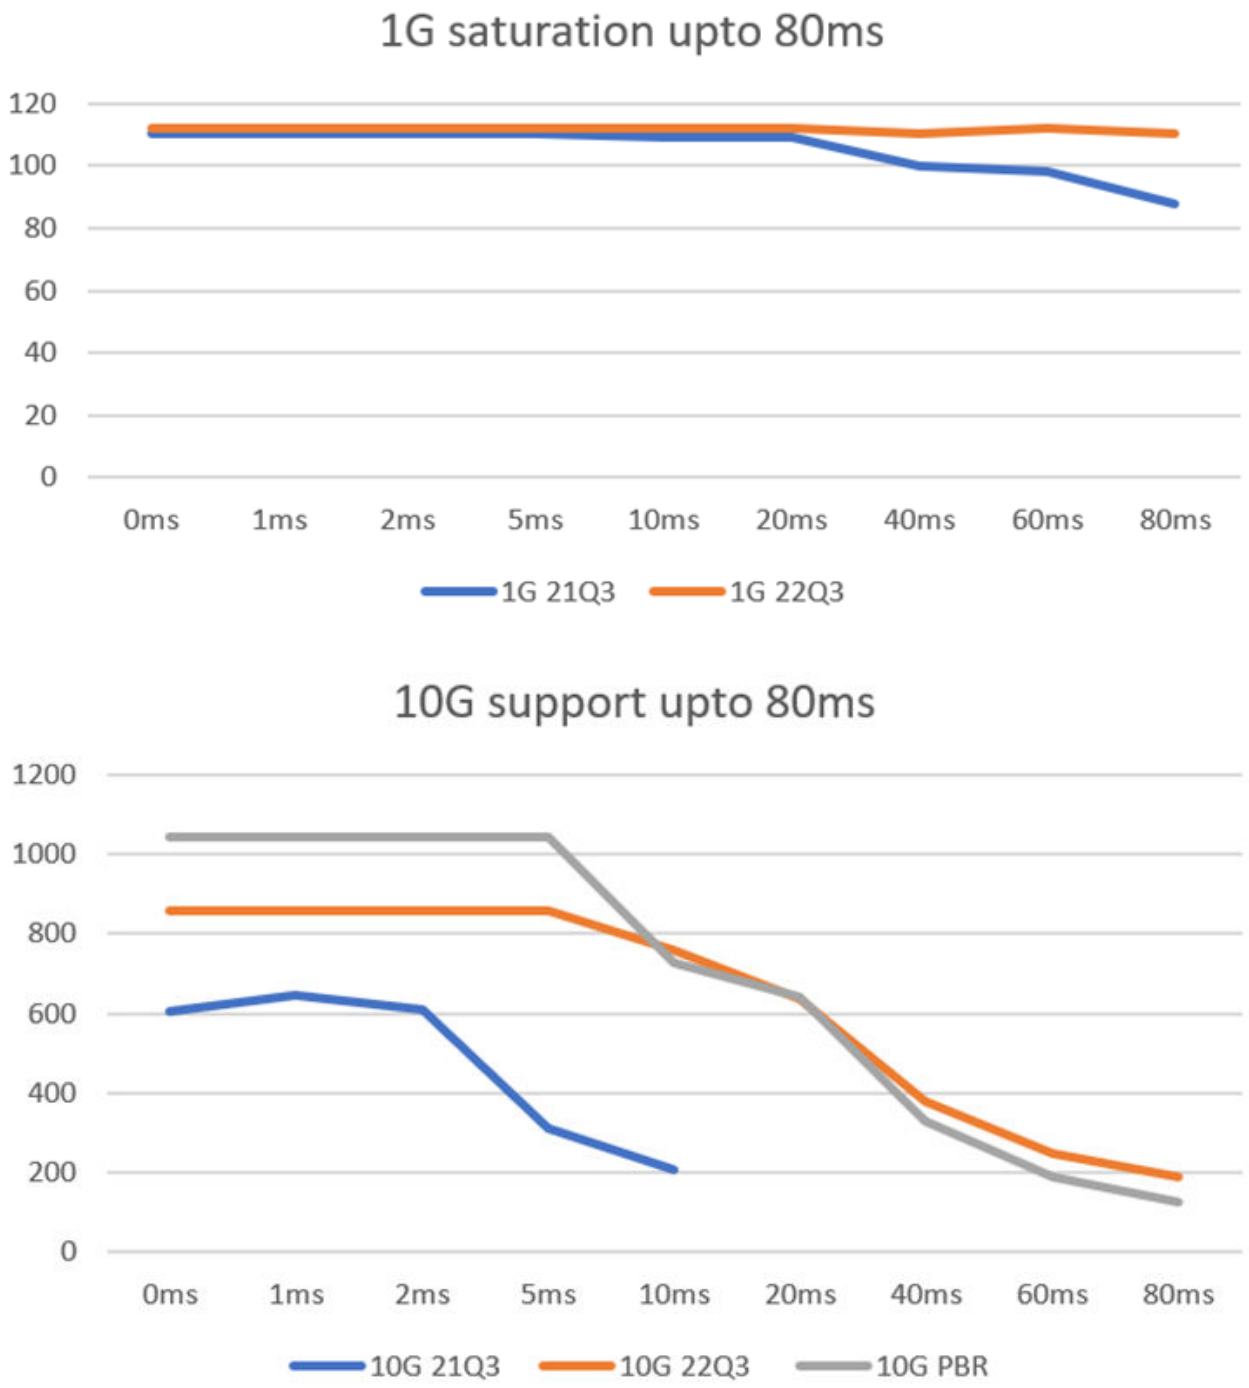


*Figure 14. Product transfer rate comparison for systems built with IBM Storage Virtualize – Latest data for 10G link (Grey line)*

*Figure 13. Product transfer rate comparison for systems built with IBM Storage Virtualize – Latest data for 1G link (orange line)*

## Long-distance requirements for partnerships

The links between systems in a partnership that are used for replication must meet specific configuration, latency, and distance requirements.

The following table lists the maximum round-trip latency for each type of partnership.

| *Table 3. Maximum supported round-trip latency between sites* | |
|---|---|
| **Partnership** | |
| **FC** | **IP** |
| 250 ms | 80 ms |

- Replication requires a specific amount of bandwidth for intersystem heartbeat traffic. When the system uses a Fibre Channel partnership, the amount of traffic depends on the number of nodes that are in both the local system and the remote system. Table 4 on page 28 provides a guideline for the intersystem heartbeat traffic between the systems. These numbers represent the total traffic between two systems when no I/O operations run on the copied volumes. Half of the data is sent by the local system and half of the data is sent by the remote system. Therefore, traffic is evenly divided between all of the available intersystem links. If you have two redundant links, half of the traffic is sent over each link.

| *Table 4. Intersystem heartbeat traffic in Mbps* | | | | |
|---|---|---|---|---|
| **System 1** | **System 2** | | | |
| | **2 nodes** | **4 nodes** | **6 nodes** | **8 nodes** |
| 2 nodes | 5 | 6 | 6 | 6 |
| 4 nodes | 6 | 10 | 11 | 12 |
| 6 nodes | 6 | 11 | 16 | 17 |
| 8 nodes | 6 | 12 | 17 | 21 |

- In a Metro Mirror or non-cycling Global Mirror relationship, the bandwidth between two sites must meet the peak workload requirements and maintain the maximum round-trip latency between the sites. When you evaluate the workload requirement in a multiple-cycling Global Mirror relationship, you must consider the average write workload and the required synchronization copy bandwidth. If there are no active synchronization copies and no write I/O operations for volumes that are in the Metro Mirror or Global Mirror relationship, the system protocols operate with the bandwidth that is indicated in Intersystem heartbeat traffic in Mbps. However, you can determine only the actual amount of bandwidth that is required for the link by considering the peak write bandwidth to volumes that are participating in Metro Mirror or Global Mirror relationships and then adding the peak write bandwidth to the peak synchronization bandwidth.
- If the link between two sites is configured with redundancy so that it can tolerate single failures, the link must be sized so that the bandwidth and latency statements are correct during single failure conditions.

## Configuration requirements for remote copy partnerships over extended distances

If you use remote mirroring between systems with 80 - 250-ms round-trip latency, you must meet the following extra requirements:

- Both the local and remote systems must support the higher round-trip latency.
- A Fibre Channel partnership must exist between the systems, not an IP partnership.
- The RC buffer size setting must be 512 MB on each system in the partnership. This setting can be accomplished by running the **`chsystem -rcbuffersize 512`** command on each system.

  **Note:** Changing the RC buffer size is disruptive to replication operations. Change this settings before partnerships are created or stop all partnerships before you change this value.
- SAN zoning should be applied to provide separate intersystem zones for each local-remote I/O group pair that is used for replication. Figure 15 on page 29 illustrates this type of configuration.

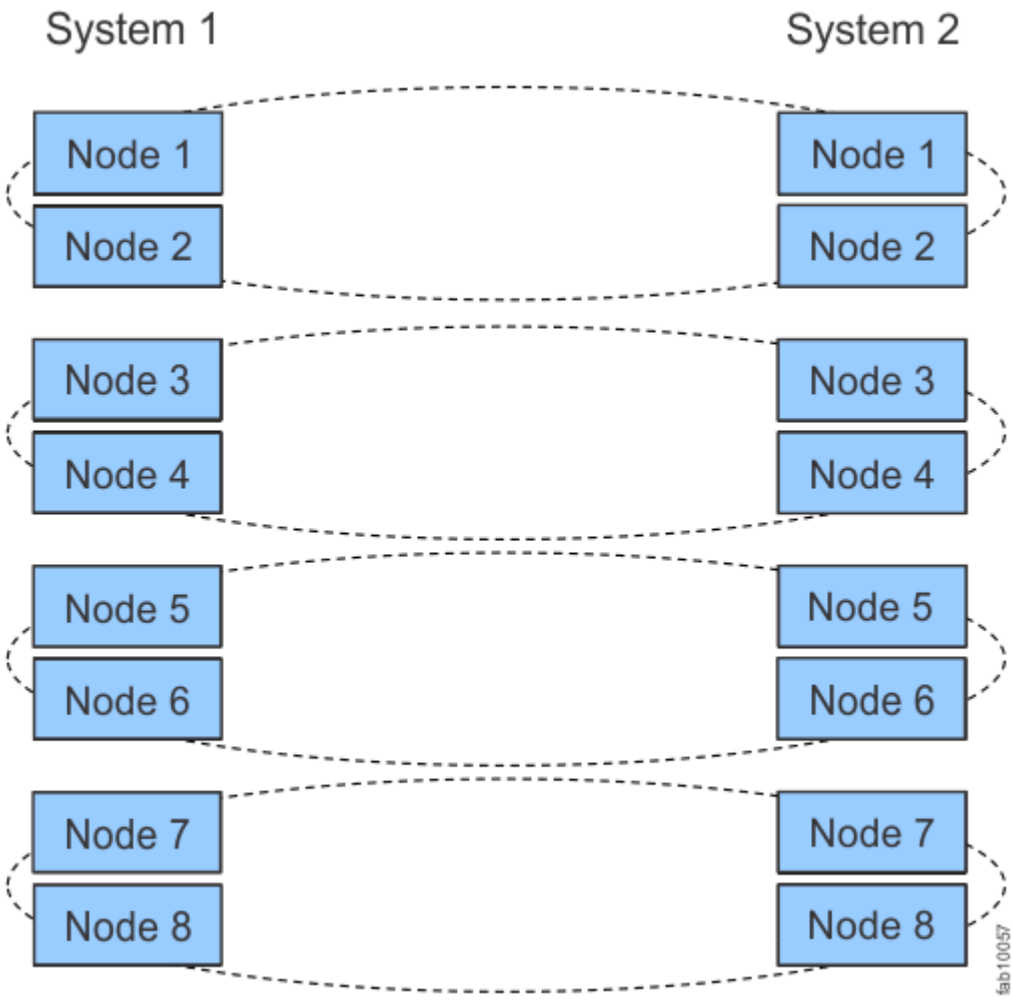


*Figure 15. Intersystem SAN zoning*

In addition to the preceding list of requirements, the following guidelines are provided for optimizing performance:

- Partnered systems should use the same number of nodes in each system for replication.
- For maximum throughput in replication that uses Global Mirror, all nodes in each system should be used for replication, both in terms of balancing the preferred node assignment for volumes and for providing intersystem Fibre Channel connectivity.
- On the system, provisioning dedicated node ports for local node-to-node traffic (by using port masking) isolates replication node-to-node traffic between the local nodes from other local SAN traffic. As a result, optimal response times can be achieved.
- Where possible, use the minimum number of partnerships between systems. For example, assume site A contains systems A1 and A2, and site B contains systems B1 and B2. In this scenario, creating separate partnerships between pairs of systems (such as A1-B1 and A2-B2) offers greater performance for replication between sites than a configuration with partnerships that are defined between all four systems.

## Pools

A pool or storage pool is an amount of capacity that is allocated to volumes that are created in that pool. Use the Pools page in the management GUI to configure and manage storage pools, internal and external storage, MDisks, and to migrate existing storage to the system.

This figure shows a basic parent pool with associated child pools. In this graphic, the *usable capacity* for the parent group is divided between two child pools. Usable capacity is the amount of capacity that is

available for storing data on a parent pool after formatting and RAID techniques are applied. Volumes can then be created by using either the capacity from the MDisks through the parent pool or from the child pool.

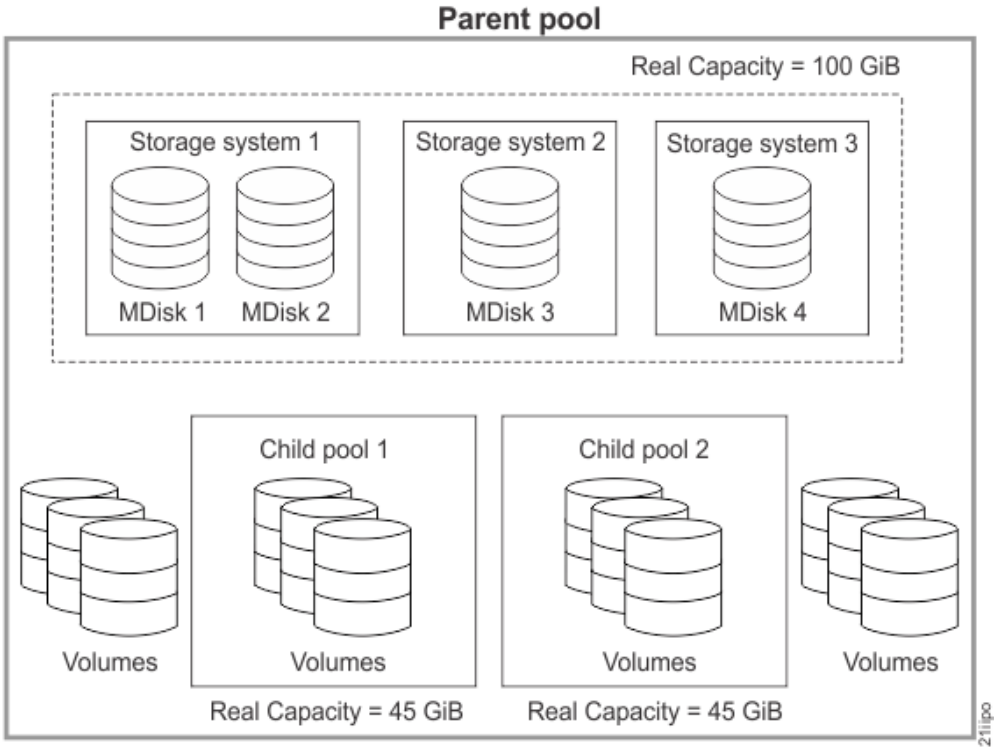


*Figure 16. Storage pool*

## Parent Pools

A parent pool is a collection of managed disks (MDisks). When the pool is created, the managed disks are split into extents. Volumes are created from those extents with the data striped across managed disks.

To track the space that is available on an MDisk, the system divides each MDisk into chunks of equal size. These chunks are called *extents* and are indexed internally. Extent sizes can be 16, 32, 64, 128, 256, 512, 1024, 2048, 4096, or 8192 MB. The choice of extent size affects the total amount of storage that is managed by the system.

You specify the extent size when you create a new parent pool. You cannot change the extent size later; it must remain constant throughout the lifetime of the parent pool.

You cannot use the data migration function to migrate volumes between parent pools that have different extent sizes. However, you can use volume mirroring to move data to a parent pool that has a different extent size.

Use volume mirroring to add a copy of the disk from the destination pool. After the copies are synchronized, you can free up extents by deleting the copy of the data in the source pool. The FlashCopy® function and Metro Mirror can also be used to create a copy of a volume in a different pool.

A system can manage 2^22 extents. For example, with a 16 MB extent size, the system can manage up to 16 MB x 4,194,304 = 64 TB of storage.

When you choose an extent size, consider your future needs. For example, if you currently have 40 TB of storage and you specify an extent size of 16 MB for all parent pools, the capacity of the system is limited to 64 TB of storage in the future. If you select an extent size of 64 MB for all parent pools, the capacity of the system can grow to 256 TB.

Using a larger extent size can waste storage. When a volume is created, the storage capacity for the volume is rounded to a whole number of extents. If you configure the system to have many small volumes and you use a large extent size, storage can be wasted at the end of each volume.

When you create or manage a parent pool, consider the following general guidelines:

- The **Pools** page in the management GUI displays the **Usable Capacity** and **Capacity Details**. *Usable capacity* indicates the amount of capacity that is available for volumes before any capacity savings methods are applied.
- An MDisk can be associated with just one parent pool.
- You can specify a warning threshold for a pool. A warning event is generated when the amount of used capacity in the pool exceeds the warning threshold. The warning threshold is especially useful with thin-provisioned volumes that are configured to automatically use capacity from the pool.
- Volumes can have one or two volume copies. A volume copy is associated with just one pool, except when you migrate a volume copy between parent pools. A volume with two volume copies can have each volume copy in a different pool.
- You can only add MDisks that are in unmanaged mode. When MDisks are added to a parent pool, their mode changes from unmanaged to managed.
- You can delete MDisks from a parent pool under the following conditions:
  - Volumes are not using any of the extents that are on the MDisk.
  - Enough free extents are available elsewhere in the pool to move any extents that are in use from this MDisk.
  - The system ensures that all extents that are used by volumes in the child pool are migrated to other MDisks in the parent pool to ensure that data is not lost.
- You can delete an array MDisk from a parent pool when:
  - Volumes are not using any of the extents that are on the MDisk.
  - Enough free extents are available elsewhere in the parent pool to move any extents that are in use from this MDisk.
- If the volume is mirrored and the synchronized copies of the volume are all in one pool, the mirrored volume is destroyed when the storage pool is deleted. If the volume is mirrored and there is a synchronized copy in another pool, the volume remains after the pool is deleted.

## Child Pools

Instead of being created directly from MDisks, child pools are created from existing usable capacity that is assigned to a parent pool. Volumes can be created that specifically use the usable capacity that is assigned to the child pool.

When a standard child pool is created, the usable capacity for a child pool is reserved from the usable capacity of the parent pool. The usable capacity for the child pool must be smaller than the usable capacity in the parent pool. After the child pool is created, the amount of usable capacity that is specified for the child pool is no longer reported as usable capacity of its parent pool. When a data reduction child pool is created, the usable capacity for the child pool is the entire usable capacity of the data reduction parent pool without limit. After a data reduction child pool is created, the usable capacity of the child pool and the usable capacity of the parent pool are reported as the same.

When you create or work with a child pool, consider the following general guidelines:

- You can create and manage child pools in the management GUI.
- As with parent pools, you can specify a warning threshold that alerts you when the used capacity of the child pool is reaching its upper limit. Use this threshold to ensure that access is not lost when the used capacity of the child pool is close to its usable capacity.
- You cannot shrink the usable capacity of a child pool below its used capacity. The system also resets the warning level when the child pool is shrunk and issues a warning if the level is reached when the usable capacity is shrunk.
- On systems with encryption enabled, you can create the standard child pools to migrate existing volumes in a non-encrypted pool to encrypted child pools only when you virtualize external storage. When you create a standard child pool after encryption is enabled, an encryption key is created for the

child pool even when the parent pool is not encrypted. You can then use volume mirroring to migrate the volumes from the non-encrypted parent pool to the encrypted child pool. Encrypted data reduction child pools can be created only if the parent pool is encrypted. The data reduction child pool inherits an encryption key from the parent pool.

- Ensure that any child pools that are associated with a parent pool have enough usable capacity for the volumes that are in the child pool before removing MDisks from a parent pool. The system automatically migrates all extents that are used by volumes to other MDisks in the parent pool to ensure that data is not lost.
- You cannot shrink the usable capacity of a child pool below its used capacity. The system also resets the warning level when the child pool is shrunk and issues a warning if the level is reached when the usable capacity is shrunk.
- The system supports migrating a copy of volumes between child pools within the same parent pool or migrating a copy of a volume between a child pool and its parent pool. Migrations between a source and target child pool with different parent pools are not supported. However, you can migrate a copy of the volume from the source child pool to its parent pool. The volume copy can then be migrated from the parent pool to the parent pool of the target child pool. Finally, the volume copy can be migrated from the target parent pool to the target child pool.

Child pools can also be assigned to an ownership group. An *ownership group* defines a subset of users and objects within the system. You can create ownership groups to further restrict access to specific resources that are defined in the ownership group. Only users with Security Administrator roles can configure and manage ownership groups. Restricted users are those users who are defined to a specific ownership group and can only view or manage specific resources that are assigned to that ownership group. Unrestricted users are not defined to an ownership group and can manage any objects on the system based on their role on the system.

Ownership can be defined explicitly or it can be inherited from the user, user group, or from other parent resources, depending on the type of resource. Ownership of child pools must be assigned explicitly, and they do not inherit ownership from other parent resources. New or existing volumes that are defined in the child pool inherit the ownership group that is assigned for the child pool.

## Data Reduction Pools

To use data reduction technologies on the system, you need to create a data reduction pool, create volumes with the data reduction pool, and map these volumes to hosts that support SCSI unmap commands.

Data reduction is a set of techniques that can be used to reduce the amount of usable capacity that is required to store data. An example of data reduction includes data deduplication. Data reduction can increase storage efficiency and performance and reduce storage costs, especially for flash storage. Data reduction reduces the amount of data that is stored on external storage systems and internal drives by reclaiming previously used capacity that are no longer needed by host systems. To estimate potential capacity savings that data reduction can provide on the system, use the Data Reduction Estimation Tool (DRET). This tool analyzes existing user workloads that are being migrated to a new system. The tool scans target workloads on all attached storage arrays, consolidates these results, and generates an estimate of potential data reduction savings for the entire system.

For more information about DRET, see https://www.ibm.com/support/pages/node/6217841. For more information about Compretimator, see https://www.ibm.com/support/pages/node/6209688.

The system supports data reduction pools which can use different capacity savings methods simultaneously, increasing the capacity savings across the entire pool. Data reduction pools also support deduplication. When deduplication is specified for a volume, duplicate versions of data are eliminated and not written to storage, thus saving more usable capacity. Some models or software versions require specific hardware or software to use this function.

When you create a data reduction pool, ensure that the usable capacity of the pool includes *overhead capacity*. Overhead capacity is an amount of usable capacity that contains the metadata for tracking unmap and reclaim operations within the pool. A general guideline is to ensure that the provisioned

capacity with the data reduction pool does not exceed 85% of the total usable capacity of the data reduction pool. Table 5 on page 33 includes the minimum data reduction pool capacity that is required to be able to create a volume within the pool.

*Table 5. Minimum overhead capacity requirements for data reduction pools*

| **Extent size (in gigabytes)** | **Overhead capacity requirements (in terabytes)[1]** |
|---|---|
| 1 GB or smaller | 1.1 TB |
| 2 GB | 2.1 TB |
| 4 GB | 4.2 TB |
| 8 GB | 8.5 TB |

[1]Standard-provisioned volumes are not included into the minimum overhead capacity values. When you are planning usable capacity for data reduction pools, determine the usable capacity that is needed for any standard-provisioned volumes first, then ensure that the minimum usable capacity values for the data reduction pools are included.

## Pool states

This table describes the operational states of a pool. Child pools adopt the state of the parent pool. States that indicate an error must be resolved on the parent pool.

*Table 6. Pool states*

| **State** | **Description** |
|---|---|
| Online | The pool is online and available. All the MDisks in the pool are available. |
| Degraded paths | This state indicates that one or more nodes in the system cannot access all the MDisks in the pool. A degraded path state is most likely the result of incorrect configuration of either the storage system or the Fibre Channel fabric. However, hardware failures in the storage system, Fibre Channel fabric, or node might also be a contributing factor to this state. To recover from this state, follow these steps:<br>1. Verify that the fabric configuration rules for storage systems are correct.<br>2. Ensure that you configured the storage system properly.<br>3. Correct any errors in the event log. |

| *Table 6. Pool states (continued)* | |
|---|---|
| **State** | **Description** |
| Degraded ports | This state indicates that one or more 1220 errors were logged against the MDisks in the pool. The 1220 error indicates that the remote Fibre Channel port was excluded from the MDisk. This error might cause reduced performance on the storage system and usually indicates a hardware problem with the storage system. To fix this problem, you must resolve any hardware problems on the storage system and fix the 1220 errors in the event log. To resolve these errors in the log, click **Monitor** > **Events** in the management GUI. This action displays a list of unfixed errors that are currently in the event log. For these unfixed errors, select the error name to begin a guided maintenance procedure to resolve them. Errors are listed in descending order with the highest priority error listed first. Resolve highest priority errors first. |
| Offline | The pool is offline and unavailable. No nodes in the system can access the MDisks. The most likely cause is that one or more MDisks are offline or excluded.<br>  **Attention:** If a single MDisk in a pool is offline and cannot be seen by any of the online nodes in the system, the pool of which this MDisk is a member goes offline. This causes all of the volume copies that are being presented by this pool to go offline. Take care when you create pools to ensure an optimal configuration. |

### Easy Tier

A child pool inherits the Easy Tier® settings from its parent pool. You cannot change the Easy Tier settings on a child pool. You can only change them on a parent pool.

## Linked Pools

With policy-based replication, storage pool links define the pool that is used on the remote system to create the replicated volumes when replication policies exist. Replication policies can only be assigned to volume groups containing volumes in linked storage pools. When linking pools on a stretched topology system, only the pools that are in site 1 are required to be linked.

If the storage pools exist on the production and recovery systems, you can add a link between the pools from either system. If child pools currently exist on a single system only, you can use the management GUI on the partnered system to create and link a child pool in a single step. The management GUI simplifies the process of creating a linked pool on the partnered system. The management GUI automatically displays the properties such as name, capacity, and provisioning policy from the system where the child pools already exist. You can use these values to create the new linked child pool on the partnered system without logging in to the other system.

When you create pool links, you can assign a provisioning policy to the pools that you are linking. Provisioning policies control how capacity is provisioned on all volumes in the linked pool. Pools on each partnered system can be assigned provisioning policies with different capacity savings methods.

The system creates two default provisioning policies when the first parent pool is created. You can create more user-defined policies to specify alternative capacity savings. If a provisioning policy is not configured, the system automatically creates fully provisioned volumes.

### Linking pools in topologies with more than two systems

If a pool on one of the systems has existing links to another partnered system, you must add the link from the unlinked system. The existing link between pools for other partnerships is not affected.

### Creating and modifying links between pools

Using the **Pools** panel on the management GUI, you can:

- Create links between existing pools
- Modify links between existing pools
- Create and link child pools, where child pools already exist on one system.

To create linked child pools, use the GUI on the system that does not have the child pools created.

**Note:** In the **Properties** column on the **Pools** page, linked pools are marked with a link icon.

## Quorum

A quorum device is used to break a tie when a SAN fault occurs, when exactly half of the nodes that were previously a member of the system are present. A quorum device is also used to store a backup copy of important system configuration data.

Just over 256 MB is reserved to store backup copy of important system configuration data on each quorum device.

It is possible for a system to split into two groups where each group contains half the original number of nodes in the system. A quorum device determines which group of nodes stops operating and processing I/O requests. In this tie-break situation, the first group of nodes that accesses the quorum device is marked as the owner of the quorum device and as a result continues to operate as the system, handling all I/O requests. If the other group of nodes cannot access the quorum device or finds that the quorum device is owned by another group of nodes, it stops operating as the system and does not handle I/O requests.

A system can have only one active quorum device that is used for a tie-break situation. However, the system uses up to three quorum devices to record a backup of system configuration data to be used in the event of a disaster. The system automatically selects one quorum device to be the active quorum device. The other quorum devices provide redundancy if the active quorum device fails before a system is partitioned. To avoid the possibility of losing all the quorum devices with a single failure, assign quorum disk candidates on multiple storage systems or run IP quorum applications on multiple servers.

### Single site configurations

The normal configuration is to use a managed drive or an MDisk as the quorum device when the system is not configured as a stretched or HyperSwap system. A system automatically assigns quorum disk candidates. When you add new storage to a system or remove existing storage, however, it is a good practice to review the quorum disk assignments. Optionally an IP quorum device can be configured either as an alternative to using quorum disks or to provide additional redundancy.

### Stretched or HyperSwap configurations

To provide protection against failures that affect an entire location, such as a power failure, you can use a configuration that splits a single system across three physical locations.

A stretched or HyperSwap system has system nodes divided between two sites. If a SAN fault causes loss of connectivity between sites or a fault causes a site wide outage then the quorum configuration

determines which site continues operating and processing I/O requests. A high availability solution has the active quorum device configured at a third site so that the system will continue to operate after any single-site failure.

Generally, when the nodes in a system are split among sites, configure the system this way:

- Site 1: Half of system nodes + one quorum device
- Site 2: Half of system nodes + one quorum device
- Site 3: Active quorum device

Typically the quorum devices at site 1 and site 2 are quorum disks and the quorum device at site 3 is an IP quorum application. However, the system can be configured to use either quorum disks or IP quorum applications at any site. This configuration ensures that a quorum device is always available, even after a single-site failure.

When you are using an IP quorum application at a third site, you can configure a preference for which site continues operation if there is a loss of connectivity between the two sites. If only one site runs critical applications, you can configure this site as preferred. If a preferred site is configured and a failure causes an outage at the preferred site, the other site wins the tie-break and continues operating and processing I/O requests.

A stretched or HyperSwap system can be configured without a quorum device at a third site. If there is no third site, then quorum must be configured to select a site to always win a tie-break. If there is a loss of connectivity between the sites, then the site that is configured as the winner continues operating and processing I/O requests and the other site stops until the fault is fixed. If there is a site outage at the wining site, then the system stops processing I/O requests until this site is recovered or the manual quorum override procedure is used.

Generally, when the nodes in a system are split between two sites and there is no third site quorum, configure the system this way:

- Site 1: Half of system nodes + one or two quorum devices
- Site 2: Half of system nodes + one quorum device

Typically, the quorum devices at site 1 and site 2 are both quorum disks and are automatically configured by the system. It is possible to configure IP quorum applications as an alternative to using quorum disks. When a winner site has been configured and both sites are operational, there is no active quorum device. The quorum devices at site 1 and site 2 are only used to retain a backup copy of important system configuration data. If a failure results in just the nodes at the winner site continuing operation, then the system automatically selects one of the quorum devices at that site to be the active quorum device to protect against further failures.

## IP quorum application

The IP quorum application is a Java™ application that runs on a server that is separate from the storage system. A policy-based High Availability partnership requires an IP quorum application to arbitrate in case of a site or system loss. The IP quorum application can be generated on either system in the partnership. The partnership must be created before generating and deploying the IP quorum application.

The maximum number of IP quorum applications that can be deployed on a single system is five. This enables multiple servers to be used to provide redundancy. Only one instance of the IP quorum application per server, per system, is supported. For example, a server can run two IP quorum instances if each instance is connected to a different nonpartnered storage system. Ensure that bandwidth is available to support multiple IP quorum instances.

Do not deploy the IP quorum application on a server that depends on storage that is presented by the system. This action can result in a situation where the nodes need to detect the IP quorum application to process I/O, but cannot because the IP quorum application cannot access storage.

An Ethernet connectivity issue can prevent an IP quorum application from accessing a node that is still online and an event is raised on the system if this occurs.

Use the following support article to understand IP network requirements: https://www.ibm.com/support/pages/node/7013877

**Warning:** For the IP quorum application to be able to connect using TLS 1.3 (SSL Security Levels 6 and 7), the version of Java running the application must also support TLS 1.3.

## Configuring IP Quorum on the storage system

You can configure IP quorum on the storage system by using the GUI or CLI:

**Using the management GUI**
In the management GUI, the Policy-based replication setup procedure includes generating the IP quorum application.

**Using the CLI**
On the CLI, enter the command: **mkquorumapp -partnersystem <remote system ID or name>**

## Deploying the IP Quorum Application on the server

Follow these steps to deploy the IP quorum application on the server:

1. Create a separate directory that is dedicated to the IP quorum application.
2. Transfer the IP quorum application from the storage system to the dedicated directory.
3. Use the **ping** command on the server to verify that it can establish a connection with the service IP address of each node in both systems.
4. Enter the command **java -jar ip_quorum.jar** to initialize the IP quorum application.

**Note:** The IP quorum application needs to be running always.

The options available when running the IP quorum application can be listed with **-help**:

```
1. $ java -jar ip_quorum.jar -help
2. -name (optional) to help identify the IP quorum app instance in the GUI/lsquorum.  Can only
contain 1-20 characters that are A-Z, a-z or 0-9.
3. -debug (optional) run the app in debug mode to display verbose messages on stdout and in the
log file.
4. -emit (optional) display T3 metadata header information
5. -location (optional) default = ip_quorum.log.16845080043810. Specify the full, or relative
path to store your log files. Allowed characters: [a-z A-Z 0-9 .-_/]
6. -rotation (optional) default = 5. Specify the maximum number of log files. Allowed values:
[1-10]
7. -size (optional) default = 5120. Specify the log file size in kb. Allowed values:
[1024-10240]
8. -version / -v (optional) display the vrmf of the system when the app was generated,
timestamp and generation id.
```

For more information on configuring IP quorum on a Linux host, see Helpful resource and publications.

## Reviewing and monitoring the IP Quorum Application

You can monitor and review the IP quorum application from each system by using the management GUI or CLI:

**Using the management GUI**
In the management GUI, select **Settings** > **System** > **IP Quorum**.

**Using the CLI**
Enter the **lsquorum** command. For more information, see **lsquorum**.

The output will display application_type: partnership.

# Volumes

A volume is a logical disk that the system presents to attached hosts.

You can create different types of volumes, depending on the type of topology that is supported and configured on your system. All systems support standard topology, which is a single-site configuration. For systems with single-site configuration, you can create basic, mirrored, or custom volumes.

Volumes can be assigned to an ownership group. An *ownership group* defines a subset of users and objects within the system. You can create ownership groups to further restrict access to specific resources that are defined in the ownership group. Only users with Security Administrator roles can configure and manage ownership groups. Restricted users are those users who are defined to a specific ownership group and can only view or manage specific resources that are assigned to that ownership group. Unrestricted users are not defined to an ownership group and can manage any objects on the system based on their role on the system.

- The volume inherits the ownership group of the child pools that provide capacity for the volume and its copies. Volume copies can be created in different ownership groups for backup up scenarios. However, this value must be set intentionally by users that are not defined in ownership groups. When you create a volume copy or migrate volumes to other pools, you can specify child pools that are defined in different ownership groups in the management GUI, which establishes inconsistent ownership. However, it is not recommended to leave volumes or volume copies in different ownership groups. After the migration, the user with Security Administrator role needs to ensure all volumes or copies are within the same ownership group as the users who need access.
- With volume groups, the volume group and its volumes can belong to different ownership groups. However, the ownership of a volume group does not impact the ownership of the volumes that it contains.

## Types

Each volume copy can be one of the following types:

**Striped**
A volume copy that was striped is at the extent level. One extent is allocated, in turn, from each MDisk that is in the storage pool. For example, a storage pool that has 10 MDisks takes one extent from each MDisk. The 11th extent is taken from the first MDisk, and so on. This procedure, which is known as a round-robin, is similar to RAID-0 striping.

You can also supply a list of MDisks to use as the stripe set. This list can contain two or more MDisks from the storage pool. The round-robin procedure is used across the specified stripe set.

**Attention:** By default, striped volume copies are striped across all MDisks in the storage pool. If some of the MDisks are smaller than others, the extents on the smaller MDisks are used up before the larger MDisks run out of extents. Manually specifying the stripe set in this case might result in the volume copy not being created.

If you are unsure if sufficient available capacity is available to create a striped volume copy, select one of the following options:

- Check the available capacity on each MDisk in the storage pool by using the **`lsfreeextents`** command.
- Let the system automatically create the volume copy by not supplying a specific stripe set.

This figure shows an example of a storage pool that contains three MDisks. This figure also shows a striped volume copy that is created from the extents that are available in the storage pool.

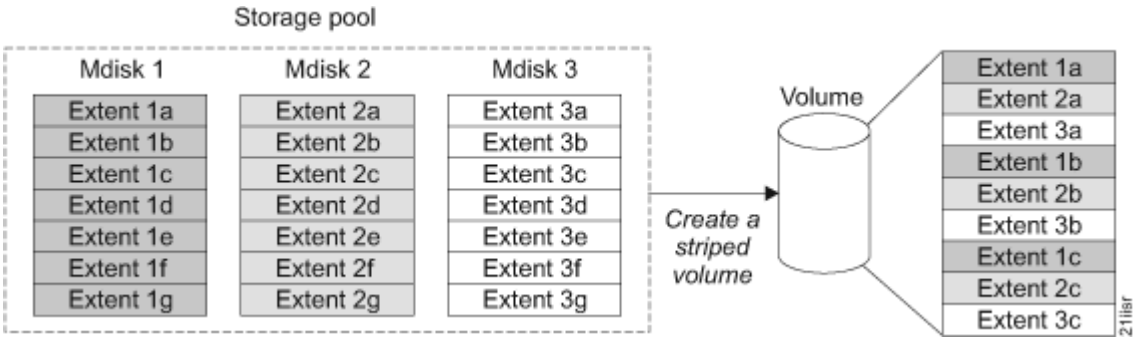


*Figure 17. Storage pools and volumes*

**Sequential**

When extents are selected, they are allocated sequentially on one MDisk to create the volume copy if enough consecutive free extents are available on the chosen MDisk.

**Image**

Image-mode volumes are special volumes that have a direct relationship with one MDisk. If you have an MDisk that contains data that you want to merge into the clustered system, you can create an image-mode volume. When you create an image-mode volume, a direct mapping is made between extents that are on the MDisk and extents that are on the volume. The MDisk is not virtualized. The logical block address (LBA) *x* on the MDisk is the same as LBA *x* on the volume.

When you create an image-mode volume copy, you must assign it to a storage pool. An image-mode volume copy must be at least one extent in size. The minimum size of an image-mode volume copy is the extent size of the storage pool to which it is assigned.

The extents are managed in the same way as other volume copies. When the extents are created, you can move the data onto other MDisks that are in the storage pool without losing access to the data. After you move one or more extents, the volume copy becomes a virtualized disk, and the mode of the MDisk changes from image to managed.

**Attention:** If you add a managed mode MDisk to a storage pool, any data on the MDisk is lost. Ensure that you create image-mode volumes from the MDisks that contain data before you start adding any MDisks to storage pools.

MDisks that contain existing data have an initial mode of unmanaged, and the clustered system cannot determine whether it contains partitions or data.

You can use more sophisticated extent allocation policies to create volume copies. When you create a striped volume, you can specify the same MDisk more than once in the list of MDisks that are used as the stripe set. This allocation is useful if you have a storage pool in which not all the MDisks are of the same capacity. For example, if you have a storage pool that has two 18 GB MDisks and two 36 GB MDisks, you can create a striped volume copy by specifying each of the 36 GB MDisks twice in the stripe set so that two-thirds of the storage are allocated from the 36 GB disks.

If you delete a volume, you destroy access to the data that is on the volume. The extents that were used in the volume are returned to the pool of free extents that is in the storage pool. The deletion might fail if the volume is still mapped to hosts. The deletion might also fail if the volume is still part of a FlashCopy, Metro Mirror, or Global Mirror mapping. If the deletion fails, you can specify the force-delete flag to delete both the volume and the associated mappings to hosts. Forcing the deletion deletes the Copy Services relationship and mappings.

## States

This table describes the different possible states of a volume.

| *Table 7. Volume states* | |
|---|---|
| **State** | **Description** |
| Online | At least one synchronized copy of the volume is online and available if both nodes in the I/O group can access the volume. A single node can access a volume only if it can access all the MDisks in the storage pool that are associated with the volume. |
| Offline | The volume is offline and unavailable if both nodes in the I/O group are missing, or if none of the nodes in the I/O group that are present can access any synchronized copy of the volume. The volume can also be offline if the volume is the secondary of a Metro Mirror or Global Mirror relationship that is not synchronized. A thin-provisioned volume goes offline if a user attempts to write an amount of data that exceeds the available disk space. |
| Degraded | The status of the volume is degraded if one node in the I/O group is online and the other node is either missing or cannot access any synchronized copy of the volume.<br>**Note:** If a volume is degraded and all of the associated nodes and MDisks are online, call your support center for assistance. |
| Deleting | For thin-provisioned or compressed volume copies in data reduction pools, the deleting status indicates that copies are being deleted. All volume copies, including fully allocated copies, are not accessible until the delete operation completes. In addition, several operations cannot be started until all copies are deleted. The following commands are restricted if one copy of a volume is in the process of being deleted:<br>• **expandvdisksize**<br>• **migratevdisk**<br>• **rmvdiskcopy**<br>• **rmvolumecopy**<br>• **shrinkvdisksize**<br>• **splitvdiskcopy** |

## Cache modes

You can select to have read and write operations that are stored in cache by specifying a cache mode. You can specify the cache mode when you create the volume. After the volume is created, you can change the cache mode.

This table describes the types of cache modes for a volume.

| *Table 8. Volume cache modes* | |
|---|---|
| **Cache mode** | **Description** |
| readwrite | All read and write I/O operations that are performed by the volume are stored in cache. This is the default cache mode for all volumes. A volume or volume copy created from a data reduction pool must have a cache mode of readwrite. If you try to create a thin provisioned or compressed volume copy from a data reduction pool and the volume cache mode is not `readwrite`, the operation fails. |
| readonly | All read I/O operations that are complete by the volume are stored in cache. |
| none | All read and write I/O operations that are complete by the volume are not stored in cache. |

## Basic volumes

A basic volume is the simplest type of volume, consisting of a copy in a single storage pool.

A basic volume is a volume that has only one physical copy. Basic volumes reside in a single pool on one site. In addition, basic volumes are supported in any system topology and are common to all configurations. A basic volume can be of any type of virtualization: striped, sequential, or image.

## Mirrored volumes

By using volume mirroring, a volume can have two copies. Each volume copy can belong to a different pool, and each copy has the same provisioned capacity as the volume. In the management GUI, an asterisk (*) indicates the primary copy of the mirrored volume. The primary copy indicates the preferred volume for read requests.

When a server writes to a mirrored volume, the system writes the data to both copies. When a server reads a mirrored volume, the system picks one of the copies to read. If one of the mirrored volume copies is temporarily unavailable; for example, because the storage system that provides the pool is unavailable, the volume remains accessible to servers. The system remembers which areas of the volume are written and resynchronizes these areas when both copies are available.

You can create a volume with one or two copies, and you can convert a non-mirrored volume into a mirrored volume by adding a copy. When a copy is added in this way, the system synchronizes the new copy so that it is the same as the existing volume. Servers can access the volume during this synchronization process.

You can convert a mirrored volume into a non-mirrored volume by deleting one copy or by splitting one copy to create a new non-mirrored volume.

The volume copy can be any type: image, striped, or sequential. The volume copy can use thin-provisioning or compression to save capacity. If the copies are located in data reduction pools, you can also use deduplication to the volume copies to increase the capacity savings. If you are creating a new volume, the two copies can be of different types, but to use deduplication, both copies must reside in a data reduction pool. You can add a deduplicated volume copy in a data reduction pool to an existing volume with a copy in a standard pool. You can use this method to migrate existing volume copies to data migration pools.

You can use mirrored volumes for the following reasons:

- Improving availability of volumes by protecting them from a single storage system failure.
- Providing concurrent maintenance of a storage system that does not natively support concurrent maintenance.
- Providing an alternative method of data migration with better availability characteristics. While a volume is migrated by using the data migration feature, it is vulnerable to failures on both the source and target pool. Volume mirroring provides an alternative because you can start with a non-mirrored volume

in the source pool, and then add a copy to that volume in the destination pool. When the volume is synchronized, you can delete the original copy that is in the source pool. During the synchronization process, the volume remains available even if there is a problem with the destination pool.

- Converting standard-provisioned volumes to use data reduction technologies, such as thin-provisioning, compression, or deduplication.
- Converting compressed or thin-provisioned volumes in standard pools to data reduction pools to improve capacity savings.

When you use volume mirroring, consider how quorum candidate disks are allocated. Volume mirroring maintains some state data on the quorum disks. If a quorum disk is not accessible and volume mirroring is unable to update the state information, a mirrored volume might need to be taken offline to maintain data integrity. To ensure the high availability of the system, ensure that multiple quorum candidate disks are allocated and configured on different storage systems.

When a volume mirror is synchronized, a mirrored copy can become unsynchronized if it goes offline and write I/O requests need to be processed, or if a mirror fast failover occurs. The fast failover isolates the host systems from temporarily slow-performing mirrored copies, which affect the system with a short interruption to redundancy.

**Note:** If the capacity is standard-provisioned, the primary volume formats before synchronizing to the volume copies. The **-syncrate** parameter on the **mkvdisk** command controls the format and synchronization speed.

## Write fast failovers

With write fast failovers, during processing of host write I/O, the system submits writes (with a timeout value of 10 seconds) to both copies. If one write succeeds and the other write takes longer than 10 seconds, the slower request times-out and ends. The duration of the ending sequence for the slow copy I/O depends on the backend from which the mirror copy is configured. For example, if the I/O occurs over the Fibre Channel network, the I/O ending sequence typically completes in 10 to 20 seconds. However, in rare cases, the sequence can take more than 20 seconds to complete. When the I/O ending sequence completes, the volume mirror configuration is updated to record that the slow copy is now no longer synchronized. When the configuration updates finish, the write I/O can be completed on the host system.

The volume mirror stops using the slow copy for 4 - 6 minutes; subsequent I/O requests are satisfied by the remaining synchronized copy. During this time, synchronization is suspended. Additionally, the volume's synchronization progress shows less than 100% and decreases if the volume receives more host writes. After the copy suspension completes, volume mirroring synchronization resumes and the slow copy starts synchronizing.

If another I/O request times out on the unsynchronized copy during the synchronization, volume mirroring again stops using that copy for 4 - 6 minutes. If a copy is always slow, volume mirroring attempts to synchronize the copy again every 4 - 6 minutes and another I/O timeout occurs. The copy is not used for another 4 - 6 minutes and becomes progressively unsynchronized. Synchronization progress gradually decreases as more regions of the volume are written.

If write fast failovers occur regularly, there can be an underlying performance problem within the storage system that is processing I/O data for the mirrored copy that became unsynchronized. If one copy is slow because of storage system performance, multiple copies on different volumes are affected. The copies might be configured from the storage pool that is associated with one or more storage systems. This situation indicates possible overloading or other back-end performance problems.

When you enter the **mkvdisk** command to create a new volume, the **mirror_write_priority** parameter is set to **latency** by default. Fast failover is enabled. However, fast failover can be controlled by changing the value of the **mirror_write_priority** parameter on the **chvdisk** command. If the **mirror_write_priority** is set to **redundancy**, fast failover is disabled. The system applies a full SCSI initiator-layer error recovery procedure (ERP) for all mirrored write I/O. If one copy is slow, the ERP can take up to 5 minutes. If the write operation is still unsuccessful, the copy is taken offline. Carefully consider whether maintaining redundancy or fast failover and host response time (at the expense of a temporary loss of redundancy) is more important.

**Attention:** Mirrored volumes can be taken offline if no quorum disk is available. This behavior occurs because synchronization status for mirrored volumes is recorded on the quorum disk. To protect against mirrored volumes being taken offline, follow the guidelines for setting up quorum disks.

### Read fast failovers

Read fast failovers affect how the system processes read I/O requests. A read fast failover determines which copy of a volume the system tries first for a read operation. The *primary-for-read copy* is the copy that the system tries first for read I/O; it is determined by a user implicated read algorithm.

The system submits host read I/O request to one copy of a volume at a time. If that request succeeds, then the system returns the data. If it is not successful, the system retries the request to the other copy volume.

With read fast failovers, when the primary-for-read copy goes slow for read I/O, the system fails over to the other copy. This means that the system tries the other copy first for read I/O during the following 4 - 6 minutes. After that, the system reverts to read the original primary-for-read copy. During this period, if read I/O to the other copy also goes slow, the system reverts immediately. Also, if the primary-for-read copy changes, the system reverts to try the new primary-for-read copy. This can happen when the system topology changes or when the primary or local copy changes. For example, in a standard topology, the system normally tries to read the primary copy first. If you change the volume's primary copy during a read fast failover period, the system reverts to read the newly set primary copy immediately.

The read fast failover function is always enabled on the system. During this process, the system does not suspend the volumes or make the copies out of sync.

### Maintaining data integrity of mirrored volumes during storage system maintenance

Volume mirroring improves data availability by allowing hosts to continue I/O to a volume even if one of the backend storage systems failed. However, this mirroring does not affect data integrity. If either of the backend storage systems corrupts the data, the host is at risk of reading that corrupted data in the same way as for any other volume. Therefore, before you perform maintenance on a storage system that might affect the data integrity of one copy, it is important to check that both volume copies are synchronized. Then, remove that volume copy before you begin the maintenance. For example, the scenario would apply if you need to zero the data on the disks that the storage system is providing.

## Custom volumes

Custom volumes create volumes that are based on user-defined customization rather than taking the standard default settings for each of the options under quick volume creation.

Under most circumstances, the preset volume types work well in creating new volumes quickly and efficiently by setting the standard defaults for options. However, you can customize settings to create volumes that are specific to your environment.

## Standard-provisioned volumes

A standard-provisioned volume completely uses storage at creation.

A volume's capacity can be considered in terms of its provisioned capacity and its real capacity. *Provisioned capacity* is the volume storage capacity that is available to a host. *Real capacity* is the storage capacity that is allocated to a volume from a pool. *Standard-provisioned volumes* are created with the real capacity equal to the provisioned capacity.

The extra metadata that describes the contents of thin-provisioned volumes is not required for standard-provisioned volumes. As a result, the I/O rates that are obtained from standard-provisioned volumes can be higher than the rates obtained from thin-provisioned volumes that are provisioned on the same MDisks.

By default, standard-provisioned volumes are automatically formatted by a background process after the volume is created. During this process, some system resources are used to make the volume available for immediate use and ensures the volume is clear of any prior data. The time taken to format is governed by the size of the volume and the mirror synchronization rate of the volume. Increasing the synchronization rate reduces the time taken to format the volume by using additional system resources. The progress and estimated completion time of each volume format process is reported by the system.

Each formatting process uses some system resources, so there is a limit on the number of volumes that can be formatted at the same time. In addition, some volume actions such as adding a volume copy, or moving, expanding, or shrinking the volume, are disabled until the system completes formatting the volume.

If it is not necessary for a standard-provisioned volume to be cleared of data when it is created, the formatting process can be skipped. Skipping formatting can also be useful if the volume will be the target of a Copy Services function because the Copy Services operation formats the volume.

## Thin-provisioned volumes

A *thin-provisioned* volume presents a different capacity to mapped hosts than the capacity that the volume consumes in the storage pool.

The system supports thin-provisioned volumes in standard pools and in data reduction pools. In standard pools, thin-provisioned volumes are created as a specific volume type, that is based on capacity savings criteria. These properties are managed at the volume level. With data reduction pools, all the benefits of thin-provisioning are available to all the volumes that are assigned to the pool. Only standard-provisioned volumes do not gain these benefits. Data reduction pools enhance capacity efficiency for thin-provisioned volumes by monitoring the hosts use of capacity. When the host indicates that the capacity is no longer needed, the capacity is released and can be reclaimed by the data reduction pool to be redistributed automatically. Standard pools do not have the ability to reclaim capacity. For added capacity savings, you can specify compression and deduplication on volumes in data reduction pools.

Thin-provisioned volumes can also help simplify server administration. Instead of assigning a volume with some capacity to an application and increasing that capacity as the needs of the application change, you can configure a volume with a large virtual capacity for the application. You can then increase or shrink the real capacity as the application needs change, without disrupting the application or server.

*Provisioned capacity* is the volume storage capacity that is available to a host. *Real capacity* is the storage capacity that is reserved to a volume copy from a pool. In a standard-provisioned volume, the provisioned capacity and real capacity are the same. However, in a thin-provisioned volume, the provisioned capacity can be much larger than the real capacity.

The provisioned capacity of a thin-provisioned volume is typically significantly larger than its real capacity. Each system uses the real capacity to store data that is written to the volume, and metadata that describes the thin-provisioned configuration of the volume. As more information is written to the volume, more of the real capacity is used. The system identifies read operations to unwritten parts of the virtual capacity and returns zeros to the server without using any real capacity.

The system must maintain extra metadata that describes the contents of thin-provisioned volumes. As a result, the I/O rates that are obtained from thin-provisioned volumes can be lower than the rates obtained from standard-provisioned volumes that are allocated on the same MDisks.

When you configure a thin-provisioned volume in a standard pool, you can use the warning level attribute to generate a warning event when the used real capacity exceeds a specified amount or percentage of the total provisioned capacity. You can also use the warning event to trigger other actions, such as taking low-priority applications offline or migrating data into other storage pools. For thin-provisioned volumes in data reduction pools, the warning level value cannot be set because capacity reporting is handled at the pool level.

If a thin-provisioned volume in a standard pool does not have enough real capacity for a write operation, the volume is taken offline and an error is logged (error code 1865, event ID 060001). Access to the thin-provisioned volume is restored by either increasing the real capacity of the volume or increasing the size of the storage pool that it is allocated on.

When you create a thin-provisioned volume in standard pools, you can choose the grain size for allocating space in 32 KB, 64 KB, 128 KB, or 256 KB chunks. The grain size that you select affects the maximum provisioned capacity for the thin-provisioned volume in standard pools. The default grain size is 256 KB. If you select 32 KB for the grain size, the volume size cannot exceed 260,000 GB. The grain size cannot be changed after the thin-provisioned volume is created in a standard pool. Generally, smaller grain sizes save space but require more metadata access, which can adversely impact performance. If you are not going to use the thin-provisioned volume as a FlashCopy source or target volume, use 256 KB to maximize performance. If you are going to use the thin-provisioned volume as a FlashCopy source or target volume, specify the same grain size for the volume and for the FlashCopy function. Grain size cannot be set on thin-provisioned volume copies in data reduction pools. The grain size of 8 KB is the default size for thin-provisioned volume copies in data reduction pools.

When you create a thin-provisioned volume, set the cache mode to `readwrite` to maximize performance. If the cache mode is set to `none`, the system cannot cache the thin-provisioned metadata, which decreases performance. A volume or volume copy created from a data reduction pool must have a cache mode of readwrite. If you try to create a thin provisioned or compressed volume copy from a data reduction pool and the volume cache mode is not `readwrite`, the operation fails.

The autoexpand feature prevents a thin-provisioned volume from using up its capacity and going offline. As a thin-provisioned volume uses capacity, the autoexpand feature maintains a fixed amount of unused real capacity, called the *contingency capacity*. For thin-provisioned volumes in data reduction pools, the autoexpand feature is always enabled to maintain contingency capacity. For thin-provisioned volumes in standard pools, the autoexpand feature is optional. However without this feature enabled, the contingency capacity can get used up, causing the volume to go offline. If you are using standard pools and want to determine whether an application requires the autoexpand feature, you can create a test thin-provisioned volume with the autoexpand feature turned off. If the application causes the volume to run out of capacity and go offline, you can then create a vdisk with the autoexpand feature turned on.

## Compressed volumes

When you create volumes, you can specify compression as a method to save capacity for the volume. With compressed volumes, data is compressed as it is written to disk, saving more space. When data is read to hosts, the data is decompressed.

Compression is available through data reduction support as part of the system. If you want volumes to use compression as part of data reduction support, compressed volumes must belong to data reduction pools. Data reduction pools also support reclaiming unused capacity automatically after mapped hosts no longer need the capacity for operations. These hosts issue SCSI unmap commands and the released capacity is reclaimed by the data reduction pool for redistribution. For compressed volumes in data reduction pools, the used capacity before compression indicates the total amount of data that is written to volume copies in the storage pool before data reduction occurs.

You can also monitor information on compression usage to determine the savings to your storage capacity when volumes are compressed. To monitor system-wide compression savings and capacity, select **Dashboard** to view the **Capacity Saving**. You can compare the amount of capacity that is used before compression is applied to the capacity that is used for all compressed volumes. In addition, you can view the total percentage of capacity savings when compression is used on the system. For systems that use multiple compression technologies, such as data reduction pools or self-compressing drives, the total percentage of capacity savings cannot be determined at a system level. To view each of the capacity savings for all the supported compression technologies for the system, select **View Compression Details**. You can also monitor compression savings across individual pools. You can also monitor compression savings across individual pools.

If your system currently does not use compression, the system automatically analyzes your configuration to determine the potential storage savings if compression is used. The management GUI incorporates the Comprestimator utility that uses mathematical and statistical algorithms to create potential compression savings for the system. The analysis for potential savings can be used to determine whether purchasing a compression license for the system is necessary to reduce cost of extra storage devices. To estimate compression savings in the management GUI, select **Volumes** > **Actions** > **Space Savings** > **Estimate Compression Savings**.

After the analysis completes, you can download a savings report that shows estimated savings for all the volumes with enough data to be analyzed. This report lists all currently configured volumes on the system and their potential compressions savings. To download a report, select **Volumes** > **Volumes** > **Actions** > **Capacity Savings** > **Download Savings Report**.

## Deduplicated volumes

Deduplication can be configured with volumes that use different capacity saving methods, such as thin-provisioning. Deduplicated volumes must be created in data reduction pools for added capacity savings. Deduplication is a type of data reduction that eliminates duplicate copies of data. Deduplication of user data occurs within a data reduction pool and only between volumes or volume copies that are marked as deduplicated. Some models or software versions require specific hardware or software to use this function.

With deduplication, the system identifies unique chunks of data, called *signatures*, to determine whether new data is written to the storage. Deduplication is a hash-based solution, which means chunks of data are compared to their signatures rather than to the data itself. If the signature of the new data matches an existing signature that is stored on the system, then the new data is replaced with a reference. The reference points to the stored data, instead of writing the data to storage. This process saves capacity on the backend storage by not writing new data to storage and might improve performance on read operations to data with an existing signature. The same data pattern can occur many times and deduplication decreases the amount of data that needs to be stored on the system. A part of every hash-based deduplication solution is a repository that supports looking up matches for incoming data. The system contains a database that maps the signature of the data to the volume and its virtual address. If an incoming write operation does not have a signature that is stored in the database, then a duplicate is not detected and the incoming data is stored on backend storage. To maximize the space that is available for the database, the system distributes this repository between all nodes in the I/O groups that contain deduplicated volumes. Each node carries a distinct portion of the records that are stored in the database. If nodes are removed or added to the system, the database is redistributed between the nodes to ensure full use of available memory. Only certain models with specific hardware support deduplication. Verify your model and hardware components to use these functions.

When you create a volume, you can specify to include deduplication with other supported capacity savings methods. Deduplicated volumes must be created in data reduction pools. If you have existing volumes in standard pools, you can migrate them to data reduction pools to add deduplication to increase capacity savings for the volume.

## Volume protection

Volume protection prevents active volumes or host mappings from being deleted inadvertently if the system detects recent I/O activity. This global setting is enabled by default on new systems. You can either set this value to apply to all volumes that are configured on your system, or control whether the system-level volume protection is enabled or disabled on specific pools.

To prevent an active volume from being deleted unintentionally, administrators can use the system-wide setting to enable volume protection. They can also specify a time period that the volume must be idle before it can be deleted.

The system-wide volume protection and the pool-level protection must both be enabled for protection to be active on a pool. The pool-level protection depends on the system-level setting to ensure that protection is applied consistently for volumes within that pool. If system-level protection is enabled, but pool-level protection is not enabled, any volumes in the pool can be deleted even when the setting is configured at the system level.

The following commands are affected by this setting:

- **`rmvdisk`**
- **`rmvdiskcopy`**
- **`rmvvolume`**

- **rmvdiskhostmap**
- **rmvolumehostclustermap**
- **rmmdiskgrp**
- **rmhostiogrp**
- **rmhost**
- **rmhostcluster**
- **rmhostport**
- **mkrcrelationship**

## Volume groups

A *volume group* is a container for managing a set of related volumes as a single object. The volume group provides consistency across all volumes in the group.

Volume groups can be used with the following functions:

**Safeguarded Copy**
One implementation of volume groups is to group volumes to be configured as Safeguarded. Safeguarded copy function is a cyber-resiliency feature that creates immutable copies of data that cannot be changed or manipulated.

A Safeguarded volume group describes a set of source volumes that can span different pools and are backed up collectively with the Safeguarded Copy function. Safeguarded snapshots are supported on the system through an internal scheduler that is defined in the snapshot policy or can be configured with an external snapshot scheduling application such as IBM Copy Services Manager.

**Policy-based High Availability replication**

One implementation of volume groups is to add the volume group to a storage partition that is associated with a high availability replication policy. Volumes in that volume group, hosts, and host-to-volume mappings contained in the partition are automatically configured across both systems associated with the replication policy. For more information, see .

**Asynchronous policy-based replication**

Asynchronous policy-based replication is configured on all volumes in a volume group by assigning an asynchronous replication policy to that volume group. The system automatically replicates the data and configuration for volumes in the group based on the values and settings in the replication policy. As part of asynchronous replication, a recovery volume group is created automatically on the recovery system. Recovery volume groups cannot be created, changed, or deleted. A single replication policy can be assigned to multiple volume groups to simplify replication management. When additional volumes are added to the group, replication is automatically configured for these new volumes.

**Snapshot**
*Snapshots* are the read only point-in-time copies of a volume group that cannot be directly accessible from the hosts. To access the snapshot contents, you can create a clone or thin clone of a volume group snapshot. You can use the management GUI to configure volume groups to use snapshot policies for multiple volumes for consistent management. Safeguarded snapshot with internal scheduler can be created by using snapshot function.

The volumes in a volume group are supposed to be mutually consistent. This means that volume group only make sense as a group. When a group of thin-clone or clone is populated, it is snapshot function's responsibility to ensure that the images are mutually consistent. When volumes are added or removed from a group, the host applications ensures the volume groups are mutually consistent.

# Host attachment

This section covers information about system host connectivity.

You can attach hosts to the system by using the following methods:

- Fibre Channel
- NVMe over Fibre Channel
- Fibre Channel over Ethernet (FCoE)
- NVMe over RDMA
- NVMe over TCP
- Serial-attached SCSI (SAS)
- Ethernet iSCSI and iSER

For more information about the supported host configurations on each product, see IBM System Storage Interoperation Center (SSIC).

Refer to the subtopics of this section for information about setting host attachment parameters.

# NVMe over RDMA and NVMe over TCP host attachments

You can attach NVM Express (NVMe) over RDMA or NVMe over TCP hosts to the system.

For more information about NVMe over RDMA and NVMe over TCP, such as interoperability requirements, see IBM System Storage Interoperation Center (SSIC).

## Linux

You can attach NVM Express (NVMe) over RDMA or NVMe over TCP hosts that run the Linux operating system.

### Attachment requirements for hosts that are running Linux

Ensure that the host is running a supported operating system and adapter cards. Refer to vendor documentation if updates are needed.

**Note:** For supported levels, see IBM System Storage Interoperation Center (SSIC).

NVM-Express user space tools for Linux are required. Install these tools by using the following commands:

- SLES: **zypper install nvme-cli**
- RHEL: **yum install nvme-cli**

Use the following version of the nvme-cli tool:

- On SLES, only nvme-cli version 1.13 is supported.
- On RHEL 8.4, use nvme-cli version 1.12.
- On RHEL 8.6 or later, use nvme-cli version 1.16.

### Configuring the host attachment

1. For connectivity and performance benefits, the following sysctl settings on host are recommended:

   net.ipv4.conf.all.arp_filter is enabled
   net.ipv4.conf.all.arp_ignore is set to 2
   TCP Delayed ACK is enabled

   Refer to Operating System documentation for configuration instructions.
2. Use the following command to identify the host NVMe Qualified Name (NQN) address:

   ```
   cat /etc/nvme/hostnqn
   ```
3. On the storage system, create the NVMe host object using the **Hosts** panel on the GUI, or the **mkhost** command on the CLI.

4. On the storage system, map volumes to the NVMe host using the **Volumes** panel on the GUI, or the **mkvdiskhostmap** command on the CLI.
5. To discover and connect to NVMe targets on the host, use the **nvme discover** and **nvme connect** commands. Refer to your operating system documentation for more information.

## VMware ESXi

You can attach NVM Express (NVMe) over RDMA or NVMe over TCP hosts that run the VMware ESXi operating system.

VMware ESXi contains two levels of connectivity. The first level of VMware ESXi is the host (ESXi machine), which is the actual server that includes the host bus adapter and the NVMe support. The host manages both hardware resource and virtualized components by using the GUI and CLI (esxcli).

The other level includes one or more guest virtual machines with an operating system. The virtual machines detect the defined volumes as VMware data stores and can read or write to them.

### Configuring the VMware ESXi

Ensure that the host is running a supported operating system and adapter cards. Refer to vendor documentation if updates are needed.

**Note:** For supported levels, see IBM System Storage Interoperation Center (SSIC).

*Table 9. NVMe over RDMA minimum supported Ethernet cards and firmware versions for VMware ESXi 7.0 UP2*

| | **CX4** | **CX5** | **CX6** | **N2100G** |
|---|---|---|---|---|
| **Firmware version** | 14.26.1040 | 16.32.1010 | 14.26.1040 | 219.0.144.0 |
| **Ethernet Card version** | 4.21.71.101-1OEM.702.0.0.17630552 | 4.21.71.101-1OEM.702.0.0.17630552 | 4.21.71.101-1OEM.702.0.0.17630552 | 219.0.29.0-1OEM.700.1.0.15843807<br>219.0.13.0-1OEM.700.1.0.15843807 |

For more information about Ethernet cards and latest firmware levels, see ConnectX Ethernet Cards for VMware ESXi Server and Broadcom Ethernet Network Adapters.

Configure the VMware ESXi 7.0 or ESXi 8.0 operating system before connecting hosts that run VMware ESXi data stores with the system.

**Configure the adapter software**

Before configuring the host operating system, the following tasks must be completed:

1. If you're using RDMA or TCP with this ESXi for the first time, enable the adapter software from the VMware ESXi GUI.
   a. Verify that Adapter MTUs are consistently set end-to-end.
   b. Validate that RDMA vSphere Installation Bundles (VIBs) are loaded with the correct vSphere version (OS dependent).
   c. Create a standard virtual switch per RDMA or TCP adapter.
      - For more information about configuring RDMA adapters, see **Configure Adapters for NVMe over RDMA (RoCE v2)** and **Configure VMkernel Binding with a vSphere Standard Switch** in VMware vSphere Storage product documentation.
      - For more information about configuring TCP adapters, see **Configure Adapters for NVMe over TCP Adapter Storage** and **Configure VMkernel Binding for a TCP Adapter with a vSphere Standard Switch** in VMware vSphere Storage product documentation.
   d. It is recommended that delayed ack be enabled. Refer to VMware ESXi vendor guides for more details.

**Configure the host software**

After the prerequisite tasks are complete, use the following general steps to configure your host.

1. Define the host object by using the NQN identifiers for each physical ESXi host in the VMware vSphere cluster.

   Find the NQN using the following command:

   ```
   esxcli nvme info get
   ```

2. Discover the vmhba value to identify each RDMA or TCP adapter.

   ```
   esxcfg-scsidevs -a
   ```

3. Use the **`nvme fabrics connect`** command.

   This command both discovers and connects within the same command.

   - Identify the discovered controller through command:

     ```
     esxcli nvme controller list
     ```

   - To discover and connect use command:

     ```
     esxcli nvme fabrics connect -i <ip_address_of_target_rdma/tcp> -a <vmhba> -c -t 1800
     ```

4. Run the following command to validate that the multipath and device mapper has loaded successfully:

   ```
   esxcli storage core plugin list
   ```

   Validating with High Performance Path (HPP) plug-in for NVMe:

   ```
   esxcli storage core claimrule list -c MP
   Rule Class   Rule  Class    Type       Plugin
   Matches                                                                  XCOPY Use Array
   Reported Values  XCOPY Use Multiple Segments  XCOPY Max Transfer Size KiB  Config String
   ----------  -----  -------  ---------  ---------
   ------------------------------------------------------------------------
   ------------------------------  --------------------------
   --------------------------  -------------
   MP             50  runtime  transport  NMP
   transport=usb
          false                          false                            0
   MP             51  runtime  transport  HPP
   transport=sata
          false                          false                            0  pss=FIXED
   MP             52  runtime  transport  NMP
   transport=ide
          false                          false                            0
   MP             53  runtime  transport  NMP
   transport=block
          false                          false                            0
   MP             54  runtime  transport  NMP
   transport=unknown
          false                          false                            0
   MP            101  runtime  vendor     MASK_PATH  vendor=DELL model=Universal Xport
   pci_vendor_id=* pci_sub_vendor_id=*
   false                         false                                 0
   MP            101  file     vendor     MASK_PATH  vendor=DELL model=Universal Xport
   pci_vendor_id=* pci_sub_vendor_id=*
   false                         false                                 0
   MP          65531  runtime  transport  HPP
   transport=sas
          false                          false                            0  pss=FIXED
   MP          65532  runtime  transport  HPP
   transport=parallel
          false                          false                            0  pss=FIXED
   MP          65533  runtime  transport  HPP
   transport=pcie
          false                          false                            0  pss=FIXED
   MP          65534  runtime  vendor     HPP        nvme_controller_model=*
   pci_vendor_id=* pci_sub_vendor_id=*
   false                         false                                 0
   ```

```
MP          65535  runtime  vendor     NMP        vendor=* model=* pci_vendor_id=*
pci_sub_vendor_id=*
false                        false                              0
```

**Configure the attachment**

To validate that the namespaces are retrieved, complete the following steps:

1. Run the following command:

   ```
   esxcli nvme namespace list
   ```

   This provides an empty output.
2. Map all volumes to the host by using the **mkvdiskhostmap** command or the management GUI. For more information about **mkvdiskhostmap** command, see **mkvdiskhostmap**.
3. Rerun the **esxcli nvme namespace list** command.

## VMware Virtual Volumes (vVols)

The system provides native support for VMware vSphere APIs for Storage Awareness (VASA) through a VASA Provider (also known as a Storage Provider), which sends and receives information about storage that is used by VMware vSphere to the vCenter Server. Through VASA, the system also supports VMware Virtual Volumes (also known as *vVols*), which allows VMware vCenter to automate the creation, deletion and mapping of volumes.

Before vSphere can be used to provision vVols, the storage system administrator must first enable vVols on the storage system. It is recommended that vVols are enabled by using the management GUI as this simplifies the setup process.

When vVols are enabled, the following objects are configured on the system:

- An ownership group
- A user account and user group
- A metadata volume
- A child pool with an associated provisioning policy

An ownership group is created that separates the resources that belong to VMware from the other resources in the system, and ensures that only these resources can be managed by VMware vCenter.

The storage system administrator delegates ownership of Virtual Volumes to VMware vCenter and the VASA Provider by creating a user account with the VASA Provider role. Although the storage system administrator can complete certain actions on volumes and pools that are owned by the VASA Provider, the VMware environment retains management responsibility for Virtual Volumes.

A metadata volume is created to store metadata for Virtual Volumes and vSphere storage policies. The system administrator selects a storage pool to provide capacity for the metadata volume. With each new volume created by the VASA Provider, VMware vCenter defines a small amount of metadata that is stored on the metadata volume. The metadata volume is exclusively used by the VASA Provider and cannot be deleted while vVols exist or mapped to hosts.

The storage system administrator decides what storage to allocate for Virtual Volumes by creating a child pool and assigning it to the VASA ownership group. Each vVols child pool is presented as a storage container in vSphere, from which a vVols datastore can be created. The storage system administrator must also associate a provisioning policy with the child pool, to specify how data Virtual Volumes, known as vmdk Virtual Volumes, are provisioned.

Storage Virtualize supports asynchronous replication for Virtual Volumes between two storage systems in different locations. The storage administrator preconfigures replication groups for vSphere by creating volume groups for vVols and assigning a replication policy to each group. The VMware administrator can use storage policies in vCenter to define the replication requirements for virtual machines, such as the second site where the VMs should be replicated and a recovery point objective. Virtual Volumes for multiple virtual machines can be grouped to manage them as a single unit for failover and disaster recovery.

## Planning vVols

VMware Virtual Volumes (vVols) are not supported on all Storage Virtualize systems. Therefore, plan how to provision your vVols.

VMware Virtual Volumes (vVols) are not supported on all Storage Virtualize systems. The following systems do not support vVols:

- IBM Storage Virtualize for Public Cloud
- IBM Storage FlashSystem 5015
- IBM Storage FlashSystem 5035
- IBM Storage FlashSystem 5045
- IBM Storage FlashSystem 5200 with less than 128GiB memory in each node canister (256 GiB per control enclosure).

The system must be configured with the standard topology. Using stretched or HyperSwap topology systems is not supported.

The vCenter Server and ESXi host management networks require TCP port 8440 to be accessible on the storage system's management IP address.

The storage system, vCenter Server and ESXi hosts must all be configured to use a Network Time Protocol (NTP) service. This ensures that the time settings between the storage system and the VMware environment remain consistent.

At least one standard storage pool is required. Data reduction pools are not supported for storing vVols or the vVols metadata volume.

Do not run vCenter Servers on vVols. This avoids the possibility of complicated failure scenarios with cyclic dependancies in the VMware environment.

## Planning vVols replication

To implement vVols replication between two systems, each system must support vVols and asynchronous policy-based replication.

The following restrictions apply when you use vVols replication:

- A limit of one I/O group can be configured in each system.
- A limit of one partnership can be configured for vVols replication. Additional partnerships can be configured in each system but only a single partnership can be used for vVols replication.
- The size of a virtual volume cannot be changed while it is part of a vSphere replication group.

For more information, see "Policy-based replication: Asynchronous" on page 65.

Each storage system is a separate fault domain for vVols replication. The storage systems must be sufficiently far apart such that anticipated disruptions affect only one location.

Use vCenter Server and ESXi versions that support vVols replication. vCenter Server and ESXi can discover replication groups, but do not manage their lifecycle. When you use vVols replication, the failover and disaster recovery operations are managed through vCenter Server integrations such as VMware Site Recovery Manager (SRM) and PowerCLI. A vCenter Server is required in each location.

Storage Virtualize supports the use of preconfigured replication groups only. The automatic creation of replication groups by using vSphere is not supported.

When you use replication with virtual volumes, you can apply a replication storage policy only to a configuration virtual volume and a data virtual volume. The swap virtual volume and any memory vVols are excluded from replication. Virtual machine snapshots are also excluded from replication.

When planning the VMware environment, if you intend to create virtual machine snapshots, you must check that the vCenter Server integrations you intend to use to manage vVol replication support replication without snapshots.

# Implementing vVols

You can use the management GUI to enable Virtual Volumes.

You can use the management GUI to enable Virtual Volumes. Navigate to **Settings** > **System** > **VMware Virtual Volumes (vVols)** to get started.

Before enabling vVols, ensure that the system certificate contains one of the following in the Subject Alternative Name field:

- IP address
- Fully-qualified domain name (FQDN)
- Hostname

To specify a hostname or FQDN, a DNS server must be configured.

The properties defined in the Subject Alternative Name field are used by the vCenter Server to establish a secure connection to the storage system. The IP address, hostname or FQDN as specified in the certificate must be used to register the storage provider in the vCenter Server.

For more information, see System Certificates.

The storage system, vCenter Server and ESXi hosts must all be configured to use a Network Time Protocol (NTP) service. This ensures that the time settings between the storage system and the VMware environment are synchronized.

At least one standard storage pool is required. Data reduction pools are not supported for storing vVols or the vVols metadata volume.

When enabling vVols, the user account created for the vCenter Server is initially configured with a username and password. These credentials are only used when first registering the Storage Provider in vCenter. After successful registration, the password is removed and the user account is automatically reconfigured with a certificate for authentication.

After enabling vVols on the storage system, complete the remaining steps required in VMware vCenter to finish the configuration. If using an internally or externally signed certificate, ensure that the root or external CA certificates have been imported into the vCenter trust store and pushed to the ESXi hosts before registering the storage provider.

Configure a host object for each ESXi host that will access storage from the system. Protocol Endpoints (PEs) are automatically configured by setting the host type to ‘vVol’. When using clustered ESXi configurations, create individual host objects for each ESXi server, and use a host cluster object to group the hosts and simplify storage management.

Each vVols child pool is presented as a storage container in vSphere, on which a vVols datastore can be created. The initial setup creates a single child pool for vVols and additional pools can be created later if required.

The management GUI indicates when vVols are enabled on the storage system.

# Implementing vVols Replication

After you enable vVols, you can implement vVols replication.

Implement vVols replication by completing the following steps:

1. Enable vVols on each system and register the storage system in each location in the corresponding vCenter Server. See Implementing vVols.
2. Create a 2-site partnership using policy-based replication between the two storage systems.
3. Enable vVols replication on the system that will be used for creating production virtual machines.

## Creating a partnership for policy-based replication

Use the management GUI to create a 2-site partnership between systems using policy-based replication.

- Navigate to **Copy Services** > **Partnerships and Remote Copy** to get started.
- If a partnership using policy-based replication already exists between the two systems, no additional partnership configuration is needed to use vVols replication.
- The GUI indicates when the partnership is configured and ready for use with policy-based replication.

### Enabling vVols replication

Use the management GUI to enable vVols replication.

- Navigate to **Settings** > **System** > **VMware Virtual Volumes (vVols)** to begin the setup.
- Configuring vVols replication links the VASA ownership groups and child pools between the two systems in order for replicated Virtual Volumes to be manageable using the vCenter Server in the second site. The initial setup also creates a volume group with a replication policy. Additional volume groups and replication policies can be created later on either of the systems.
- After enabling vVols replication on the storage system, rescan the Storage Provider in vSphere to discover the new replication capabilities. If the storage system has already been registered with vSphere on an earlier version, an additional step may be required to discover the new replication capabilities. See Updating from version 8.6.1 or earlier, below.
- You can now create VM storage policies that define the replication requirements for virtual machines.
- When a new virtual machine is created using the policy, each of the compatible preconfigured volume groups on the storage system is available as a replication group choice in vSphere.
- Storage Virtualize supports the use of preconfigured replication groups only. The automatic creation of replication groups using vSphere is not supported.

### Updating from version 8.6.1 or earlier

If a system with vVols configured on version 8.6.1 or earlier is updated to version 8.6.2 or later, an additional step is required to discover the new capabilities of the storage system in the vSphere environment.

After the storage system software has been updated to version 8.6.2 or later, perform the following steps:

1. Using the vSphere Client, unregister the storage provider on the vCenter Server:
   a. Navigate to **VCSA** > **Configure** > **Storage Providers**.
   b. Select the storage provider and select **Remove**.
2. Using the storage system command-line interface, reset the user account used by the vCenter server:
   a. Use the following command to remove the stored certificate and specify a temporary password for the VASA user account: `chuser -nocertuid -password <new_password> <vCenter_user>`
3. Using the vSphere Client, reregister the storage provider on the vCenter Server:
   a. Navigate to **VCSA** > **Configure** > **Storage Providers**.
   b. Select **Add** and read the storage provider.

## Managing vVols

Create vVols child pools, resize the storage capacity of a vVols child pool, register the storage provider in vCenter, and perform other tasks to manage vVols.

### Creating a vVols child pool

Additional child pools can be created for vVols by using the management GUI or the command-line interface. Each vVols child pool is presented as a storage container in vSphere, on which a vVols data store can be created.

To use the GUI to create a child pool for vVols, specify the VASA ownership group and provisioning policy when you create the child pool.

To use the command-line interface, use the following command:

```
mkmdiskgrp -name <vVol_child_pool_name> -owner vvol_child_pool -ownershipgroup VASA
-parentmdiskgrp <parent_pool> -provisioningpolicy <policy> -size <size> -unit <unit>
```

The parent pool must be a standard storage pool. Data reduction pools are not supported for vVols.

If the new child pool will be used with vVols replication, add a pool link between the new pools on each system. Use the **Pools** window on the management GUI to add pool links for replication from either system.

After creating a new vVol child pool, the vCenter Server administrator must rescan the Storage Provider in vSphere to refresh the available storage containers before creating a new vVols data store.

## Resizing a vVols child pool

The storage administrator can increase or decrease the amount of storage capacity that is allocated to a vVols child pool by using the management GUI or the command-line interface.

After changing the capacity of a child pool, the vCenter Server administrator must refresh the datastore capacity information in vSphere to detect the new capacity.

## Considerations before updating the system software

The system operates with limited functionality during the software update process. During this time, no configuration changes can be made to the system and vVols functions are reduced.

In vSphere, consider the following:

- Any virtual machines that are powered on continue to function.
- The Storage Provider may appear as offline.
- The vVols datastores may appear as inaccessible.
- Any powered off virtual machines on vVols datastores appear as inaccessible.
- Any attempt to perform VM-level management operations fail. These include, for example: power on and off, snapshots, vMotion, Storage vMotion, and cloning.

## Re-registering the Storage Provider in vCenter

If the Storage Provider needs to be removed and re-registered in vCenter, use the following command to remove the stored certificate and specify a temporary password for the VASA user account.

```
chuser -nocertuid -password <new_password> <VASA user ID>
```

These credentials are only used when registering the Storage Provider in vCenter. After successful registration, the password is removed and the user account is automatically reconfigured with a certificate for authentication.

## Changing the system management IP address or DNS name

If the IP address, hostname or fully-qualified domain name (FQDN) specified in the Subject Alternative Name field of the system certificate is changed, a new system certificate must be generated using either the internal or external certificate authority. After generating the certificate, the storage Provider must be reregistered in vCenter (see Reregistering the Storage Provider in vCenter).

### User accounts

The user account created when enabling vVols is initially configured with a username and password. These credentials are only used when first registering the Storage Provider in vCenter. After successful registration, the password is removed and the user account is instead, automatically reconfigured with a certificate.

If any manual vVols-related service or recovery tasks are required on the storage system, the storage administrator must log in using a user account with the VASA Provider role. This can either be a new user in the same user group, or the same user account that the storage provider was registered with.

To create a new user account, use the following command:

```
mkuser –usergrp VASAProvider –name <new user> -password <password>
```

To configure a password for the existing user account, use the following command:

```
chuser -password <password> <vvol_user_name>
```

## Managing vVols Replication

Create and remove a replication group, modify the existing storage policy or assign a different policy in vSphere, and manage failover operations.

### Creating a Replication Group

Storage Virtualize supports the use of preconfigured replication groups only. The automatic creation of replication groups using vSphere is not supported.

To create a replication group for use with vSphere Virtual Volumes, you must first create a volume group for vVols and assign a replication policy on the storage system. To create a new volume group for vVols replication using the management GUI, go to the **Volumes** > **Volume groups** page.

To use the command-line interace, use the following command:

```
mkvolumegroup -owner vvol -ownershipgroup VASA -replicationpolicy <policy> -name <name>
```

After creating volume groups for vVols, rescan the Storage Provider in vSphere to refresh the available replication groups.

### Removing a Replication Group

When a replication group is no longer required in vSphere, the corresponding volume group can be removed on the storage system. The volume group must be empty before it can be removed. When the volume group is removed on the production system, the volume group is automatically removed on the recovery system.

After removing volume groups for vVols, rescan the Storage Provider in vSphere to refresh the available replication groups.

### Assigning or Changing the Replication Group for a Virtual Machine using vSphere

If the storage requirements for a virtual machine change, you can modify the existing storage policy or assign a different policy in vSphere. If the new storage policy will use a different replication group, the existing storage policy must first be removed from the virtual machine, or replaced with a policy that does not include replication, before the new policy can be assigned.

When editing the settings for a virtual machine with a storage policy that includes replication, the replication group selection is mandatory. If the replication group drop-down is initially not displayed in the vSphere Client GUI, you must wait for the GUI to update to select a replication group before applying the changes to the VM.

### Changing the size of a Virtual Volume using vSphere

The size of a virtual volume cannot be changed if the volume is in a replication group. You must remove the storage policy from the virtual machine, or assign a different policy that does not include replication, before resizing the VMDK for a virtual machine using vSphere. After applying the new size, the original storage policy can be reassigned to the virtual machine.

### Managing Failover Operations

When using vVols replication, the failover and disaster recovery operations are managed through the vCenter Server integrations such as VMware Site Recovery Manager (SRM) and PowerCLI.

Volume groups used for vVols replication can be monitored using the storage system GUI and command-line interface. However, the management of the volume groups is delegatated to the VMware environment.

vCenter Server and ESXi can discover replication groups, but do not manage their life cycle. When using vVols replication, the failover and disaster recovery operations are managed through vCenter Server integrations such as VMware Site Recovery Manager (SRM) and PowerCLI.

VMware can perform three types of failovers on vVol-based VMs:

**Planned Failover**
Movement of a VM from one site to another for a planned migration. The systems at both sites are required to be accessible throughout the failover process. Once a planned failover is complete, replication can optionally be reversed so that the failed over VM can be reprotected.

**Unplanned Failover**
Movement of a VM in response to a failure in the production site. Only the recovery site is required to be accessible to perform an unplanned failover. If the original production datacenter recovers after the failover, replication can optionally be reversed so that the failed over VM can be reprotected. In the event that the original production environment cannot be recovered, a new storage policy with alternative replication settings can be configured.

**Test Failover**
Allows the recovery copy of a VM to be brought up for testing, without taking down the production VM. Test failover enables temporary access to the recovery environment to allow a failover plan to be verified before an actual disaster or planned migration. When the test failover is stopped, any changes made to the recovery copy during the test are discarded.

### Performing a Test Failover

When performing a failover test in the VMware environment, the following restrictions apply to the replication groups being tested on the recovery site:

- Creating and deleting vVols is not supported in the test copy of the replication group. These operations are still supported on the production copy of the replication group and any changes made to the production copy while the test is active will be replicated to the recovery copy when the test is stopped.
- Any snapshots created from virtual machines in the replication group must be removed before stopping the test.
- Promoting a replication group directly from test to production is not supported. The test must first be stopped before performing a failover.

## Storage partitions

Storage partitions are used to implement Policy-based High Availability solution. Partitions contain volumes, volume groups, hosts, and host-to-volume mappings.

Within a partition:

- All volumes are in volume groups.

- Mappings can only be created between volumes and hosts in the same partition.

Each partition that is associated with an HA replication policy has two properties - the preferred management system and the active management system.

All configuration actions on a storage partition must be performed on the active management system. The storage partition can be monitored on either system.

The preferred management system is the system that you would like to be the active management system under ideal conditions. In the event of a situation where the active management system and the preferred management system are not the same system, the system will automatically failover the active management system back to the preferred management system when it is able. The preferred management system can be changed by the user.

You can configure additional volumes, volume groups, hosts, and host-to-volume mappings at any time, either by adding to an existing partition or by creating a new one.

**Related concepts**

"Policy-based High Availability" on page 65
Policy-based High Availability (HA) provides a solution for two Storage Systems in different locations where the storage will automatically remain accessible to hosts if there is an event that impacts the infrastructure and makes one of the systems unavailable. Hosts will seamlessly fail over to the other system. The solution ensures that both data and configuration are kept consistent on both systems.

## Snapshots

A volume group snapshot is an object that holds a consistent point-in-time copy of a volume group or a set of volumes.

A volume group snapshot can be used to create a clone or thin-clone volume group or volume pre-populated with the contents of the volume group snapshot. It can also be used to restore a parent volume group or volumes or a subset of the parent volumes.

Volume group snapshots are immutable, meaning that their content is protected from manipulation by the host. The volume group snapshots consume space in the parent volume's storage pool either, by default, in the same pool or a child pool of the parent volume's pool. Volume group snapshots can be added and removed to a volume group automatically by using the snapshot scheduler.

**Restriction:**

On the SAN Volume Controller platform, volume group snapshots can be added to groups of volumes that are mirrored. On other platforms, volume mirroring is not generally permitted. It can only be used for the volume migration use-case, that is, with the auto delete option.

Volume group snapshots can be made cyber-resilient by marking the snapshot as Safeguarded. Safeguarded snapshots automatically remove themselves based on retention period specified at creation time. Only users with Security Administrator access can remove these snapshots before this point. Like a standard volume group snapshot, Safeguarded snapshots can be automatically added and removed by configuring a snapshot policy to be safeguarded.

Volume group snapshots cover some of the use cases that have historically been covered by FlashCopy. The two mechanisms are not compatible, but a migration path exists that allows the user to add volume groups snapshots for volumes that are already part of a FlashCopy mapping. After volume group snapshots have been added, the user cannot create or start any new FlashCopy mappings that are associated with those volumes. This restriction does not include FlashCopy used for features like replication or cloud tiering, those features are compatible with volume group snapshots.

## Snapshot policies

A *snapshot policy* is a set of rules that controls the creation, retention, and expiration of snapshots.

With snapshot policies, administrators can schedule the creation of snapshots for volumes in a volume group at specific intervals and retain based on their security and recovery point objectives (RPO). A snapshot policy has following properties:

- It can be assigned to one or more volume groups.
- 
- Only one *snapshot policy* can be scheduled to one volume group.
- The system supports a maximum number of 32 snapshot policies.

The system supports an internal scheduler to manage and create snapshot policies on the system. The management GUI supports selecting either a user-defined policy or a predefined policy and the user-defined policies can be created by using the management GUI. Predefined policies contain specific retention and frequency values for common use-cases. Both predefined and user-defined policies are managed on the IBM Storage Virtualize system. The following predefined policies are supported:

**Predefinedsspolicy0**
: Select this policy for the most frequent copies and retention. For this policy, snapshots are created every six hours and retained for a week. Use this policy for volume data that requires the highest recovery point objective (RPO). For example, volume data that is frequently updated and critical to your business can benefit from frequent copies and retention. Customer accounts, orders, or proprietary information are examples of data that can need more frequent backups. For more information, refer to your organization's business continuity plan.

**Predefinedsspolicy1**
: Select this policy for less frequent copies and medium retention. For this policy, snapshots are created weekly and retained for a month. Use this policy for application data that is updated frequently and requires a high RPO, but might not contain business-critical data.

**Predefinedsspolicy2**
: Select this policy for less frequent copies and longer retention. For this policy, snapshots are created monthly and retained for a year. Use this policy for older data that is not updated frequently but still requires retention, such as past customer accounts or employee records.

## Safeguarded snapshots

Safeguarded snapshots are point-in-time cyber-resilient copy of volumes groups.

The system supports an internal scheduler and external scheduling applications to create and manage Safeguarded snapshots based on frequency and retention. An internal scheduler creates immutable point-in-time snapshots of volume groups that are created automatically based on the schedule that is defined in the snapshot policy. Before you create a volume group, determine whether source volumes that you want to create Safeguarded snapshots are supported. For Safeguarded snapshots, source volumes cannot have the following properties:

- The volume is of size zero.
- The volume is target of flash copy mapping.
- The ownership group of the pool for volume does not match with the ownership group of the volume group.
- The volume is mirrored.

Safeguarded snapshots are created with either an internal scheduler or an external scheduler.

- The internal scheduler uses either the default or a user-defined snapshot policy that is configured on the system. Safeguarded snapshots that are created with the internal scheduler use the same features as the snapshot function and are created and managed on the IBM Storage Virtualize system. You can configure Safeguarded snapshots with the internal scheduler if the volume group is not assigned to a Safeguarded backup policy.

- Safeguarded snapshots with the external scheduler are created and managed with IBM Copy Services Manager. You can configure Safeguarded snapshots with the external scheduler if the volume group is not assigned to a snapshot policy.
- Only a Security administrator can delete Safeguarded snapshots.

### Safeguarded snapshots with internal scheduler

Safeguarded snapshot operates similar to snapshots and can use internal scheduler. A new volume group can be created and prepopulated from another volume group's snapshot or Safeguarded snapshot. It inherits data, volumes, and volume groups from the snapshot. The use cases of this feature are data analysis, data cloning, testing, and development. For more information, see “Snapshot policies” on page 59 and “Prepopulated volume groups” on page 61.

### Safeguarded snapshots with IBM Copy Services Manager

Safeguarded snapshots with IBM Copy Services Manager simplifies the overall configuration and management of snapshots. IBM Copy Services Manager controls both scheduling and creating Safeguarded snapshots that eliminate configuration steps on the IBM Storage Virtualize system. If using the previous version of Safeguarded Copy function that uses a Safeguarded backup policy to manage Safeguarded backups, you must remove the Safeguarded backup policy from the system to use Safeguarded snapshots. .

To configure Safeguarded snapshots with IBM Copy Services Manager, you only need to create an administrator user and assign volumes to a volume group. You can also use the create clone or thin-clone feature to create copies an existing Safeguarded snapshot to determine whether the source volumes were compromised. When a clone or thin-clone is created, you can examine the data and determine whether the data is usable. You can also map the cloned volume group to host applications in IBM Copy Services Manager. With the thin-clone and clone feature, you can test and recover data on a volume group basis.

Safeguarded snapshot does not support automatic restore-in-place. However, you can create the clone or thin-clone of Safeguarded snapshot to recover data for analysis and testing.

**Related information**

Managing snapshot policies

Configuring Safeguarded snapshots with IBM Copy Services Manager

## Safeguarded snapshot user roles

Safeguarded snapshot user roles explains exclusively all the responsibilities and roles that are given to the Safeguarded snapshot function users.

The following section explains all the responsibilities and roles that are given to the Safeguarded Copy function users.

**System administrator**

The system administrator can do the following actions on the Safeguarded copy objects:

- Create a Safeguarded volume group.
- Assign volumes to the Safeguarded volume group.
- Delete a Safeguarded source volume.
- Delete a Safeguarded volume group.
- Associate a Safeguarded snapshot policy to a volume group.

The system administrator cannot do the following actions on the Safeguarded Copy function objects:

- Delete Safeguarded snapshot.
- Remove Safeguarded snapshot policy association from a volume group.

**Security administrator**

In addition to the system administrator, the security administrator can do the following actions on the Safeguarded snapshots:

- Delete Safeguarded snapshots.
- Remove Safeguarded snapshot policy association from a volume group.

### Suspending snapshot policy

You can suspend a Safeguarded snapshot policy to stop creating and removing copies for a volume group.

If the source volumes are compromised during a cyberattack or during a disaster recovery scenario, the system supports suspending all Safeguarded snapshots policies until the compromised data can be determined. During an attack, Safeguarded snapshot can contain malicious data or software. You can suspend Safeguarded snapshots at the system level with the **chsystem snapshotpolicysuspended** *yes* command then you can examine the existing uncompromised Safeguarded snapshots.

You can use uncompromised Safeguarded snapshots to recover data after the snapshot policy is suspended. When the policy is suspended, the internal scheduler does not create new snapshots and prevents expired snapshots from being deleted. Safeguarded snapshot policy is suspended at the system level. You can suspend the policy by **chsystem safeguardedcopysuspended** *yes* command, and the **snapshot_policy_suspended** parameter in the **lssystem** command displays the suspension status.

In addition to responding to a potential breach, suspending the Safeguarded Copy function can also be used as part of a disaster recovery response. However, suspending Safeguarded Copy function prevents any new snapshots being created which results in a lack of new snapshots that can be used for future recovery. Before suspending the Safeguarded Copy function, contact your support representative for assistance.

Suspending snapshot policy is not applicable to Safeguarded snapshot with IBM Copy Services Manager.

## Prepopulated volume groups

A new volume group can be created and prepopulated from another volume group's snapshot or Safeguarded snapshot. It inherits data, volumes, and volume groups from the snapshots. The use cases of this feature are data analysis, data cloning, testing, and development.

The snapshots are used to prepopulate a new volume group for various use cases, such as creating a test environment for application development. The system supports creating either a thin clone or a clone copy. Administrators can use the management GUI or the command-line interfaces to create a copy of volume data in a volume group and add the copy to a new volume group. The system also supports the ability to create these copies from a single volume.

The volume groups are populated in following snapshot types:

### Thin clone

A thin clone is a host-accessible copy of a specified snapshot. The thin clone will always depend on its source volume but can be modified by the host. It is normally created in the same storage pool as the source volume, but the thin clone can also be created in a separate pool if needed. Similarly, the thin clone will normally be created in the same I/O group as the source volume, but it can be created in a different I/O group if required.

Thin clone have the following properties:

- A thin clone is always dependent on the source snapshot.
- A thin clone can have its own snapshots.
- The capacity of a thin clone can be expanded or shrunk. It can not be shrunk to a size that is smaller than any of its snapshots or its source volume.
- A thin clone can be restored from one of its own snapshots.

- A thin clone snapshot can be used to create and populate a clone.

### Clone

A clone is a host-accessible copy of a snapshot. Once all the data from its source volume has been copied to the clone it becomes independent of its source volume. During the copy process the volume will display as type 'clone', this will change once the copy is complete and the volume becomes a standard volume.

Clone have following properties:

- A clone is populated from a volume group snapshot.
- A clone can have its own snapshots.
- While the copy is in progress, a clone can be expanded or shrunk. The volume can not be shrunk to a size smaller than any of its snapshots or its source volume. Once the copy is complete the volume behaves as a standard volume and the limitation on shrinking the volume no longer applies.
- A clone can populate a thin clone with one of its snapshots.
- A clone can populate another clone with one of its snapshots.

Thin clone and clone copies can be used for the following use cases.

**Analysis And Design**
Creates copies of application data for data mining type activities without impacting the production data.

**Data Cloning**
Instantly creates a new user environment from a previously defined image.

**Test And Development**
Create the copies of application data for test and development tasks without impacting the production data.

## Restoring a volume group

Snapshot restore provide the user with the ability to restore a production volume group with the contents of one of the volume groups snapshots. This feature is sometimes referred to as *restore in place*.

This type of restore operation overwrites the contents of the production system with a previously saved version of the data. As opposed to recovery, which allows the user to access the previously saved version and copy the relevant data from the saved copy back to production system. This latter type of recovery is available with the existing snapshot solution by creating a thin clone volume group of the saved copy or snapshot.

Snapshot restore allows restores at the volume group level or to a user specified subset of volumes within a volume group.

Snapshot restore is not supported on volume groups using policy-based replication.

A snapshot refresh helps to refresh the data on a set of volumes, whether the volumes are in a volume group or in a prepared list of volumes. For more information on snapshot refresh, see “Refresh a prepopulated volume group” on page 62.

## Refresh a prepopulated volume group

Snapshot refresh can be used to refresh the data on a set of thin clone volumes from a volume group snapshot.

Any existing data on the set of volumes are discarded and the volumes are populated with the data from the specified snapshot.

Snapshot refresh operates at the volume group level or to a user specified subset of volumes within a volume group. Only thin clone volumes can be refreshed.

Snapshot refresh allows the user to:

- Examine the data on different snapshots without having to map a new set of volumes to any hosts.
- Discard any changes made to a set of thin-clone volumes and repopulate them with the original data from a snapshot.

You can use the command-line interface (CLI) to refresh the snapshot. For more information, see **refreshfromsnapshot** command.

# Managing snapshots

You can use the management GUI or the command-line interface (CLI) to manage snapshots, safeguarded snapshots, and the volume groups.

To work with snapshots, safeguarded snapshots, and the volume group in the management GUI, select **Volumes** > **Volume Groups** > **Existing volume group** > **Snapshots**.

## Adding snapshots to volume groups

Use the management GUI to create snapshots of a volume group. Each volume snapshot occupies capacity from a specific pool, by default the parent volume is the default pool.

To add snapshots to volume groups, the following requirements must be met:

- Snapshots cannot be added to the volumes that are target volumes of legacy FlashCopy mappings.
- Snapshots cannot be added to mirrored volumes.

To add snapshots to a volume group, use the **Volumes** > **Volume Groups** panel in management GUI. Select the volume group that needs to be associated with the snapshots and under the **Snapshots** tab, select **Take Snapshot**.

## Creating volume groups from snapshots

Volume groups can be created from a snapshot. Snapshots themselves are not host accessible so creating a volume group from a snapshot and creates a host accessible copy of that snapshot.

To create a volume group from snapshots, use the **Volumes** > **Volume Groups** > **Create Volume Group** panel in management GUI. Select **Choose existing snapshot from a volume group**, select the required snapshot and specify the copy type that is required.

## Restoring a volume group from snapshots

Volume groups can be restored from one of its snapshots.

To restore a volume group from snapshots, the following requirements must be met:

- When requesting a restore operation, the volume group specified by the **volumegroup** parameter must be the same parent group from which the snapshot was originally created.
- The composition of the volume group must be the same at the time of the restore as it was at the time the snapshot was taken.
- If volumes have been added to or removed from the volume group in the time between the snapshot being taken and the restore being requested. Then those volumes must be removed from or added back into the volume group before the restore can be performed.

  **Note:** The actual volumes involved in the creation of the snapshot must be in the volume group when starting a restore, for example, it is not possible to create new volumes within the volume group in order to 'replace' volumes that have been moved.
- If a volume has been deleted after the snapshot was taken, it will have been removed from the volume group and will be in the deleting state. By requesting a snapshot restore, assuming all other prerequisites have been met and the restore operation proceeds. The volume will be added back to the volume group and put into the active state.

- If the volumes have been expanded between the time that the snapshot was taken and the restore requested, then the restore will fail. The user will have to shrink the parent volumes back to the size they were when the snapshot was taken before a restore will go ahead.
- This has implications if new snapshots have been taken since a volume was expanded, as you will not be able to shrink the parent volumes without cleaning out any new snapshots, and their dependent volume groups.

To restore a volume group from snapshots, use the **Volumes** > **Volume Groups** panel in management GUI. Select the volume group that needs to be associated with the snapshots and under the **Snapshots** tab, select the **Restore** option on the overflow menu for the snapshot being used for the restore. There is an option to either restore the entire volume group or a subset of volumes within that group. Select the applicable option and the subset of volumes if applicable and select **Restore**.

### Deleting snapshots

To delete snapshots, use the **Volumes** > **Volume Groups** panel in management GUI. Select the volume group that needs to be associated with the snapshots and under the **Snapshots** tab, select the snapshot to delete and click overflow menu to select **Delete**.

## Replication policies

A replication policy defines how replication is configured between systems.

A replication policy is defined between two systems in a partnership that supports policy-based replication. The partnership must be fully-configured and each system must be connected. Each replication policy identifies a single I/O group on each system. To replicate volume groups from multiple I/O groups, you will need to create another replication policy. You can use the management GUI interface and command-line interface (CLI) for creating replication policies.

Policies that use asynchronous replication must also specify an RPO (recovery point objective) alert. The system always tries to achieve the lowest possible recovery point. However, if the recovery point increases beyond this value an alert is raised. If the RPO exceeds that value, an alert is sent.

A replication policy defines three key attributes:

**A set of locations**
Defines the I/O groups on the partnered systems that contain a replicated copy of the volume group or storage partition. The location defines *where* data is replicated.

**A topology**
Represents organization of the systems and the type of replication that is completed between each location. The topology defines *how* data is replicated between the locations.

**A recovery point objective (RPO) for asynchronous replication topologies**
Defines a maximum acceptable RPO for the asynchronous replication between locations.

The following rules apply to replication policies:

- An asynchronous replication policy can be assigned to one or more volume groups.
- A high availability replication policy can be assigned to one or more storage partitions.
- A replication policy cannot be assigned to a volume group that is in a storage partition.
- Replication policies cannot be changed after they are created. If changes are required, a new policy can be created and assigned to a volume group.
- Each system supports up to a maximum of 32 replication policies.

There are two uses cases for replication policy:

**Policy-based replication: Asynchronous**

For more information, refer to the "Policy-based replication: Asynchronous" on page 65 section.

**Policy-based High Availability**

For more information, refer to the “Policy-based High Availability” on page 65 section.

### Creating a replication policy

To create a replication policy, use the **Policies** > **Replication Policies** panel in the management GUI. You can also create replication policy by following the **Copy services** > **Partnerships and remote copy** panel in the management GUI and click **Setup policy-based replication**

On the left under **Location 1**, select the first system to use in the replication policy. This system always appears on the left when managing and monitoring replication. On the right under **Location 2**, select the second system to use in the replication policy. This system always appears on the right when managing and monitoring replication. If either system contains multiple I/O groups, select the I/O group that will be the caching I/O group for all volumes replicated with this policy.

### Assigning an asynchronous replication policy to a volume group

To assign an asynchronous replication policy to a volume group, use the **Volumes** > **Volume Groups** panel in the management GUI. Select the volume group that needs to be associated with the replication policy, and under the **Policies** tab, click **Assign Replication Policy**.

### Assigning a high availability replication policy to an existing storage partition

To configure a high availability replication policy on a new storage partition, use the **Copy services** > **Partnerships and remote copy** panel in the management GUI and click **Setup policy-based replication**.

To assign a high availability replication policy to an existing storage partition, use the **Storage partitions** panel in the management GUI. Choose a partition that does not already have a replication policy assigned and click **Add High availability replication**.

For more information on the command-line interface, see **`mkreplicationpolicy`** command.

## Policy-based replication: Asynchronous

*Policy-based replication* uses volume groups and replication policies to automatically deploy and manage replication. Policy-based replication significantly simplifies configuring, managing, and monitoring replication between two systems.

With policy-based replication, you can replicate data between systems with minimal management, significantly higher throughput and reduced latency compared to the remote-copy function. A replication policy has following properties:

- A replication policy can be assigned to one or more volume groups.
- Replication policies cannot be changed after they are created. If changes are required, a new policy can be created and assigned to the associated volume group.
- Each system supports up to a maximum of 32 replication policies.

For more information, refer to Getting started with policy-based replication.

## Policy-based High Availability

Policy-based High Availability (HA) provides a solution for two Storage Systems in different locations where the storage will automatically remain accessible to hosts if there is an event that impacts the infrastructure and makes one of the systems unavailable. Hosts will seamlessly fail over to the other system. The solution ensures that both data and configuration are kept consistent on both systems.

## Planning High Availability

Review the topics in this section to understand what is covered in this scenario, the reasons why a business might want to follow the scenario, and to see an overview of the solution proposed by the scenario.

### *Environment*

Each location contains a system that supports Policy-based High Availability (HA) with Fibre Channel network connectivity between them.

The locations are sufficiently far apart such that anticipated disruptions affect only one location. The locations are sufficiently close together to not cause the response time for I/O to exceed the maximum that an application can tolerate because of latency in replication link communications. For policy-based High Availability, a maximum of 1 ms RTT (Round-Trip Time) inter site latency is supported.

Both systems have access to an external IP quorum app to arbitrate if a site or system loss occurs. Multiple quorum applications can be deployed for redundancy.

Hosts in a high-availability solution must use an ALUA-compliant multipath policy. For more information, see Host attachment.

### *Solution overview*

The High Availability (HA) solution uses Storage Partitions, a configuration object that is the single point of management for HA objects. These partitions contain volumes, volume groups, hosts, and host-to-volume mappings.

Within a partition:

- All volumes are in volume groups.
- Mappings can be created only between volumes and hosts in the same partition.

An HA replication policy can be associated with a partition. The HA replication policy defines the two systems that are connected by a partnership that are providing the HA solution.

When the partitions are associated with an HA replication policy, all the objects that are contained in the partition are automatically configured across both systems that are associated with the replication policy. Hosts that are created in a partition will be able to discover paths to the mapped volumes through both systems.

Before creation of storage partitions, the systems are configured with “Partnerships” on page 9, Storage pools that are linked between systems (see “Linked Pools” on page 34), “Replication policies” on page 64, and “IP quorum application ” on page 36. You can configure multiple partitions for HA.

Each partition that is associated with an HA replication policy has two properties - the preferred management system and the active management system.

The active management system is the system from which all configuration tasks must be performed. If an outage or other failure happens on the current active management system, the active management system will automatically fail over to the other system.

The preferred management system is the system that you would like to be the active management system under ideal conditions. If the active management system and the preferred management system are not the same system, the system will automatically failover the active management system back to the preferred management system when it is able. The preferred management system can be changed by the user.

Everything in a configured Storage Partition is highly available. When required, both management and I/O fail over between systems to maintain access.

You can configure more volumes, volume groups, hosts, and host-to-volume mappings at any time, either by adding to an existing partition or by creating a new one.

The IP Quorum application determines which system becomes the active management system and avoids a scenario where both halves of the High Availability solution continue to process the same partition.

## Implementing High Availability

This guide explains how to configure a Policy-based High Availability (HA) solution. You can setup Policy-based High Availability (HA) solution using the management GUI or the CLI.

### *Configuring and monitoring High Availability by using the GUI*

High Availability should be configured from the **Storage Partitions** or the **Copy Services** > **Partnerships** panel in the management GUI.

Follow the steps to configure Policy-based High Availability (HA) by using the management GUI:

1. If you do not already have a partnership, refer to .
2. Select a partnership that is ready for use with policy-based replication, and select **Setup policy-based replication**.
3. Choose the High-availability replication type.
4. The management GUI guides you through the required steps to:
   - Configure an IP quorum application.
   - Link storage pools between systems.
   - Create an HA replication policy and a storage partition.
   - Select existing volume groups and mapped hosts to add to the partition.
   - Create new hosts, volume groups, and host-to-volume mappings on both systems.

Use the **Storage Partition Overview** panel to monitor connectivity between the two systems and the IP quorum applications, and the health of the Hosts and Volumes associated with the partition.

### *Configuring and monitoring High Availability by using the CLI*

It is recommended that Policy-based High Availability (HA) is configured and managed by using the GUI. Configuration can also be performed by using the REST API or the CLI.

Prerequisite to configure your systems to support the High Availability solution:

- If not already configured, follow the link to set up a partnership: .
- If not already configured, follow the link to set up a linked storage pool: .
- If not already configured, follow the link to set up an IP quorum application: “IP quorum application ” on page 36.

Follow the steps to configure Policy-based High Availability (HA) configuration by using the CLI:

1. To create a replication policy with the *2-site-ha* topology, enter the following command.

   ```
   mkreplicationpolicy -topology 2-site-ha -location1system <name or ID of one system>
   -location1iogrp 0 -location2system <name or ID of other system> -location2iogrp 0
   ```

   if necessary, use **lspartnership** to get the name or ID of each system. For more information, see **mkreplicationpolicy**.
2. To create a Storage Partition and associate it with the replication policy, enter the following command.

   ```
   # Creating a single-site storage partition that is later configured for HA
   mkpartition -name my_partition
   ...
   chpartition -replicationpolicy <HA policy ID or name> my_partition

   # Creating an HA storage partition
   mkpartition -name ha_partition -replicationpolicy <HA policy ID or name>
   ```

   For more information, see **mkpartitions**.
3. To create a logical host object that is maintained across both systems by specifying that the host should be associated with the Storage Partition, enter the following command.

   ```
   mkhost -partition <storage partition ID or name>
   ```

For more information, see **mkhost**.

4. To create a volume group that is maintained across both systems by specifying that the volume group should be associated with the Storage Partition, enter the following command.

```
mkvolumegroup -partition <storage partition ID or name>
```

   For more information, see **mkvolumegroup**.

5. To create a volume in this volume group, enter the following command.

```
mkvolume -volumegroup <volume group ID or name>
```

   For more information, see **mkvolume**.

6. To map the created volume to the created host, enter the following command.

```
mkvdiskhostmap -host <host ID or name> <volume ID or name>
```

   For more information, see **mkvdiskhostmap**.

7. Hosts can now discover paths to volumes at both locations. Monitor and change the active and preferred management systems that use **lspartition** and **chpartition**. Monitor connectivity between the systems that use **lspartnership**. Monitor the health of the HA solution using **lspartition**.

# System Monitoring

The system supports sending notifications for remote monitoring. The system supports sending event log entries using Simple Network Management Protocol (SNMP). The system also supports sending event log entries, audit log entries, and authentication attempts using syslog.

## Statistics collection

The system collects statistics over an interval and creates files that can be viewed.

### Introduction

For each collection interval, the management GUI creates five statistics files: one for managed disks (MDisks), which is named **Nm_stats**; one for volumes and volume copies, which is named **Nv_stats**; one for volumes groups, which is named **Ng_stats**; one for nodes, which is named **Nn_stats**; and one for drives, which is named **Nd_stats**. The files are written to the `/dumps/iostats` directory on the node. To retrieve the statistics files from the non-configuration nodes onto the configuration node, use the **svctask cpdumps** command.

A maximum of 16 files of each type can be created for the node. When the 17th file is created, the oldest file for the node is overwritten.

### Fields

The following fields are available for user definition:

**Interval**
Specify the interval in minutes between the collection of statistics. You can specify 1 - 60 minutes in increments of 1 minute. Statistics must be set to 1 or 5 minute intervals when sending data to Storage Insights.

### Tables

The following tables describe the information that is reported for individual nodes and volumes.

Table 10 on page 69 describes the statistics collection for MDisks, for individual nodes.

| *Table 10. Statistics collection for MDisks for individual nodes* | |
|---|---|
| **Statistic name** | **Description** |
| id | Indicates the name of the MDisk for which the statistics apply. |
| idx | Indicates the identifier of the MDisk for which the statistics apply. |
| pre | Indicates the peak of read external response time in milliseconds for each MDisk. The external response time for disk reads is calculated by starting a timer when a SCSI read command is issued and stopped when the command completes successfully. |
| pro | Indicates the peak of read queued response time in milliseconds for each MDisk. The value means the peak elapsed time that is taken for read commands to complete from the time they join the queue. |
| pwe | Indicates the peak of write external response time in milliseconds for each MDisk. The external response time for disk writes is calculated by starting a timer when a SCSI write command is issued and stopped when the command completes successfully. |
| pwo | Indicates the peak of write queued response time in milliseconds for each MDisk. The value means the peak elapsed time that is taken for write commands to complete from the time they join the queue. |
| rb | Indicates the cumulative number of blocks of data that was read (since the node started). |
| re | Indicates the cumulative read external response time in milliseconds for each MDisk. The cumulative response time for disk reads is calculated by starting a timer when a **SCSI read** command is issued and stopped when the command completes successfully. The elapsed time is added to the cumulative counter. |
| ro | Indicates the cumulative number of MDisk read operations that were processed (since the node started). |
| rq | Indicates the cumulative read queued response time in milliseconds for each MDisk. This response is measured from above the queue of commands to be sent to an MDisk because the queue depth is already full. This calculation includes the elapsed time that is taken for **read** commands to complete from the time they join the queue. |
| ure | Indicates the cumulative read external response time in microseconds for each MDisk. The cumulative response time for disk reads is calculated by starting a timer when a SCSI read command is issued and stopped when the command completes successfully. The elapsed time is added to the cumulative counter. |
| urq | Indicates the cumulative read queued response time in microseconds for each MDisk. This response is measured from above the queue of commands to be sent to an MDisk because the queue depth is already full. This calculation includes the elapsed time that is taken for read commands to complete from the time they join the queue. |
| uwe | Indicates the cumulative write external response time in microseconds for each MDisk. The cumulative response time for disk writes is calculated by starting a timer when a SCSI write command is issued and stopped when the command completes successfully. The elapsed time is added to the cumulative counter. |
| uwq | Indicates the cumulative write queued response time in microseconds for each MDisk. This time is measured from above the queue of commands to be sent to an MDisk because the queue depth is already full. This calculation includes the elapsed time that is taken for write commands to complete from the time they join the queue. |
| wb | Indicates the cumulative number of blocks of data written (since the node started). |

| *Table 10. Statistics collection for MDisks for individual nodes (continued)* | |
|---|---|
| **Statistic name** | **Description** |
| we | Indicates the cumulative write external response time in milliseconds for each MDisk. The cumulative response time for disk writes is calculated by starting a timer when a **SCSI write** command is issued and stopped when the command completes successfully. The elapsed time is added to the cumulative counter. |
| wo | Indicates the cumulative number of MDisk write operations that were processed (since the node started). |
| wq | Indicates the cumulative write queued response time in milliseconds for each MDisk. This time is measured from above the queue of commands to be sent to an MDisk because the queue depth is already full. This calculation includes the elapsed time that is taken for write commands to complete from the time they join the queue. |

**Note:** MDisk statistics files for nodes are written to the `/dumps/iostats` directory on the individual node.

Table 11 on page 70 describes the volume information that is reported for individual nodes.

**Note:** The system only collects a subset of statistics for volumes used for snapshot function and policy-based replication.

| *Table 11. Statistic collection for volumes for individual nodes* | |
|---|---|
| **Statistic name** | **Description** |
| id | Indicates the volume name for which the statistics apply. |
| idx | Indicates the volume for which the statistics apply. |
| rb | Indicates the cumulative number of blocks of data read (since the node started). |
| rl | Indicates the cumulative read response time in milliseconds for each volume. The cumulative response time for volume reads is calculated by starting a timer when a SCSI read command is received and stopped when the command completes successfully. The elapsed time is added to the cumulative counter. |
| rlw | Indicates the worst read response time in microseconds for each volume since the last time statistics were collected. This value is reset to zero after each statistics collection sample. |
| ro | Indicates the cumulative number of volume read operations that were processed (since the node started). |
| rxl | Indicates the cumulative read data transfer response times in milliseconds for each volume since the last time the node was reset. A high value for this statistic indicates that latency is likely caused by the fabric or the host, which is submitting more read commands than it can process. |
| ub | Indicates the cumulative number of blocks of data unmapped (since the node started). |
| ul | Indicates the cumulative unmap response time in milliseconds for each volume. The cumulative response time for volume unmaps is calculated by starting a timer when a SCSI unmap command is received and stopped when the command completes successfully. The elapsed time is added to the cumulative counter. |
| ulw | Indicates the worst unmap response time in microseconds for each volume. The worst response time for volume unmaps is calculated by starting a timer when a SCSI unmap command is received and stopped when the command completes successfully. |

| *Table 11. Statistic collection for volumes for individual nodes (continued)* | |
|---|---|
| **Statistic name** | **Description** |
| uo | Indicates the cumulative number of volume unmap operations that were processed (since the node started). |
| uou | Indicates the cumulative number of volume unmap operations that are not aligned on an 8 K boundary (according to the alignment/granularity setting in Block Limits VPD Page (0xb0). |
| wb | Indicates the cumulative number of blocks of data written (since the node started). |
| wl | Indicates the cumulative write response time in milliseconds for each volume. The cumulative response time for volume writes is calculated by starting a timer when a SCSI write command is received and stopped when the command completes successfully. The elapsed time is added to the cumulative counter. |
| wlw | Indicates the worst write response time in microseconds for each volume since the last time statistics were collected. This value is reset to zero after each statistics collection sample. |
| wo | Indicates the cumulative number of volumes write operations that were processed (since the node started). |
| wou | Indicates the cumulative number of volumes write operations that are not aligned on a 4 K boundary. |
| wxl | Indicates the cumulative write data transfer response times in milliseconds for each volume since the last time the node was reset. A high value for this statistic indicates that latency is likely caused by the fabric and/or the host responding slowly when the node requests the data to write to the volume. |
| xl | Indicates the cumulative read and write data transfer response time in milliseconds for each volume since the last time the node was reset. When this statistic is viewed for multiple volumes and with other statistics, it can indicate whether the latency is caused by the host, fabric, or the FlashSystem 9500 . |

**Note:** For unmap statistics, it is where an unmap operation is a **SCSI unmap** or **Write same with unmap** command.

Table 12 on page 72 describes the VDisk information that is related to Metro Mirror or Global Mirror relationships that is reported for individual nodes.

| *Table 12. Statistic collection for volumes that are used in Metro Mirror and Global Mirror relationships for individual nodes* | |
|---|---|
| **Statistic name** | **Description** |
| gwl | Primary VDisk: This statistic accumulates the primary write lag, which is the time taken after a primary global mirror write completes until the data has been received and hardened in the secondary cluster and notification of that has been sent back to and received by the primary cluster. The average value for one operation can be found by dividing `gwl` by `gws` (the number of global mirror writes in the sample period). While running, the value of `gwl/gws` is approximately the Recovery Point Objective (RPO). This calculation does not work when the relationship is stopped.<br><br>Secondary VDisk: This statistic accumulates the secondary write lag, which is the time taken after a write is received on the secondary cluster before it is submitted to be written to disk.<br>The average value for one operation can be found by dividing `gwl` by `gws` (the number of global mirror writes in the sample period). |
| gwo | Indicates the total number of overlapping volume writes. An overlapping write is when the logical block address (LBA) range of write request collides with another outstanding request to the same LBA range and the write request is still outstanding to the secondary site. |
| gwot | Indicates the total number of fixed or unfixed overlapping writes. |
| gws | Indicates the total number of write requests that are issued to the secondary site. |

Table 13 on page 72 describes the port information that is reported for individual nodes.

| *Table 13. Statistic collection for node ports* | |
|---|---|
| **Statistic name** | **Description** |
| bbcz | Indicates the total time in microseconds for which the buffer credit counter was at zero. That this statistic is only reported by 8 Gbps Fibre Channel ports. For other port types, this statistic is 0. |
| cbr | Indicates the bytes received from controllers. |
| cbt | Indicates the bytes transmitted to disk controllers. |
| cer | Indicates the commands that are received from disk controllers.<br>**Note:** The cer metric is always 0. |
| cet | Indicates the commands that are initiated to disk controllers. |
| dtdc | Indicates the number of transfers that experienced excessive data transmission delay. |
| dtdm | Indicates the number of transfers that had their data transmission delay measured. |
| dtdt | Indicates the total time in microseconds for which data transmission was excessively delayed. |
| har | Indicates the count aborted host read operations aborted while a data transfer was in progress. This count includes I/O operations aborted by hosts, and those operations aborted internally by the system. |

| *Table 13. Statistic collection for node ports (continued)* | |
|---|---|
| **Statistic name** | **Description** |
| haw | Indicates the count aborted host write operations aborted while a data transfer was in progress. This count includes I/O operations aborted by hosts, and those operations aborted internally by the system. |
| hbr | Indicates the bytes received from hosts. |
| hbt | Indicates the bytes transmitted to hosts. |
| her | Indicates the commands that are received from hosts. |
| het | Indicates the commands that are initiated to hosts.<br>**Note:** The het metric is always 0. |
| hsr | Indicates the count of data transfers which were considered slow for host read operations. |
| hsw | Indicates the count of data transfers which were considered slow for host write operations. |
| icrc | Indicates the number of CRC that is not valid. |
| id | Indicates the port identifier for the node. |
| itw | Indicates the number of transmission word counts that are not valid. |
| lf | Indicates a link failure count. |
| lnbr | Indicates the bytes received to other nodes in the same cluster. |
| lnbt | Indicates the bytes transmitted to other nodes in the same cluster. |
| lner | Indicates the commands that are received from other nodes in the same cluster. |
| lnet | Indicates the commands that are initiated to other nodes in the same cluster. |
| lsi | Indicates the lost-of-signal count. |
| lsy | Indicates the loss-of-synchronization count. |
| pspe | Indicates the primitive sequence-protocol error count. |
| tmp | Indicates the SFP temperature in degrees Celsius. |
| tmpht | Indicates the SFP temperature high alarm threshold in degrees Celsius. |
| txpwr | Indicates the TX power in microwatts (µW). |
| txpwrlt | Indicates the TX power low alarm threshold in microwatts (µW). |
| rmbr | Indicates the bytes received to other nodes in the other clusters. |
| rmbt | Indicates the bytes transmitted to other nodes in the other clusters. |
| rmer | Indicates the commands that are received from other nodes in the other clusters. |
| rmet | Indicates the commands that are initiated to other nodes in the other clusters. |
| rxpwr | Indicates the RX power in microwatts (µW). |
| rxpwrlt | Indicates the RX power low alarm threshold in microwatts (µW). |
| wwpn | Indicates the worldwide port name for the node. |

Table 14 on page 74 describes the node information that is reported for each node.

| *Table 14. Statistic collection for nodes* | |
|---|---|
| **Statistic name** | **Description** |
| cluster_id | Indicates the identifier of the cluster. |
| cluster | Indicates the name of the cluster. |
| cpu | `busy` - Indicates the total CPU average core busy milliseconds since the node was reset. This statistic reports the amount of the time the processor spends polling, waiting for work versus doing work. This statistic accumulates from zero. |
| | `system` - Indicates the total CPU average core busy milliseconds since the node was reset. This statistic reports the amount of the time the processor spends polling, waiting for work versus doing work. This statistic accumulates from zero. This statistic is the same information as the information provided with the `cpu busy` statistic and eventually replaces the `cpu busy` statistic. |
| cpu_core | `id` - Indicates the CPU core ID. |
| | `system` - Indicates the per-core CPU average core busy milliseconds for system process cores since node was reset. |
| env | `obj` - Indicates whether enclosure (e) or node (n). |
| | `id` - Indicates the enclosure or node ID. |
| | `p` - Indicates the power consumption of the enclosure or node. |
| | `t` - Indicates the temperature (Celsius) of the enclosure or node. |
| dimm | `id` - Indicates the memory module ID. |
| | `loc` - Indicates the location of the memory module. |
| | `manu` - Indicates the manufacturer of the memory module. |
| | `sn` - Indicates the serial number of the memory module. |
| | `ce` - Indicates the number of corrected errors in the memory module. |
| id | Indicates the name of the node. |
| lrb | Indicates the number of logical bytes received from the other node. |
| lwb | Indicates the number of logical bytes sent to the other node. |
| node_id | Indicates the unique identifier for the node. |
| rb | Indicates the number of physical bytes received from the other node. |
| re | Indicates the accumulated receive latency, excluding inbound queue time. This statistic is the latency that is experienced by the node communication layer from the time that an I/O is queued to cache until the time that the cache gives completion for it. |
| ro | Indicates the number of messages or bulk data received. |
| rq | Indicates the accumulated receive latency, including inbound queue time. This statistic is the latency from the time that a command arrives at the node communication layer to the time that the cache completes the command. |
| wb | Indicates the number of physical bytes sent to the other node. |
| we | Indicates the accumulated send latency, excluding outbound queue time. This statistic is the time from when the node communication layer issues a message out onto the Fibre Channel until the node communication layer receives notification that the message arrived. |

| *Table 14. Statistic collection for nodes (continued)* | |
|---|---|
| **Statistic name** | **Description** |
| wo | Indicates the number of messages or bulk data sent. |
| wq | Indicates the accumulated send latency, including outbound queue time. This statistic includes the entire time that data is sent. This time includes the time from when the node communication layer receives a message and waits for resources, the time to send the message to the remote node, and the time that is taken for the remote node to respond. |

There are three types of data reduction properties per data reduction pool.

- dca - these statistics are related to the data stored within the data reduction pool.
- rca - these statistics are related to I/O to manage the background garbage collection processes of the data reduction pool.
- jca - these statistics are related to journaling operations for the metadata that manages the data reduction pool.

Table 15 on page 75 provides details about the statistic collection for volume performance.

| *Table 15. Performance statistics for volume.* This table lists the performance statistics that is reported for individual volume. | |
|---|---|
| **Statistic name** | **Description** |
| entav | Internal cyber resiliency statistic average |
| entcn | Internal cyber resiliency statistic count |

| *Table 16. Statistic collection for volume cache per individual nodes.* This table describes the volume cache information that is reported for individual nodes. | |
|---|---|
| **Statistic name** | **Description** |
| cm | Indicates the number of sectors of modified or dirty data that are held in the cache. |
| ctd | Indicates the total number of cache destages that were initiated writes, submitted to other components as a result of a volume cache flush or destage operation. |
| ctds | Indicates the total number of sectors that are written for cache-initiated track writes. |
| ctp | Indicates the number of track stages that are initiated by the cache that are prestage reads. |
| ctps | Indicates the total number of staged sectors that are initiated by the cache. |
| ctrh | Indicates the number of total track read-cache hits on prestage or non-prestage data. For example, a single read that spans two tracks where only one of the tracks obtained a total cache hit, is counted as one track read-cache hit. |
| ctrhp | Indicates the number of track reads received from other components, which are treated as cache hits on any prestaged data. For example, if a single read spans two tracks where only one of the tracks obtained a total cache hit on prestaged data, it is counted as one track that is read for the prestaged data. A cache hit that obtains a partial hit on prestage and non-prestage data still contributes to this value. |
| ctrhps | Indicates the total number of sectors that are read for reads received from other components that obtained cache hits on any prestaged data. |
| ctrhs | Indicates the total number of sectors that are read for reads received from other components that obtained total cache hits on prestage or non-prestage data. |

| *Table 16. Statistic collection for volume cache per individual nodes.* This table describes the volume cache information that is reported for individual nodes. *(continued)* | |
|---|---|
| **Statistic name** | **Description** |
| ctr | Indicates the total number of track reads received. For example, if a single read spans two tracks, it is counted as two total track reads. |
| ctrs | Indicates the total number of sectors that are read for reads received. |
| ctwft | Indicates the number of track writes received from other components and processed in flush through write mode. |
| ctwfts | Indicates the total number of sectors that are written for writes that are received from other components and processed in flush through write mode. |
| ctwfw | Indicates the number of track writes received from other components and processed in fast-write mode. |
| ctwfwsh | Indicates the track writes in fast-write mode that were written in write-through mode because of the lack of memory. |
| ctwfwshs | Indicates the track writes in fast-write mode that were written in write through due to the lack of memory. |
| ctwfws | Indicates the total number of sectors that are written for writes that are received from other components and processed in fast-write mode. |
| ctwh | Indicates the number of track writes received from other components where every sector in the track obtained a write hit on already dirty data in the cache. For a write to count as a total cache hit, the entire track write data must already be marked in the write cache as dirty. |
| ctwhs | Indicates the total number of sectors that are received from other components where every sector in the track obtained a write hit on already dirty data in the cache. |
| ctw | Indicates the total number of track writes received. For example, if a single write spans two tracks, it is counted as two total track writes. |
| ctws | Indicates the total number of sectors that are written for writes that are received from components. |
| ctwwt | Indicates the number of track writes received from other components and processed in write through write mode. |
| ctwwts | Indicates the total number of sectors that are written for writes that are received from other components and processed in write through write mode. |
| cv | Indicates the number of sectors of read and write cache data that is held in the cache. |

Table 17 on page 76 describes the XML statistics specific to an IP Partnership port.

| *Table 17. XML statistics for an IP Partnership port* | |
|---|---|
| **Statistic name** | **Description** |
| ipbz | Indicates the average size (in bytes) of data that is being submitted to the IP partnership driver since the last statistics collection period. |
| iprc | Indicates the total bytes that are received before any decompression takes place. |
| ipre | Indicates the bytes retransmitted to other nodes in other clusters by the IP partnership driver. |

| *Table 17. XML statistics for an IP Partnership port (continued)* | |
|---|---|
| **Statistic name** | **Description** |
| iprt | Indicates the average round-trip time in microseconds for the IP partnership link since the last statistics collection period. |
| iprx | Indicates the bytes received from other nodes in other clusters by the IP partnership driver. |
| ipsz | Indicates the average size (in bytes) of data that is being transmitted by the IP partnership driver since the last statistics collection period. |
| iptc | Indicates the total bytes that are transmitted after any compression (if active) takes place. |
| iptx | Indicates the bytes transmitted to other nodes in other clusters by the IP partnership driver. |

Table 18 on page 77 describes the offload data transfer (ODX) Vdisk and node level I/O statistics.

| *Table 18. ODX VDisk and node level statistics* | | |
|---|---|---|
| **Statistic name** | **Acronym** | **Description** |
| Read cumulative ODX I/O latency | orl | Cumulative total read latency of ODX I/O per VDisk. The unit type is micro-seconds (US). |
| Write cumulative ODX I/O latency | owl | Cumulative total write latency of ODX I/O per VDisk. The unit type is micro-seconds (US). |
| Total transferred ODX I/O read blocks | oro | Cumulative total number of blocks that are read and successfully reported to the host, by ODX WUT command per VDisk. It is represented in blocks unit type. |
| Total transferred ODX I/O write blocks | owo | Cumulative total number of blocks that are written and successfully reported to the host, by ODX WUT command per VDisk. It is represented in blocks unit type. |
| Wasted ODX I/Os | oiowp | Cumulative total number of wasted blocks that are written by ODX WUT command per node. It is represented in blocks unit type. |
| WUT failure count | otrec | Cumulative total number of failed ODX WUT commands per node. It includes WUT failures due to a token revocation and expiration. |

Table 19 on page 78 describes the statistics collection for cloud per cloud account ID.

| *Table 19. Statistics collection for cloud per cloud account ID* | | |
|---|---|---|
| **Statistic name** | **Acronym** | **Description** |
| id | id | Cloud account ID |
| Total Successful Puts | puts | Total number of successful PUT operations |
| Total Successful Gets | gets | Total number of successful GET operations |
| Bytes Up | bup | Total number of bytes successful transferred to the cloud |
| Bytes Down | bdown | Total number of bytes successful downloaded/read from the cloud |
| Up Latency | uplt | Total time that is taken to transfer the data to the cloud |
| Down Latency | dwlt | Total time that is taken to download the data from the cloud |
| Down Error Latency | dwerlt | Time that is taken for the GET errors |
| Part Error Latency | pterlt | Total time that is taken for part errors |
| Persisted Bytes Down | prbdw | Total number of bytes successfully downloaded from the cloud and persisted on the local storage that were part of successful GET operation |
| Persisted Bytes Up | prbup | Total number of bytes successfully transferred to the cloud and persisted on the cloud that were part of successful PUT operation. The difference is that you might have a 100 bytes file, of which you successfully had 80 bytes sent to the cloud through a PUT operation, but the last data transfer cycle carrying 20 bytes errored out, and the entire request failed. In that case, the statistics indicates: BYTES_UP = 80 and PERSISTED_BYTES_UP = 0 |
| Persisted Down Latency | prdwlt | Total time that is taken to download the data from the cloud that were part of successful GET operation |
| Persisted Up Latency | pruplt | Total time that is taken to transfer the data to the cloud that were part of successful PUT operation |
| Failed Gets | flgt | Total number of failed GET operations |
| Failed Puts | flpt | Total number of failed PUT operations |

| *Table 19. Statistics collection for cloud per cloud account ID (continued)* | | |
|---|---|---|
| **Statistic name** | **Acronym** | **Description** |
| Get Errors | gter | Total number of times a read from the cloud failed (including the last retry that failed the GET request) |
| Get Retries | gtrt | Total number of GET retries |
| Part Errors | pter | Total number of part errors. It is the count if multi part upload occurs. The part refers to the multi-part upload scenario. |
| Parts Put | ptpt | Total number of parts that are successfully transferred to the cloud |
| Persisted parts | prpt | Total number parts successfully persisted on the cloud that were part of successful put operation |
| Put retries | ptrt | Total number of PUT retries |
| Throttle upload latency | tuplt | Average delay introduced due to setting upload bandwidth limit |
| Throttle download latency | tdwlt | Average delay introduced due to setting download bandwidth limit |
| Throttle upload bandwidth utilization percentage | tupbwpc | Bandwidth utilization in percentage of configured upload bandwidth limit |
| Throttle download bandwidth utilization percentage | tdwbwpc | Bandwidth utilization in percentage of configured download bandwidth limit |

Table 20 on page 79 describes the statistics collection for cloud per VDisk.

| *Table 20. Statistics collection for cloud per VDisk* | | | |
|---|---|---|---|
| **SNo** | **Statistic name** | **Acronym** | **Description** |
| 1 | blocks up | bup | Number of blocks that are uploaded in cloud. |
| 2 | blocks down | bdn | Number of blocks that are downloaded from cloud. |

**Note:** A block is 512 bytes.

XML formatting information

The XML is more complicated now, as seen in this raw XML from the volume (Nv_statistics) statistics. Notice how the names are similar but because they are in a different section of the XML, they refer to a different part of the VDisk.

```
<vdsk idx="0"
ctrs="213694394" ctps="0" ctrhs="2416029" ctrhps="0"
ctds="152474234" ctwfts="9635" ctwwts="0" ctwfws="152468611"
ctwhs="9117" ctws="152478246" ctr="1628296" ctw="3241448"
ctp="0" ctrh="123056" ctrhp="0" ctd="1172772"
ctwft="200" ctwwt="0" ctwfw="3241248" ctwfwsh="0"
ctwfwshs="0" ctwh="538" cm="13768758912876544" cv="13874234719731712"
gwot="0" gwo="0" gws="0" gwl="0"
```

```
id="Master_iogrp0_1"
ro="0" wo="0" rb="0" wb="0"
rl="0" wl="0" rlw="0" wlw="0" xl="0">
Vdisk/Volume statistics
<ca r="0" rh="0" d="0" ft="0"
wt="0" fw="0" wh="0" ri="0"
wi="0" dav="0" dcn="0" pav="0" pcn="0" teav="0"  tsav="0"  tav="0"
pp="0"/>
```

```
<cpy idx="0">
```

```
volume copy statistics
<ca r="0" p="0" rh="0" ph="0"
d="0" ft="0" wt="0" fw="0"
wh="0" pm="0" ri="0" wi="0"
dav="0" dcn="0" sav="0" scn="0"
pav="0" pcn="0" teav="0"  tsav="0"
tav="0"  pp="0"/>
```

```
</cpy>
<vdsk>
```

The `<cpy idx="0">` means it is in the volume copy section of the VDisk, whereas the statistics shown under `Vdisk/Volume statistics` are outside of the `cpy idx` section and therefore refer to a VDisk/volume.

Similarly, the following text is the output for the volume cache statistics for node and partitions:

```
<uca><ca dav="18726" dcn="1502531" dmx="749846" dmn="89"
sav="20868" scn="2833391" smx="980941" smn="3"
pav="0" pcn="0" pmx="0" pmn="0"
wfav="0" wfmx="2" wfmn="0"
rfav="0" rfmx="1" rfmn="0"
pp="0"
hpt="0" ppt="0" opt="0" npt="0"
apt="0" cpt="0" bpt="0" hrpt="0"
/><partition id="0"><ca dav="18726" dcn="1502531" dmx="749846" dmn="89"
fav="0" fmx="2" fmn="0"
dfav="0" dfmx="0" dfmn="0"
dtav="0" dtmx="0" dtmn="0"
pp="0"/></partition>
```

This output describes the volume cache node statistics where `<partition id="0">` indicates that the statistics are described for partition 0.

The following text shows the cache statistics for data reduction pools and volume copy cache statistics nodes and partitions:

```
<lca><ca dav="18726" dcn="1502531" dmx="749846" dmn="89"
sav="20868" scn="2833391" smx="980941" smn="3"
pav="0" pcn="0" pmx="0" pmn="0"
wfav="0" wfmx="2" wfmn="0"
rfav="0" rfmx="1" rfmn="0"
pp="0"
hpt="0" ppt="0" opt="0" npt="0"
apt="0" cpt="0" bpt="0" hrpt="0"
/>
<dca p="2089792" rh="305754" ph="178873" d="0"
ft="0" wt="0" fw="0" wh="0"
v="10348585" m="3334742" pm="1120" ri="10720"
wi="0" r="3923240" dav="0" dcn="0"
sav="59926" scn="6045" pav="48350" pcn="2723"
teav="0"  tsav="0"  tav="0"  pp="0"/>
<rca p="2089792" rh="305754" ph="178873" d="0"
ft="0" wt="0" fw="0" wh="0"
v="10348585" m="3334742" pm="1120" ri="10720"
wi="0" r="3923240" dav="0" dcn="0"
sav="59926" scn="6045" pav="48350" pcn="2723"
teav="0"  tsav="0"  tav="0"  pp="0"/>
<jca p="2089792" rh="305754" ph="178873" d="0"
ft="0" wt="0" fw="0" wh="0"
v="10348585" m="3334742" pm="1120" ri="10720"
wi="0" r="3923240" dav="0" dcn="0"
sav="59926" scn="6045" pav="48350" pcn="2723"
```

```
teav="0"  tsav="0"  tav="0"  pp="0"/>
</partition>
```

# Notifications

The system supports sending notifications for remote monitoring. The system supports sending event log entries using Simple Network Management Protocol (SNMP). The system also supports sending event log entries, audit log entries, and authentication attempts using syslog.

## Syslog notifications

The syslog protocol is a standard protocol for forwarding log messages from a sender to a receiver on an IP network. The system can send syslog messages that notify personnel about an event. You can set up syslog event notifications with either the management GUI or the command-line interface (CLI).

The system can transmit syslog messages in either expanded or concise format. Servers configured with facility values of 0 - 3 receive syslog messages in concise format. Servers configured with facility values of 4 - 7 receive syslog messages in fully expanded format. The default value is 0. The facility number that is used in syslog messages also identifies the origin of the message to the receiving server. You can use a syslog manager to view the syslog messages that the system sends. For error, warning, and information notifications the format that messages are sent in depends on the facility setting. Audit (-audit) and authentication (-login) messages are sent in a single format so for these messages there is no distinction between concise and expanded format. The system supports both TCP and UDP transmission protocols to send the syslog message to the specified syslog servers. You can specify up to a maximum of six syslog servers with either an IP address or a fully qualified domain name and its corresponding port. The default port for the TCP protocol is port 6514, and the default port for UDP transmissions is 514. If you are using a domain name to identify a syslog server, ensure that a DNS server is configured on the system. Domain names cannot exceed 40 characters.

The system supports the following syslog notifications and message types:

**Error notifications**
Select this option to send error notifications that can indicate a serious problem with the system.

**Warning notifications**
Select this option to send warning notifications that can indicate a problem or unexpected condition with the system. Always immediately investigate this type of notification to determine the effect that it might have on your operation, and make any necessary corrections.

**Information notifications**
Select this option to send information messages that indicate an expected operation has completed on the system.

**Audit log messages**
Select this option to include any CLI or management GUI operations on the specified syslog servers.

**Authentication log messages**
Select this option to send successful and failed authentication attempts to the specified syslog servers.

## Using the management GUI

To configure or work with syslog notification settings in the management GUI, select **Settings** > **Notifications** > **Syslog**.

## Using the Command Line interface

The following commands can be used to configure and manage syslog notifications:

1. See **mksyslogserver** command to create a syslog notification.
2. See **chsyslogserver** command to modify a syslog notification.
3. See **rmsyslogserver** command to delete a syslog notification.

4. See **lssyslogserver** command to display a concise list of syslog notification.

## SNMP notifications

Simple Network Management Protocol (SNMP) is a standard protocol for managing networks and exchanging messages. The system can send SNMP messages that notify personnel about an event. You can use an SNMP manager to view the SNMP messages that the system sends. The system supports both SNMP version 2 and version 3.

### About this task

Event notifications are reported to the SNMP destinations of your choice. You can specify SNMP destinations by creating SNMP server objects on the system. You must provide a valid IP address or fully qualified domain name. A DNS server must be created on the system to use a fully qualified domain name. A maximum of six SNMP destinations can be specified.

You can use the Management Information Base (MIB) file for SNMP to configure a network management program to receive SNMP messages that are sent by the system. This file can be used with SNMP messages from all versions of the software. You can use your browser to download the MIB file by clicking **Download MIB** in the management GUI by navigating to **Settings** > **Notifications** > **SNMP**.

For version 2 SNMP servers, the community string is required and the default value is public. SNMP version 3 introduces improved encryption and authentication mechanisms. The system supports two different security models when using version 3. The first security model is the User-based Security Model (USM), which provides authentication and privacy at the message level. The second security model is the Transport Security Model (TSM), which uses Transport Layer Security (TLS) to send messages over a secure connection. For more information about using TLS, see . Table 1 displays the security levels that are supported when using SNMP version 3 with USM.

These different security levels depend on the credentials that you configure to authenticate successfully to the SNMP server.

*Table 21. Security levels and required credentials for SNMP version 3 servers using the user-based security model*

| Security Level | Description | Required security credentials |
|---|---|---|
| None | No additional authentication or encryption is used to send SNMP notifications. | • Engine ID<br>• Security Name |
| Authentication | Notifications are authenticated, but message content is not encrypted. | • Engine ID<br>• Security Name<br>• Authentication Protocol<br>• Authentication Passphrase |
| Authentication and Privacy | Notifications are authenticated and message content is encrypted. | • Engine ID<br>• Security Name<br>• Authentication Protocol<br>• Authentication Passphrase<br>• Privacy Protocol<br>• Privacy Passphrase |

For SNMP Version 3, use a unique engine ID for every system. The system does not automatically generate the engine ID. You can use an existing engine ID from an SNMP server. If you are setting up SNMP for the first time, you must generate a new engine ID.

**Note:**

- SNMP will not accept any configuration get requests from SNMP server.

**Using the management GUI**

1. To configure or manage SNMP servers in the management GUI, select **Settings** > **Notifications** > **SNMP**.
2. To configure a new server, select **Add SNMP Server** or right-click an existing instance and select **Modify**.
3. To send a test trap from the system to an SNMP server, right-click an existing instance and select **Test**. The system cannot provide any feedback about whether the test is a success, but you can check on the specified server that the trap has arrived successfully.

**Using the Command Line interface**

To configure the SNMP server, use the following commands:

- See **mksnmpserver** to add an SNMP Version 2 or Version 3 server.
- See **chsnmpserver** to change the settings of an existing SNMP server.
- See **rmsnmpserver** to remove an existing SNMP server from the system.
- See **lssnmpserver** to display either a concise list or a detailed view of the SNMP servers that are detected by the system.
- See **testsnmpserver** to send a test trap to an existing server.

### *SNMP over TLS*

The system supports SNMP v3 servers that use the Transport Security Model (TSM). Traps are sent to the SNMP server by using Transport Layer Security (TLS) and the Transmission Control Protocol (TCP). The system does not support DTLS (or TLS over UDP).

One benefit of using TLS and TCP is that the TCP protocol handles timeouts and retries to improve delivery of traps. This can be compared to the User-based Security Model (USM) where traps are sent from the system by using UDP, with no confirmation that a trap has reached the SNMP server. Using SNMP over TLS may be preferable if your organization already has an existing public-key infrastructure in place.

**Note:** IPv6 is not supported with SNMP over TLS. Ensure that the system and SNMP servers are using IPv4 addresses.

## Before you begin

The system and the SNMP server use mutual TLS to establish a secure connection. The system must verify the SNMP server certificate, and the SNMP server must verify the certificate that is presented by the system.

The SNMP server certificate must be installed in a trust store on the system, with the snmp tag turned on. If the SNMP server's certificate is signed by a Certificate Authority (CA), then the CA certificate must be installed in a trust store with the snmp tag that is turned on.

**Note:** The system only supports the use of one CA to sign the server certificates. All server certificates must be signed by the same CA. If using self-signed certificates, then all servers must use the same certificate.

The system's certificate must be installed on the SNMP server. If the system's certificate is signed by a Certificate Authority (CA), then the CA certificates must be installed on the SNMP server too.

**Note:**

The system uses the same certificate for all services that use certificate authentication. If changes are made to the system certificate, then services that use certificate authentication may be interrupted. If any services are interrupted, add the new certificate to the necessary trust stores.

Some SNMP managers may have requirements about details that must be included in the system certificate. The certificate may need to contain a username in a specific field (such as the Common

Name field). The cluster's fully qualified domain name, and the cluster IP address, should be included in the subject alternative name fields. See for more details about generating a new system certificate.

The system and the SNMP server must agree to use a cipher suite that is supported by both parties. See for more information about changing the system's list of supported cipher suites.

## Creating a new SNMP server that uses TLS using the management GUI

1. In the management GUI, select **Setting** > **Notifications** > **SNMP**.
2. To configure a new server, select **Add SNMP Server** and select the TLS checkbox. When the TLS checkbox is selected, the **SNMP certificate** box appears.
3. If the SNMP server's certificate is signed by a CA, upload the root CA certificate. If the server's certificate is self-signed, upload the self-signed certificate.

   **Note:**

   If the server certificate is signed by a chain of CAs that includes a root CA and intermediate CAs, then the server should be configured to present its server certificate and any intermediate CA certificates when establishing a connection.

   When you create an SNMP server that uses TLS, the management GUI creates a new trust store that contains the uploaded server certificates.

## Testing the SNMP server and resolving common problems

1. In order to testing the SNMP server and resolving common problems, select **Settings** > **Notifications** > **SNMP**, right-click a server and select **Test** to send a test trap to an SNMP server. When using TLS, this checks that the TLS connection can be established, and sends a trap to the server.
2. The SNMP log on the server should be examined to verify that the test trap has arrived. If the test returns an error, or the test trap has not arrived on the server, check the SNMP log on the SNMP server for the following issues:
   a. If the log shows SSL or TLS errors at the time of the test trap, check that the correct certificates have been installed and trusted on the system and on the SNMP server. Also check that the system and the SNMP server are using compatible cipher suites.
   b. If there are no errors at the time of the test trap, check that the system can ping the SNMP server. See **`ping`** for information about running the command. If the system can ping the SNMP server, check that the system certificate contains any required usernames in the correct field of the certificate. Check that the user is configured correctly on the SNMP server.

### *Examples of configuration using net-snmp*

The following demonstrates the configuration settings used to create a connection to SNMP servers using Net-SNMP on Linux.

In this example, the system certificate and the SNMP server certificate are both signed by the IBM Example CA. The storage system certificate contains the username flashsystem in the `Common Name` field. See System Certificates for more details about generating the system certificate.

## Configuring the server

A configuration file must be at a location accessible to snmptrapd, and the following example illustrates this configuration file at `/usr/local/share/snmp/snmptrapd.conf` location:

```
# Logging Settings
[snmp] logOption f /var/log/SNMP/snmptraps.log
format 1 "%02.2h:%02.2j TRAP%w.%q from %B\n"
format 2 "%02.2h:%02.2j TRAP%w.%q from %B\n"
authCommunity log public
# TLS Settings
[snmp] localCert snmpd.crt
[snmp] tlsMinVersion tls1_2
[snmp] tlsMaxVersion tls1_3
```

```
[snmp] trustCert ibm-example-ca
certSecName 10 flashsystem.crt --cn flashsystem
authUser log -s tsm flashsystem
```

- snmpd.crt is the signed server certificate and added to `/usr/local/share/snmp/tls/certs/snmpd.crt` location.

  The subject alternative name field should contain the fully qualified domain name and IP address of the SNMP server.

- The private key that matches snmpd.crt is added to `/usr/local/share/snmp/tls/private/snmpd.key` location.

- ibm-example-ca is the name of the root CA used to sign both snmpd.crt and flashsystem.crt. The root certificate has been added to `/usr/local/share/snmp/tls/ca-certs/ibm-example-ca.crt`

  If a chain of CA certificates is used in the signing process, then all of the CA certificates should be added to this directory, and an entry for each CA should be added to the configuration file.

- flashsystem.crt is the signed system certificate that was exported from the **Security** > **System Certificates** panel on the storage system, and added to `/usr/local/share/snmp/tls/certs/flashsystem.crt`. --cn flashsystem indicates that the username flashsystem is located in the Common Name field of the certificate.

## Configuring the storage system

1. In this example we are signing the SNMP server's certificate and the storage system's certificate with our organization's CA (The IBM Example CA). If the system's certificate needs to be configured, use the **Security** > **System Certificates** panel in the management GUI to generate a new certificate request.
2. In this example, we have added flashsystem to the Common Name field as username. We have also added the system's fully qualified domain name, and the cluster IP, to the subject alternative name fields. The certificate request is then signed by the IBM Example CA, and the signed certificate is installed in **Security** > **System Certificates**.
3. In the management GUI, navigate to **Settings** > **Notifications** > **SNMP**.
4. Select **Add SNMP Server**.
5. Enter the IP and port details for the server.
6. Select the TLS checkbox.
7. In the SNMP certificate box, upload the root CA certificate used to sign the SNMP server certificate. In this example, we are using ibm-example-ca.crt.

Now that the SNMP server and the storage system have both been configured, we can start snmptrapd on the server and have it listen for TLS connections. We are using the default trap port 10162

```
snmptrapd -c /usr/local/share/snmp/snmptrapd.conf -L o tlstcp:10162
```

We can now send a test trap to the server in **Settings** > **Notifications** > **SNMP** by right-clicking the server and selecting `Test`. In our example, the trap is logged in `/var/log/SNMP/snmptraps.log` on the server.

# Security

IBM Storage Virtualize based storage systems are secure storage platforms that implement various security-related features for both system-level security and data-level security.

## Security overview

IBM Storage Virtualize based storage systems are secure storage platforms that implement various security-related features for both system-level security and data-level security.

The security features are broadly categorized as *System security* and *Data security*. These security features protect and prevent unauthorized access and use of the system, its resources, and the data that is stored on the system.

IBM Storage Virtualize provides the following security features on the storage systems.

### System Security

System security describes the controls that protect the system and its resources from both internal and external disruption. In general, these features prevent unauthorized access to system resources and provide event notifications and alerts that warn security administrators of any unauthorized access attempts. IBM Storage Virtualize systems support the following system security features:

**User authentication**

The system supports both local users, and remote users who are authenticated to the system through a remote authentication service. You can create local users who can access the system. These user types are defined based on the administrative privileges that they have on the system. Local users must provide either a password, a Secure Shell (SSH) key, or both. Local users are authenticated through the authentication methods that are configured on the system. If the local user needs access to the management GUI, a password is needed for the user. If the user requires access to the command-line interface (CLI) through SSH, either a password or a valid SSH key file is necessary. Local user passwords are securely stored by using the PBKDF2 hashing algorithm. Local users must be part of a user group that is defined on the system. User groups define roles that authorize the users within that group to a specific set of operations on the system. For more information, see Configuring user authentication.

**Remote authentication**

Remote authentication allows users to authenticate to the system using credentials that are stored on an external authentication service. When you configure remote authentication, you do not need to configure users on the system or assign more passwords. Instead, you can use your existing passwords and user groups that are defined on the remote service to simplify user management and access to enforce password policies more efficiently, and to separate user management from storage management. For more information, see “Configuring remote authentication” on page 89.

**Role-based access control**

Each user of the management GUI must provide a username and a password to sign on. Each user also has an associated role, such as monitor or security administrator. These roles are defined at the system level. For example, a user can be the administrator for one system, but the security administrator for another system. For more information, see User roles.

**Default user**

When a system is created, a single local user with Security Administrator privileges, called superuser, is created. The superuser contains maximum privileges to complete system setup and configuration. For new systems, the default superuser password must be changed on first login to the system. Although the superuser cannot be deleted, you may want to lock or disable the superuser to prevent access to the system. For more information, see Locking user accounts.

**Object-based access control**

An *ownership group* defines a subset of users and objects within the system. You can create ownership groups to further restrict access to specific resources that are defined in the ownership group. Only users with Security Administrator roles can configure and manage ownership groups.

Restricted users are those users who are defined to a specific ownership group and can only view or manage specific resources that are assigned to that ownership group. Unrestricted users are not defined to an ownership group and can manage any objects on the system based on their role on the system. For more information on configuring an ownership group, see Configuring ownership groups.

**Login interfaces**

The system provides several management interfaces that allow user to authenticate and manage the system and its objects. These interfaces include the management GUI, command-line interface, service assistant interface and commands, and REST APIs. All interfaces use in-flight encryption for data to secure the login and all subsequent connections with system. You can create a user group with specific roles to allow or restrict the user group to interfaces. With this capability, you can also control access from scripts or automated services, as well as individual users. For more information, see Changing user groups.

**Password Policy**

With password policy support, system administrators can set security requirements that are related to password creation and expiration, timeout for inactivity, and actions after failed logon attempts. For more information, see “Password policy” on page 154.

**Account locking**

The Security Administrator can manually lock or unlock a user account at any time. By default, the superuser is exempt from the system-wide policy for manual or automatic account locking. For more information, see Locking user accounts.

**Session timeouts**

You can configure session timeout for both the management GUI and the CLI. Session timeouts automatically signs a user out of a session if it has been idle for a specified amount of time. The timeout values can be 5 - 240 minutes. For more information, see the **chsecurity** command.

**Login banner**

You can create or change a message that displays when users log on to the system. When users log on to the system with the management GUI, command-line interface, or service assistant, the message displays before they log on to the system. For more information, see Changing the login message.

**Multifactor authentication**

Multifactor authentication requires users to provide multiple pieces of information when they log in to the system to prove their identity. Multifactor authentication uses any combination of two or more methods, called *factors*, to authenticate users to your resources and protect those resources from unauthorized access. For more information, see “Multifactor authentication” on page 92.

**Single sign-on**

Single Sign-on (SSO) authentication requires users to register their credentials only once when the user signs on to the application for the first time. The user information is stored at the Identity Provider (IdP) that manages the user credentials and determines whether the user is required to authenticate again or not. For more information, see “Single Sign-on” on page 113.

**Secured IP partnership**

*Secured IP partnerships* secure the data as it travels through an untrusted network between production and recovery systems. Secured IP partnerships minimize the risk of hackers manipulating or intercepting data in untrusted networks as replicated data travels between partnered systems. For more information, see “Planning secured IP partnerships” on page 17.

**Auditing and reporting**

The system includes an internal tamper-proof audit log that traces all successful commands and identifies the user who issued the commands. This information includes the user details, the IP address where they connected to the system, and timestamps. For more information, see **Audt log commands**.

**Secure and trusted boot**

Secure boot encrypts the file systems and relies on a hardware root of trust that extends all the way through to the operating system and **initrd** to unlock the file system. The system supports hardware root of trust and secure boot operations, which protects against unauthorized physical access to the hardware and prevents malicious software from running on the system. For more information, see Secure boot.

**Secure sockets and Secure Shell settings**

Secure sockets (SSL/TLS) and Secure Shell (SSH) are used to establish secure connection to management interfaces, such as the management GUI, CLI, and REST APIs. These security protocols are also used to authenticate to remote servers, such as email, LDAP, and key management servers. The system supports different levels of these security protocols that define the cipher suites and key exchange algorithms that can be used. For more information, see "Changing security protocol levels" on page 156.

**SSL/TLS certificates**

During system initialization and set up an initial internally signed certificate is created to secure connections between nodes on the system and a supported web browser for management GUI access. After system setup, you can either create internally signed certificate from the native certificate authority or generate a certificate signing request to an external, third-party CA. For more information, refer to Managing certificates for secure communications.

**Disabling USB ports**

This feature can be combined with encryption key management on local USB flash drives to keep the ports disabled during day-to-day use and only reenabled when a key must be provided. For more information, see the **satask chnodeusb** command.

**Network Time Protocol**

To mitigate against attacks that are based on the system time/date being out of sync with other services, the system supports a configurable Network Time Protocol (NTP) server to control the system time. An NTP server can be configured only by a user with the Security Administrator role. For more information, see the **chsystem** command.

**IP address and network port allocation**

As you plan your installation, you must consider IP address requirements and service access for the system. For more information, see IP address allocation and usage.

**Internal protection against installing unauthorized software**

The IBM Storage Virtualize software runs on customized Linux installation that uses a bare minimum of carefully selected packages. The kernel configuration is locked and tightly controlled so that only software that is required for the hardware and software interaction is installed.

**Preventing access as root**

The management software runs as an internal Linux user without root privileges, preventing unauthorized users from accessing the system. Although external users cannot access this software, an attacker might use a weakness to compromise the system. Even in this circumstance, they do not have root privileges.

**File system protection**

The secure boot encrypts the file systems and relies on a hardware root of trust. The boot drives are tied to the Trusted Platform Module (TPM). The checks and validation of the file systems at boot time provide security against files being corrupted maliciously or by hardware or software faults. Updates to system software can only be installed with built-in update technology. All software packages are validated and signed by a private key that is on all IBM servers. This configuration ensures that the tight control on software is maintained for all systems that are deployed at customers.

**Disabling service assistant password reset**

The superuser is the only user account that is permitted to run service assistant commands on the system. You can use the command-line interface (CLI) to view and change the status of the password reset feature for the system.

## Data security

Data security protects the data that is stored on the system against theft, loss, or attack.

**Encrypting data at rest**

All NVMe drives that are supported by the system, including IBM FlashCore® Modules (FCMs) and a range of other third-party drives, are self-encrypting drives (SEDs) that encrypt data within the electrical circuit of each individual drive.

**Secure data deletion**

The system provides methods to securely erase data from a drive or from a boot drive when a node or node canister is decommissioned. For more information, refer to Secure data deletion.

**Volume protection**

Volume protection prevents active volumes or host mappings from being deleted inadvertently if the system detects recent I/O activity. This global setting is enabled by default on new systems. You can either set this value to apply to all volumes that are configured on your system, or control whether the system-level volume protection is enabled or disabled on specific pools. For more information, see “Volume protection ” on page 46.

**Safeguarded Copy function**

Safeguarded Copy function supports the ability to create cyber-resilient point-in-time copies of volumes that cannot be changed or deleted through user errors, malicious actions, or ransomware attacks. Safeguarded snapshots are supported on the system through an internal scheduler that is defined in the snapshot policy. When the policy is assigned to a volume group, you can select Safeguarded option. The policy creates immutable snapshots of all volumes in the volume group. The system supports internal and external snapshot scheduling applications such as IBM Copy Services Manager and IBM Spectrum Copy Data Management.

For more information, refer to the Safeguarded Copy function.

**Logical port isolation**

The system provides means to define logical port sets. These port sets can be used to further restrict the login traffic from a host, or set of hosts, to isolate traffic to specific SAN paths. *Portsets* are groupings of logical addresses that are associated with the specific traffic types. The system supports both Fibre Channel and IP portsets for host attachment, IP portsets for backend storage connectivity, and IP replication traffic. The system supports a maximum of 72 portsets. For more information, refer to Portsets.

**Encrypting data in flight**

If you use secured IP partnerships to secure connections between partnered systems, you also require an encryption license. If you have not purchased a license, contact a customer representative to purchase an encryption license. For more information, refer to Configuring encryption.

# Configuring remote authentication

Remote authentication allows users to authenticate to the system using credentials that are stored on an external authentication service. When you configure remote authentication, you do not need to configure users on the system or assign more passwords. Instead, you can use your existing passwords and user groups that are defined on the remote service to simplify user management and access to enforce password policies more efficiently, and to separate user management from storage management.

A remote user is authenticated on a remote LDAP server. A remote user does not need to be added to the list of users on the system, although they can be added to configure optional SSH keys. For remote users, an equivalent user group must be created on the system with the same name and role as the group on the

remote LDAP server. Remote users cannot access the system when the remote LDAP server is down. In that case, a local user account must be used until the LDAP service is restored. Remote users have their groups that are defined by the remote authentication server.

## Using the management GUI

To configure remote authentication with LDAP, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Remote Authentication**.
2. Select **Configure Remote Authentication**.
3. Select **LDAP**.
4. Select the type of LDAP server that is used for authentication.
5. Select one of the following security options:

   **LDAP with StartTLS**
   Select this option to configure extensions that upgrade the standard LDAP port (389) to an encrypted port that uses TLS or SSL. The initial connection to the directory server is decrypted but can be used on systems that do not have port 636 available.

   **LDAPS**
   Select this option to secure LDAP communication by using the default secure port (636). The connections for all transactions with the directory server are encrypted.

   **LDAP with no security**
   Select this option to transport data in clear text format without encryption.
6. Specify optional service credentials or modify advanced LDAP settings. The following LDAP attributes can be configured:

   **User attribute**
   For all server types, users are authenticated with a username that is defined with the LDAP user attribute. This attribute must exist in your LDAP schema and must be unique for each of your users. Active Directory users can also authenticate by using their user principal names (UPN) or NT login names.

   **Group attribute**
   Authenticated users are assigned roles according to their LDAP group memberships. The groups to which a user belongs are stored in the LDAP group attribute. This attribute value can be the distinguished name of each group, or a colon-separated list of user group names.

   **Audit log attribute**
   If an LDAP user completes an audited action, the contents of the audit log attribute are recorded in the audit log.
7. Click **Next**.
8. Define up to six LDAP servers to use for authentication. Multiple servers can be configured to provide access to different sets of users for redundancy. You can also configure which servers are preferred to authenticate users. You can specify either IP addresses or domain name for the LDAP servers. If you specify a fully qualified domain name, a DNS server must be configured on your system. To configure a DNS server for the system, select **Settings** > **Network** > **DNS**.

   **Note:** The system does not support using LDAP referrals to find related LDAP servers. Each required LDAP server must be explicitly configured on the system.
9. Configure user groups on the system to match groups that are configured on the remote authentication service. For each user group on the authentication service, a corresponding user group must be created with the same name. In the management GUI, complete these steps:

   a. Select **Access** > **Users** > **Create User Groups**.

   b. On the **Create User Group** page, enter the following information:

**Group Name**
Enter the name of the group that is on the remote LDAP server. The name of the group on the system must match.

**Ownership Group**
If ownership groups are configured on your system, you can select an ownership group for the user group.

**Remote Authentication**
Select **LDAP**.

**Multifactor authentication**
Select **On** to enable second-factor authentication for remote users on the system. These users authenticate with the first factors that are stored on the remote LDAP server and then are required to provide a second factor to access the system through a supported authentication service.

c. Click **Create**.

d. Remote users are defined on the remote authentication service. Only remote users who require access to the command-line interface using a Secure Shell (SSH) key, need to be created. To create remote users, who require access to the command-line interface, select **Access** > **Users by Groups** > **Create User**.

e. On the **Create User** page, enter the following information:

**Name**
Enter a name for the user.

**Authentication mode**
Select **Remote**.

**SSH key**
For remote users who require to access the system, select the public SSH key which stored locally on the system.

f. Click **Create**. Repeat these steps for all remote users.

10. Verify your LDAP configuration. To test authentication to the LDAP servers, select **Test LDAP Authentication** and enter corresponding credentials for the user.

## Using the command-line interface

To enable user authentication with LDAP by using the command-line interface, follow these steps:

1. Configure LDAP by entering the **chldap** command.

   This command provides default settings for both IBM Security Directory Server and Microsoft Active Directory. To configure authentication with IBM Security Directory Server schema defaults and Transport Layer Security (TLS), for example, enter the following command:

   ```
   chldap -type itds -security tls
   ```

   LDAP configuration can be inspected with the **lsldap** command.

   **Note:** Use TLS so that transmitted passwords are encrypted.

2. Specify the **mkldapserver** command to define up to six LDAP servers to use for authentication.

   Multiple servers can be configured to provide access to different sets of users or for redundancy. All servers must share the settings that are configured with **chldap**.

   **Note:** The system does not support using LDAP referrals to find related LDAP servers. Each LDAP server that is required must be explicitly configured on the system.

   To configure an LDAP server with an SSL certificate and users in the `cn=users,dc=company,dc=com` subtree, for example, enter the following command:

   ```
   mkldapserver -ip 9.71.45.108 -basedn cn=users,dc=company,dc=com -sslcert /tmp/sslcert.pem
   ```

If you have a DNS server configured on the system, you can also specify a domain name in the **-ip** parameter. For example:

```
mkldapserver -ip myldap.myco.com -basedn cn=users,dc=company,dc=com -sslcert /tmp/sslcert.pem
```

You can also configure which servers are preferred to authenticate users.

Specify **lsldapserver** for LDAP server configuration information. Specify **chldapserver** and **rmldapserver** to change the configured LDAP servers.

3. Configure user groups on the system by matching those user groups that are used by the authentication service.

   For each group of interest that is known to the authentication service, a system user group must be created with the same name and with the remote setting enabled. If members of a group that is called `sysadmins`, for example, require the system administrator (admin) role, enter the following command:

   ```
   mkusergrp -name sysadmins -remote -role Administrator
   ```

   If none of the user groups match a system user group, the user cannot access the system.

4. Verify your LDAP configuration by using the **testldapserver** command.

   To test the connection to the LDAP servers, enter the command without any options. A user name can be supplied with or without a password to test for configuration errors. To process a full authentication attempt against each server, enter the following commands:

   ```
   testldapserver -username username -password 'password'
   ```

5. Enter the following command to enable LDAP authentication:

   ```
   chauthservice -type ldap -enable yes
   ```

6. Configure users who do not require Secure Shell (SSH) key access.

   Delete system users who must use the remote authentication service and do not require SSH key access.

   **Remember:** A superuser cannot be deleted or use the remote authentication service.

7. Configure users who require SSH key access.

   All system users who use the remote authentication service and require SSH key access must have remote settings that are enabled and a valid SSH key that is configured on the system.

8. Specify the type of security to use when communicating with LDAP servers.

   Specify *tls* to enable TLS. Select this option to configure extensions that upgrade the standard LDAP port (389) to an encrypted port that uses TLS. The initial connection to the directory server is unencrypted but can be used on systems that do not have port 636 available.

   Specify *ssl* to enable SSL security. This option secures LDAP communication by using the default secure port (636). The connections for all transactions with the directory server are encrypted. The default value is none.

## Multifactor authentication

Multifactor authentication requires users to provide multiple pieces of information when they log in to the system to prove their identity. Multifactor authentication uses any combination of two or more methods, called *factors*, to authenticate users to your resources and protect those resources from unauthorized access.

One of the key concepts of multifactor authentication is each factor comes from a different category. These categories include the following:

**Something a user knows**
Users authenticate with information that only each individual user knows, such as a password or PIN.

**Something a user has**
The users prove their identity with information that is given to the user by a trusted authentication service, such as one-time passcodes that are generated by an application or mobile device.

**Something a user is**
Users prove their identity with biometrics, such as fingerprint or retinal scans.

With the adoption of cloud-based services, multifactor authentication increases the control over user access and security settings. First-factor authentication methods alone, such as username and password combinations, do not provide the level of protection and security that is required in cloud and hybrid-cloud environments. With multifactor authentication support, security administrators can reinforce account protection, create granular access for users and user groups, and monitor access more efficiently at a system level.

## Multifactor authentication with IBM Security Verify

With IBM Security Verify, security administrators can configure the system as an application that requires two factors for users and user groups to access the system with either the management GUI or CLI.

Multifactor authentication can be used to protect both local users, including superuser, and remote users.

Remote users are users who are defined on a remote LDAP server. For remote users that authenticate with LDAP servers, install and configure IBM Security Verify Bridge for Directory Sync on your LDAP server, such as Windows Active Directory. IBM Security Verify Bridge for Directory Sync duplicates any users and groups that are defined on the source LDAP server into the Cloud Directory in IBM Security Verify. Any subsequent changes that are made to the source LDAP server are copied automatically to the Cloud Directory in IBM Security Verify. For more information, see IBM Security Verify Bridge for Directory Sync in the IBM Security Verify documentation.

## Multifactor authentication with Duo Security

The system integrates with IBM Security Verify or Duo Security , which are cloud-based identity and access management (IAM) service providers. These services provide different factors to validate and verify users who access the system.

You can manage every aspect of your Duo two-factor authentication system from the Duo Admin Panel including creating and managing applications, enrolling and activating users, issuing and managing SMS passcodes and bypass codes, managing mobile devices, fine-tuning the user experience of your Duo installation.

IBM Security Verify configures the management GUI and the command-line interface as separate API clients that require separate credentials. For GUI-based logins, the system communicates with IBM Security Verify through the OpenID Connect (OIDC) protocol.

**Related information**
“Configuring multifactor authentication with IBM Security Verify” on page 94
The system integrates with IBM Security Verify to provide multifactor authentication for system users.

“Configuring user groups for multifactor authentication with IBM Security Verify” on page 99
After you configure multifactor authentication on the system, you must enable multifactor authentication for user groups and add users to those groups in IBM Security Verify.

IBM Security Verify documentation

“Configuring multifactor authentication with Duo Security” on page 108
The system integrates with Duo Security to provide multifactor authentication for system users.

“Configuring user groups for multifactor authentication with Duo Security” on page 110

After you configure multifactor authentication on the system, you must enable multifactor authentication for user groups and add users to those groups in Duo Security.

Duo Security documentation

## Configuring multifactor authentication

Multifactor authentication requires users to provide multiple pieces of information when they log in to the system to prove their identity. Multifactor authentication uses any combination of two or more methods, called *factors*, to authenticate users to your resources and protect those resources from unauthorized access.

With the adoption of cloud-based services, multifactor authentication increases the control over user access and security settings. First-factor authentication methods alone, such as username and password combinations, do not provide the level of protection and security that is required in cloud and hybrid-cloud environments. With multifactor authentication support, security administrators can reinforce account protection, create granular access for users and user groups, and monitor access more efficiently at a system level.

The system integrates with IBM Security Verify or Duo Security , which are cloud-based identity and access management (IAM) service providers. These services provide different factors to validate and verify users who access the system.

### *Configuring multifactor authentication with IBM Security Verify*

The system integrates with IBM Security Verify to provide multifactor authentication for system users.

With IBM Security Verify, security administrators can configure the system as an application that requires two factors for users and user groups to access the system with either the management GUI or CLI.

Multifactor authentication can be used to protect both local users, including superuser, and remote users.

Remote users are users who are defined on a remote LDAP server. For remote users that authenticate with LDAP servers, install and configure IBM Security Verify Bridge for Directory Sync on your LDAP server, such as Windows Active Directory. IBM Security Verify Bridge for Directory Sync duplicates any users and groups that are defined on the source LDAP server into the Cloud Directory in IBM Security Verify. Any subsequent changes that are made to the source LDAP server are copied automatically to the Cloud Directory in IBM Security Verify. For more information, see IBM Security Verify Bridge for Directory Sync in the IBM Security Verify documentation.

IBM Security Verify configures the management GUI and the command-line interface as separate API clients that require separate credentials. For GUI-based logins, the system communicates with IBM Security Verify through the OpenID Connect (OIDC) protocol.

**Important:**

When the system certificate added as the signer certificate expires, you cannot login to the management GUI if you are using multifactor authentication. To export the new system certificate and install it in IBM Security Verify, login using the CLI.

## Prerequisites

Ensure that the following prerequisite tasks are completed on the system before you configure multifactor authentication IBM Security Verify:

1. Ensure the system is updated to 8.5.0 or later release.
2. Configure a DNS server. To create a DNS server, select **Settings** > **Network** > **DNS**.
3. Configure an HTTP proxy server or configure your firewall to access IBM Security Verify. To create an HTTP proxy server, **Settings** > **Network** > **Internal Proxy Server**. If your system does not directly connect to the Internet, you can create a firewall exception to allow your system access to IBM Security Verify.

   **Important:** If the proxy server is not configured correctly, the system cannot communicate with IBM Security Verify and the login to the system fails.

4. For the management GUI and the command-line interface, ensure that the inactivity logout is equal to or greater than the time it takes for a user to receive a one-time passcode (OTP) from the authentication service. The default value for the inactivity timeout is 30 minutes for the management GUI and 15 minutes for the CLI. To set the inactivity timeout for both interfaces in the management GUI, select **Settings** > **Security** > **Inactivity Logout**.
5. Ensure that the SSH grace time for the command-line interface is equal to or greater than the time it takes for a user to receive a one-time passcode (OTP) from the authentication service. The default value for the SSH grace time period is 60 seconds. To set SSH grace time on the system in the command-line interface, use the **chsecurity -sshgracetime** command.

The following prerequisite steps on IBM Security Verify must be completed before you can configure multifactor authentication on the system:

1. Create a subscription for IBM Security Verify. You need an IBMid to create a subscription. A 90-day free trial subscription is also available. For more information, see Cloud identity and access management (IAM) solutions. During subscription creation, you specify a tenant that is used to create a URL to access the IBM Security Verify dashboard.
2. Access the IBM Security Verify administrator dashboard by entering the following URL in a web browser:

   ```
   https://tenant.verify.ibm.com/ui/admin
   ```

   Where `tenant` is the name of the tenant that you specified when you created your subscription. Usually this tenant name is associated with your company or organization.
3. In the IBM Security Verify interface, select **Applications** > **Applications** > **Add application**.
4. Select **IBM Storage Virtualize** > **Add Application**.

   **Note:** Each system must be added as a separate application.

   The following table shows the required fields and actions for the **General** tab in the IBM Security Verify interface.

| *Table 22. General tab* | |
|---|---|
| **Field** | **Action** |
| Name | Enter a name to identify the system on IBM Security Verify. If you are adding multiple systems, enter a unique name. |
| Description | Enter a brief description of the system. |
| Company name | Name of organization or company. |

   The following table shows the required fields and actions for the **Sign-on** tab in the IBM Security Verify interface. The Sign-on tab is used to add the management GUI as an API-based client.

| *Table 23. Sign-on tab* | | |
|---|---|---|
| **Field** | **Action** | **Details** |
| Application URL | Enter the URL for your system. | Enter the URL that is used to access the management GUI. |
| Grant type | Select **Authorization code** and **JWT bearer**. | Two grant types are required for setting up MFA for the system. **Authorization code** indicates that the client can request access to protected resources on behalf of users. |

| *Table 23. Sign-on tab (continued)* | | |
|---|---|---|
| **Field** | **Action** | **Details** |
| Client ID | This value is automatically generated when the system is saved as an application. | This value must be entered to the Multifactor authentication page in the management GUI under **OpenID Credentials**. |
| Client secret | This value is automatically generated when the system is saved as an application. | This value must be entered to the Multifactor authentication page in the management GUI under **OpenID Credentials**. |
| User consent | Select **Do not ask for consent**. | |
| Redirect URIs | Enter the locations where the authorization server sends users after they are successfully authorized and granted an authorization code or access token. | Multiple redirect URIs can be specified for both the management GUI and the service assistant GUI. For management GUI access, the redirect URI is comprised of the management IP address or hostname followed by `/mfa`. For the service assistant interface, the redirect URI is comprised of the hostname or IP address for the system followed by `service/mfa`. For example:<br><br>https://hostname.com/mfa<br>https://hostname/service/mfa |
| JWT bearer user identification | Select **Username**. | Indicates that the username field in the JWT bearer is used to find users in the Cloud Directory and determines what second factors IBM Security Verify presents to users when they log into the system. |
| JWT bearer default identity source | Ensure **Cloud Directory** is selected. | Indicates that the IBM Security Verify Cloud Directory is used to look up the second factor for the username. After you configure multifactor authentication on the system, users and user groups must be added to the Cloud Directory. |
| Generate refresh token | Ensure that this option is unchecked. | |
| Send all known user attributes in the ID token | Ensure that this option is checked. | |

| *Table 23. Sign-on tab (continued)* | | |
|---|---|---|
| **Field** | **Action** | **Details** |
| Access policies | Complete these steps:<br>a. Deselect **Use default policy**.<br>b. Click the **Edit** icon.<br>c. Select **Always require 2FA in all devices**.<br>d. Click **OK**. | This action creates an access policy which controls the authentication steps for system access. Access policies can specify different authentication requirements based on properties of the user or connection. In this case, all users must complete a second factor authentication every time they access the system from all devices. |
| Restrict custom scopes | Ensure this option is unchecked. | |

The following table shows the required fields and actions for the **API access** tab in the IBM Security Verify. The **API access** tab is used to add the command-line interface as an API-based client and create credentials for multifactor authentication access for CLI users. To add the command-line as a separate API client, click **Add API client**. On the **Add API client** page, enter the following information:

| *Table 24. Add API Client action* | | |
|---|---|---|
| **Field** | **Action** | **Details** |
| Name | Enter a name to identify the command line interface as the API client. | |
| Select the APIs to which you want to grant access | Ensure that all APIs are selected by moving the toggle to display **On** | |

5. Click **Save**. After the system is saved as a new application, the **Custom Application** reloads with the **Entitlements** tab selected.
6. On the **Entitlements** tab, select **Automatic access for all users and groups**.
7. Click **Save**.
8. Select **Applications** and select the application name that represents the system.
9. On the **Sign-on** tab, copy the **Client ID** and the **Client secret**. These values must be specified as the OpenID credentials on the Multifactor authentication page in the management GUI.
10. On the **API access** tab, click the edit icon and copy the **Client ID** and the **Client secret**. These values must be specified as the API Client credentials on the Multifactor authentication page in the management GUI.

## Using the management GUI

To configure multifactor authentication on IBM Security Verify, complete these steps:

1. Select **Settings** > **Security** > **Multifactor Authentication**.
2. Enter the host name and port of the authentication server. For IBM Security Verify, enter the following:

```
tenant.verify.ibm.com
```

Where `tenant` is the name that is associated with your subscription. Port 443 is the default for the authentication server.

3. For the **OpenID Credentials**, add the **Client ID** and **Client Secret** that you copied on from the Sign-on tab in the IBM Security Verify interface.
4. For the **API Client Credentials**, add the **Client ID** and **Client Secret** that you copied on from the API access tab in the IBM Security Verify interface.
5. On the **Multifactor Authentication** page, click **Export Certificate** to export the system certificate to your device. Copy the system ID alias that displays. This value must be used as the friendly name of the certificate in IBM Security Verify.
6. Access the IBM Security Verify administrator dashboard by entering the following URL in a web browser:

   ```
   https://tenant.verify.ibm.com/ui/admin
   ```

   Where `tenant` is the name of the tenant that you specified when you created your subscription. Usually this tenant name is associated with your company or organization.
7. Select **Security** > **Certificates**.
8. Under **Signer certificates**, select **Add signer certificate**.
9. On the **Add signer certificate** page, select **Add file** and navigate to where you exported the certificate on your device.
10. In the **Friendly name** field, copy the system ID alias that displays on the **Multifactor Authentication** page in the management GUI.
11. Click **OK**.
12. Return to the **Multifactor Authentication** page in the management GUI, and click **Save**. On the confirmation page, click **Confirm** to enable multifactor authentication for the system.

Multifactor authentication is enabled for the system. You can configure user groups to use multifactor authentication. Click **Navigate** to launch the **User Groups** page.

## Using the CLI

Before you can enable multifactor authentication on the system, ensure that the system certificate is exported and added as a signer certificate. If the certificate is not added as a signer certificate, users with multifactor authentication enabled cannot sign in to the management GUI. To export and add the system certificate, complete the following steps:

1. To view the system ID alias, enter the following command:

   ```
   lssystem | grep id_alias
   ```

   **Note:** The system ID alias must be entered in the **Friendly name** field on the **Add Signer Certificate** page in the IBM Security Verify interface.
2. Enter the following command to export the system certificate:

   ```
   chsystemcert -export
   ```

   This command exports the system certificate. Download the resulting file `/dumps/certificate.pem` to your machine and upload to IBM Security Verify. Ensure to add the system ID alias to the **Friendly name** field. For more information, see steps 7 through 11 in the management GUI section.
3. To enable multifactor authentication with IBM Security Verify, enter the following command:

   ```
   chauthmultifactorverify -hostname tenant.verify.ibm.com -openidclientid yyyyyyyy-yyyy-
   yyyy-yyyy-yyyyyyyyyyyy -openidclientsecret yyyyyyyyyy -cliclientid xxxxxxx-xxxx-xxxx-xxxx-
   xxxxxxxxxxxx -cliclientsecret xxxxxxxxxx -enable
   ```

In the example, `tenant` is the tenant name that is associated with your subscription. The values for the **`-openidclientid`** and the **`-openidclientsecret`** are the Open ID Client and Open ID Secret that are automatically generated when you created your system as a custom application in IBM Security Verify on the Sign-on tab in the IBM Security Verify interface. The values for the **`-cliclientid`** and the **`-cliclientsecret`** are the API Client ID and API Client Secret that are automatically generated when you created your system as a custom application in IBM Security Verify on the API access tab in the IBM Security Verify interface.

*Configuring user groups for multifactor authentication with IBM Security Verify*
After you configure multifactor authentication on the system, you must enable multifactor authentication for user groups and add users to those groups in IBM Security Verify.

As part of multifactor authentication configuration, you must enable the function per user group. The system supports enabling multifactor authentication for local and remote user groups.

The security administrator must define all local users manually in the Cloud Directory of IBM Security Verify. When users log into the system with multifactor authentication, their username is used to look up the required second factor in the Cloud Directory in IBM Security Verify. Users can set up multiple second factors to avoid getting locked out of the system.

**Note:** When you configure multifactor authentication for the first time, follow these guidelines to avoid getting locked out of the system unintentionally:

- Ensure at least one user with Security Administrator role does not have multifactor authentication enabled. If the security administrator gets locked out of the system because of errors in the multifactor authentication set up, the additional user can still access the system.
- Enable multifactor authentication on a user group without logged in users or leave an SSH session active to avoid locking out users on the system.

For remote users that authenticate with LDAP servers, install and configure IBM Security Verify Bridge for Directory Sync on your LDAP server, such as Windows Active Directory. IBM Security Verify Bridge for Directory Sync duplicates any users and groups that are defined on the source LDAP server into the Cloud Directory in IBM Security Verify. Any subsequent changes that are made to the source LDAP server are copied automatically to the Cloud Directory in IBM Security Verify. For local users, each user must be added manually to the Cloud Directory in IBM Security Verify.

## Using the management GUI

To enable multifactor authentication on user groups, complete these steps:

**For existing user groups:**

1. In the management GUI, select **Access** > **Users by Group**.
2. Select the user group from the left navigation and select **User Group Actions** > **Properties**.
3. On the **User Group Properties** page, select **On** under **Multifactor Authentication** to enable second-factor authentication for all users with the user group. These users authenticate with the first factors that are stored on the local system and then are required to provide a second factor to access the system through a supported authentication service.
4. Click **OK**.

**For new user groups for local users:**

1. In the management GUI, select **Access** > **Users by Group** > **Create User Group**.
2. On the **Create User Group** page, enter the following information:

   **Group Name**
   Enter a name of the user group.

   **Ownership Group**
   If ownership groups are configured on your system, you can select an ownership group for the user group.

**Multifactor Authentication**
Select **On** to enable second-factor authentication for local users on the system. These users authenticate with the first factors that are stored on the local system and then are required to provide a second factor to access the system through a supported authentication service.

**Role**
Select a role that for the user group.

3. Click **Create**.
4. Select **Access** > **Users by Group** > **Create User**.
5. On the **Create User** page, enter the following information:

**Name**
Enter a user name for the user. This user name must match the user name that is added to the Cloud Directory on IBM Security Verify.

**Authentication mode**
Select **Local**.

**User Group**
Select the name of the user group that the local user belongs to.

**Password**
Enter a password that is used as the first factor for management GUI access.

**SSH key**
For CLI users, include a public SSH key that is used as the first factor for CLI access.

6. Click **Create**. Repeat these steps for all local users.

**For new user groups for remote users:**

1. Select **Access** > **Users** > **Create User Groups**.
2. On the **Create User Group** page, enter the following information:

**Group Name**
Enter the name of the group that is on the remote LDAP server. The name of the group on the system must match.

**Ownership Group**
If ownership groups are configured on your system, you can select an ownership group for the user group.

**Remote Authentication**
Select **LDAP**.

**Multifactor authentication**
Select **On** to enable second-factor authentication for remote users on the system. These users authenticate with the first factors that are stored on the remote LDAP server and then are required to provide a second factor to access the system through a supported authentication service.

3. Click **Create**.
4. Remote users are defined on the remote authentication service. Only remote users who require access to the command-line interface using a Secure Shell (SSH) key, need to be created. To create remote users, who require access to the command-line interface, select **Access** > **Users by Groups** > **Create User**.
5. On the **Create User** page, enter the following information:

**Name**
Enter a name for the user.

**Authentication mode**
Select **Remote**.

**SSH key**
For remote users who require to access the system, select the public SSH key which stored locally on the system.

6. Click **Create**. Repeat these steps for all remote users.

## Using the CLI

To enable multifactor authentication on user groups, enter the following commands:

**For existing user groups**

```
chusergrp -multifactor yes <id or name of group>
```

**For new user groups for local users**

```
mkusergrp -name name -role role -multifactor yes
```

**For new user groups for remote users**

```
mkusergrp -name name -role role -multifactor yes -remote
```

*Adding local users to IBM Security Verify*
After user groups are enabled for multifactor authentication on the system, add local users to the Cloud Directory in IBM Security Verify.

To add local users to the Cloud Directory in IBM Security Verify, complete these steps:

1. Access the IBM Security Verify administrator dashboard by entering the following URL in a web browser:

   ```
   https://tenant.verify.ibm.com/ui/admin
   ```

   Where `tenant` is the name of the tenant that you specified when you created your subscription. Usually this tenant name is associated with your company or organization.
2. Select **Directory** > **Users & groups**.
3. On the **Users & groups** page, select **Add user**.
4. On the **Add user** page, enter the following information:

   **Identity provider**
   Select **Cloud Directory**.

   **Basic user profile**
   Enter the given name, middle name, and surname for the local user.

   **Username**
   Enter the same username that you defined for the user in the management GUI. Usernames must match otherwise authentication fails.

   **User information**
   Ensure to include an email address for the user. Each user receives instructions on login to IBM Security Verify which provides step-by-step instructions on selecting second-factors for the user's profile.
5. Click **Save**.
6. Repeat these steps for all local users that are defined on the system.

Depending on your security requirements, you can either require only a certain type of second factor or let users select from the supported list of available factors when they log in for the first time. If you want to control which factors and their settings users can specify, select **Security** from the IBM Security Verify administrator dashboard.

After you have added all users to IBM Security Verify, each user receives an email notification with instructions on setting up the second factors for their account. When these users log in to the

management interfaces that are configured to use multifactor authentication, they are directed to IBM Security Verify to present the second factor.

*Adding remote users to IBM Security Verify*
After user groups are enabled for multifactor authentication on the system, add remote users to the Cloud Directory in IBM Security Verify.

For remote users that authenticate with LDAP servers, install and configure IBM Security Verify Bridge for Directory Sync on your LDAP server, such as Windows Active Directory. IBM Security Verify Bridge for Directory Sync duplicates any users and groups that are defined on the source LDAP server into the Cloud Directory in IBM Security Verify. Any subsequent changes that are made to the source LDAP server are copied automatically to the Cloud Directory in IBM Security Verify.

To add remote users to IBM Security Verify, complete these steps:

1. Download and install IBM Security Verify Bridge for Directory Sync on your supported LDAP server. For more information, see IBM Security Verify Bridge for Directory Sync in the IBM Security Verify documentation.
2. After the installation completes, you need to update the sample JSON file and add properties that define connection settings and the LDAP server. On your supported LDAP server, open the `/DirectorySync/` directory, create a copy of the file `IcbLdapSync.json.ad-sample` and rename it `IcbLdapSync.json`.
3. Open the file in a text editor.
4. Under the `ibm-auth-api` JSON object, complete the following steps:

   **`client-id`**
   Enter the client ID that was automatically generated when you created the system as an application in IBM Security Verify.

   **`obf-client-secret`**

   a. Change the name of this JSON object to: **`client-secret`**.

   b. Enter the client secret that was automatically generated when you created the system as an application in IBM Security Verify.

   **`host`**
   Enter the hostname for the IBM Security Verify tenant.

   **`port` (optional)**
   Enter the port for your IBM Security Verify tenant.

   **`proxy` (optional)**
   Enter the hostname for the proxy server.
5. Under the `cloud-bridge` JSON object, add the following information:

   **`ldap-poll-time`**
   Enter a value in seconds to determine how frequently the LDAP directory and the Cloud Directory are synchronized. The default setting is 4 seconds. However, the example JSON file uses the value of 60 seconds. You can change this value based on your needs.
6. Under the `ldap` JSON object, add the following information:

   **`host`**
   Enter the hostname of the LDAP server.

   **`port`**
   Enter the port number for the LDAP server.

   **`user`**
   Enter the Administrator account distinguished name (DN) format for the LDAP server. Typically the Administrator account has these permissions, for example: `"user":"CN=Administrator,CN=Users,DC=mydomain,DC=com"`

**obf-password**

a. Change the name of this JSON object to: **password**.

b. Enter the password for the Administrator for the LDAP server.

7. Under the `user` JSON object, add the following information:

**realm**
Change the realm to "`cloudIdentityRealm`".

**userCategory**
Change `userCategory` to "`regular`".

**Notification**
Ensure that you include the following values:

```
 "notifyType": "EMAIL",
          "notifyPassword": true,
          "notifyManager": false
```

These values specify that an email notification is sent when they are added to the Cloud Directory inIBM Security Verify. In the email notification, they are instructed to access IBM Security Verify and select second factors to authenticate to the management interfaces.

8. Under the `ldap-to-scim` JSON object, remove or comment out the following string:

```
"append":"@cloudIdentityRealm".
```

9. Save the JSON file. The updated JSON file looks like the following example:

```
{
    "ibm-auth-api":{
        "timeout":100,
        "client-id":"xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
        "client-secret":"xxxxxxxxx",
        "protocol":"https",
        "host":"tenant.verify.ibm.com",
        "port":xxx,
        "proxy":"x.xx.xx.x:xxxx",
        "max-handles":16
    },
    "cloud-bridge":{
        "max-ops": 512,
        /* "trace-file":"c:/tmp/cloudbridge.log", */
        "enable-op-log":true,
        "op-log-rollover": 2097152,
        "do-not-sync-delete":  false,
        "ldap-search-filter":"(&(|(objectClass=user)(objectClass=group))(!
(isCriticalSystemObject=*)))",
        "ldap-is-deleted-attr":"isDeleted",
        "ldap-poll-time":60,
        /* "log-stats-interval": 30, */
        /* "ldap-base-dn": "DC=mycompany,DC=com", */
        "ldap":{
            "host":"ldapserver.example.com",
            "port":xxx,
            "user":"CN=Administrator,CN=Users,DC=mycompany,DC=com",
            "password":"password",
            "use-tls":false,
            "start-tls":false
        },
        "user":{
            "ldap-object-classes": [ "user" ],
            "scim-external-id-attr":"externalId",
            "scim-outline":{
                "urn:ietf:params:scim:schemas:extension:ibm:2.0:User":{
                    "realm":"cloudIdentityRealm",
                    "userCategory":"regular"
                },
                "urn:ietf:params:scim:schemas:extension:ibm:2.0:Notification": {
                    "notifyType": "EMAIL",
                    "notifyPassword": true,
                    "notifyManager": false
                },
                "schemas":[
                    "urn:ietf:params:scim:schemas:core:2.0:User",
```

```
                    "urn:ietf:params:scim:schemas:extension:ibm:2.0:User",
                    "urn:ietf:params:scim:schemas:extension:enterprise:2.0:User",
                    "urn:ietf:params:scim:schemas:extension:ibm:2.0:Notification"
                ]
            },
            "ldap-to-scim":[
                {
                    "ldap":"sAMAccountName",
                    "tweaks":{
                        /* "append":"@cloudIdentityRealm" */
                    },
                    "new-attr":{
                        "scim":{"userName":"{{value}}"}
                    },
                    "mod-attr":{
                        "scim":{
                            "add":{"op":"add","path":"userName","value":"{{value}}"},
                            "remove":{"op":"remove","path":"userName"},
                            "replace":{"op":"replace","path":"userName","value":"{{value}}"}
                        }
                    }
                },
/* Example of using a custom attribute: */
/*
                {
                    "ldap":"carLicense",
                    "is-multi-value":true,
                    "new-attr": {
                        "scim":{
                            "urn:ietf:params:scim:schemas:extension:ibm:2.0:User":
{"customAttributes":[{"values":["{{value}}"],"name":"myattrscim"}]}
            }
                    },
                    "mod-attr":{
                        "scim":{
                            "add":
{"op":"add","path":"urn:ietf:params:scim:schemas:extension:ibm:2.0:User:customAttributes",
                                "value":[{"name":"myattrscim","values":["{{value}}"]}]},
                            "remove-all":
{"op":"remove","path":"urn:ietf:params:scim:schemas:extension:ibm:2.0:User:customAttributes[
name eq \"myattrscim\"]"},
                            "remove":
{"op":"remove","path":"urn:ietf:params:scim:schemas:extension:ibm:2.0:User:customAttributes[
name eq \"myattrscim\"].values",
                                "value":["{{value}}"]},
                            "replace":
{"op":"add","path":"urn:ietf:params:scim:schemas:extension:ibm:2.0:User:customAttributes",
                                "value":[{"name":"myattrscim","values":["{{value}}"]}]}
                        }
                    }
                },
*/
                {
                    "ldap":"cn",
                    "new-attr":{
                        "scim":{"displayName":"{{value}}"}
                    },
                    "mod-attr":{
                        "scim":{
                            "add":{"op":"add","path":"displayName","value":"{{value}}"},
                            "remove":{"op":"remove","path":"displayName"},
                            "replace":
{"op":"replace","path":"displayName","value":"{{value}}"}
                        }
                    }
                },
                {
                    "ldap":"telephoneNumber",
                    "new-attr":{
                        "scim":{"phoneNumbers":[{"type":"work","value":"{{value}}"}]}
                    },
                    "mod-attr":{
                        "scim":{
                            "add":{"op":"add","path":"phoneNumbers","value":
[{"type":"work","value":"{{value}}"}]},
                            "remove":{"op":"remove","path":"phoneNumbers[type eq
\"work\"]"},
                            "replace":{"op":"add","path":"phoneNumbers","value":
[{"type":"work","value":"{{value}}"}]}
                        }
                    }
                },
```

```
                {
                    "ldap":"mobile",
                    "new-attr":{
                        "scim":{"phoneNumbers":[{"type":"mobile","value":"{{value}}"}]}
                    },
                    "mod-attr":{
                        "scim":{
                            "add":{"op":"add","path":"phoneNumbers","value":
[{"type":"mobile","value":"{{value}}"}]},
                            "remove":{"op":"remove","path":"phoneNumbers[type eq
\"mobile\"]"},
                            "replace":{"op":"add","path":"phoneNumbers","value":
[{"type":"mobile","value":"{{value}}"}]}
                        }
                    }
                },
                {
                    "ldap":"homePhone",
                    "new-attr":{
                        "scim":{"phoneNumbers":[{"type":"home","value":"{{value}}"}]}
                    },
                    "mod-attr":{
                        "scim":{
                            "add":{"op":"add","path":"phoneNumbers","value":
[{"type":"home","value":"{{value}}"}]},
                            "remove":{"op":"remove","path":"phoneNumbers[type eq
\"home\"]"},
                            "replace":{"op":"add","path":"phoneNumbers","value":
[{"type":"home","value":"{{value}}"}]}
                        }
                    }
                },
                {
                    "ldap":"ipPhone",
                    "new-attr":{
                        "scim":{"phoneNumbers":[{"type":"pager","value":"{{value}}"}]}
                    },
                    "mod-attr":{
                        "scim":{
                            "add":{"op":"add","path":"phoneNumbers","value":
[{"type":"pager","value":"{{value}}"}]},
                            "remove":{"op":"remove","path":"phoneNumbers[type eq
\"pager\"]"},
                            "replace":{"op":"add","path":"phoneNumbers","value":
[{"type":"pager","value":"{{value}}"}]}
                        }
                    }
                },
                {
                    "ldap":"facsimileTelephoneNumber",
                    "new-attr":{
                        "scim":{"phoneNumbers":[{"type":"fax","value":"{{value}}"}]}
                    },
                    "mod-attr":{
                        "scim":{
                            "add":{"op":"add","path":"phoneNumbers","value":
[{"type":"fax","value":"{{value}}"}]},
                            "remove":{"op":"remove","path":"phoneNumbers[type eq \"fax\"]"},
                            "replace":{"op":"add","path":"phoneNumbers","value":
[{"type":"fax","value":"{{value}}"}]}
                        }
                    }
                },
                {
                    "ldap":"givenName",
                    "new-attr":{
                        "scim":{"name":{"givenName":"{{value}}"}}
                    },
                    "mod-attr":{
                        "scim":{
                            "add":{"op":"add","path":"name.givenName","value":"{{value}}"},
                            "remove":{"op":"remove","path":"name.givenName"},
                            "replace":
{"op":"replace","path":"name.givenName","value":"{{value}}"}
                        }
                    }
                },
                {
                    "ldap":"sn",
                    "new-attr":{
                        "scim":{"name":{"familyName":"{{value}}"}}
                    },
```

```
                    "mod-attr":{
                        "scim":{
                            "add":{"op":"add","path":"name.familyName","value":"{{value}}"},
                            "remove":{"op":"remove","path":"name.familyName"},
                            "replace":
{"op":"replace","path":"name.familyName","value":"{{value}}"}
                        }
                    }
                },
                {
                    "ldap":"displayName",
                    "new-attr":{
                        "scim":{"name":{"formatted":"{{value}}"}}
                    },
                    "mod-attr":{
                        "scim":{
                            "add":{"op":"add","path":"name.formatted","value":"{{value}}"},
                            "remove":{"op":"remove","path":"name.formatted"},
                            "replace":
{"op":"replace","path":"name.formatted","value":"{{value}}"}
                        }
                    }
                },
                {
                    "ldap":"streetAddress",
                    "new-attr":{
                        "scim":{"addresses":[{"type":"work","streetAddress":"{{value}}"}]}
                    },
                    "mod-attr":{
                        "scim":{
                            "add":{"op":"add","path":"addresses","value":
[{"type":"work","streetAddress":"{{value}}"}]},
                            "remove":{"op":"remove","path":"addresses[type eq \"work\"]"},
                            "replace":{"op":"add","path":"addresses","value":
[{"type":"work","streetAddress":"{{value}}"}]}
                        }
                    }
                },
                {
                    "ldap":"manager",
                    "is-dn":true,
                    "new-attr":{
                        "scim":
{"urn:ietf:params:scim:schemas:extension:enterprise:2.0:User":{"manager":
{"value":"{{value}}"}}}
                    },
                    "mod-attr":{
                        "scim":{
                            "add":
{"op":"add","path":"urn:ietf:params:scim:schemas:extension:enterprise:2.0:User:manager.value
","value":"{{value}}"},
                            "remove":
{"op":"remove","path":"urn:ietf:params:scim:schemas:extension:enterprise:2.0:User:manager.va
lue"},
                            "replace":
{"op":"replace","path":"urn:ietf:params:scim:schemas:extension:enterprise:2.0:User:manager.v
alue","value":"{{value}}"}
                        }
                    }
                },
                {
                    "ldap":"mail",
                    "new-attr":{
                        "scim":{"emails":[{"type":"work","value":"{{value}}"}]}
                    },
                    "mod-attr":{
                        "scim":{
                            "add":{"op":"add","path":"emails","value":
[{"type":"work","value":"{{value}}"}]},
                            "remove":{"op":"remove","path":"emails[type eq \"work\"]"},
                            "replace":{"op":"add","path":"emails","value":
[{"type":"work","value":"{{value}}"}]}
                        }
                    }
                },
                {
                    "ldap":"department",
                    "new-attr":{
                        "scim":
{"urn:ietf:params:scim:schemas:extension:enterprise:2.0:User":{"department":"{{value}}"}}
                    },
                    "mod-attr":{
```

```
                        "scim":{
                            "add":
{"op":"add","path":"urn:ietf:params:scim:schemas:extension:enterprise:2.0:User:department","
value":"{{value}}"},
                            "remove":
{"op":"remove","path":"urn:ietf:params:scim:schemas:extension:enterprise:2.0:User:department
"},
                            "replace":
{"op":"replace","path":"urn:ietf:params:scim:schemas:extension:enterprise:2.0:User:departmen
t","value":"{{value}}"}
                        }
                    }
                },
                {
                    "ldap":"objectGUID",
                    "is-binary":true,
                    "new-attr":{
                        "scim":{"externalId":"{{value}}"}
                    }
                },
                {
                    "ldap":"userAccountControl",
                    "new-attr":{
                        "cond-scim":{
                            "if":{"&":["{{value}}", "2"]},
                            "then":{"active":false},
                            "else":{"active":true}
                        },
                        "mod-attr":{
                            "cond-scim":{
                                "if":{"&":["{{value}}","2"]},
                                "then":{
                                    "add":{"op":"add","path":"active","value":false},
                                    "remove":{"op":"remove","path":"active"},
                                    "replace":{"op":"replace","path":"active","value":false}
                                },
                                "else":{
                                    "add":{"op":"add","path":"active","value":true},
                                    "remove":{"op":"remove","path":"active"},
                                    "replace":{"op":"replace","path":"active","value":true}
                                }
                            }
                        }
                    }
                ]
            },
            "group":{
                "ldap-object-classes": [ "group" ],
                "scim-external-id-attr":"externalId",
                "scim-outline":{
                    "schemas":[
                        "urn:ietf:params:scim:schemas:core:2.0:Group",
                        "urn:ietf:params:scim:schemas:extension:ibm:2.0:Group"
                    ]
                },
                "ldap-to-scim":[
                    {
                        "ldap":"sAMAccountName",
                        "new-attr":{
                            "scim":{"displayName":"{{value}}"}
                        },
                        "mod-attr":{
                            "scim":{
                                "add":{"op":"add","path":"displayName","value":"{{value}}"},
                                "remove":{"op":"remove","path":"displayName"},
                                "replace":
{"op":"replace","path":"displayName","value":"{{value}}"}
                            }
                        }
                    },
                    {
                        "ldap":"objectGUID",
                        "is-binary":true,
                        "new-attr":{
                            "scim":{"externalId":"{{value}}"}
                        }
                    },
                    {
                        "ldap":"description",
                        "new-attr":{
                            "scim":{"urn:ietf:params:scim:schemas:extension:ibm:2.0:Group":
```

```
{"description":"{{value}}"}}
                    },
                    "mod-attr":{
                        "scim":{
                            "add":
{"op":"add","path":"urn:ietf:params:scim:schemas:extension:ibm:2.0:Group:description","value
":"{{value}}"},
                            "remove":
{"op":"remove","path":"urn:ietf:params:scim:schemas:extension:ibm:2.0:Group:description"},
                            "replace":
{"op":"replace","path":"urn:ietf:params:scim:schemas:extension:ibm:2.0:Group:description","v
alue":"{{value}}"}
                        }
                    }
                },
                {
                    "ldap":"member",
                    "is-dn":true,
                    "is-multi-value":true,
                    "new-attr":{
                        "scim":{"members":[{"type":"{{type}}","value":"{{value}}"}]}
                    },
                    "mod-attr":{
                        "scim":{
                            "add":{"op":"add","path":"members","value":
[{"type":"{{type}}","value":"{{value}}"}]},
                            "remove":{"op":"remove","path":"members[value eq
{{jq_value}}]"},
                            "remove-all":{"op":"remove","path":"members"},
                            "replace":{"op":"replace","path":"members","value":
[{"type":"{{type}}","value":"{{value}}"}]}
                        }
                    }
                }
            ]
        }
    }
}
```

10. To synchronize the LDAP server and the Cloud Directory for the first time, run `IcbLdapSync.exe`.

Depending on your security requirements, you can either require only a certain type of second factor or let users select from the supported list of available factors when they log in for the first time. If you want to control which factors and their settings users can specify, select **Security** from the IBM Security Verify administrator dashboard.

After you have added all users to IBM Security Verify, each user receives an email notification with instructions on setting up the second factors for their account. When these users log in to the management interfaces that are configured to use multifactor authentication, they are directed to IBM Security Verify to present the second factor.

### *Configuring multifactor authentication with Duo Security*

The system integrates with Duo Security to provide multifactor authentication for system users.

With Duo Security, security administrators can configure the system as an application that requires two factors for users and user groups to access the system with either the management GUI or CLI.

You can manage every aspect of your Duo two-factor authentication system from the Duo Admin Panel including creating and managing applications, enrolling and activating users, issuing and managing SMS passcodes and bypass codes, managing mobile devices, fine-tuning the user experience of your Duo installation.

Multifactor authentication can be used to protect both local users, including superuser, and remote users.

For information on how to configure directory synchronization for remote users with Duo Security, see Duo Directory Synchronization.

Duo Security configures the management GUI and the command-line interface as separate API clients that require separate credentials. For CLI access, Duo Security communicates with the system through standard REST API requests. For GUI-based logins, the system communicates with Duo Security through the OpenID Connect (OIDC) protocol.

## Prerequisites

Ensure that the following prerequisite tasks are completed on the system before you configure multifactor authentication on Duo Security:

1. Configure a DNS server. To create a DNS server, select **Settings** > **Network** > **DNS**.
2. Configure an HTTP proxy server or configure your firewall to access Duo Security. To create an HTTP proxy server, **Settings** > **Network** > **Internal Proxy Server**. If your system does not directly connect to the Internet, you can create a firewall exception to allow your system access to Duo Security.

   **Important:** If the proxy server is not configured correctly, the system cannot communicate with Duo Security and the login to the system fails.
3. For the management GUI and the command-line interface, ensure that the inactivity logout is equal to or greater than the time it takes for a user to receive a one-time passcode (OTP) from the authentication service. The default value for the inactivity timeout is 30 minutes for the management GUI and 15 minutes for the CLI. To set the inactivity timeout for both interfaces in the management GUI, select **Settings** > **Security** > **Inactivity Logout**.
4. Ensure that the SSH grace time for the command-line interface is equal to or greater than the time it takes for a user to receive a one-time passcode (OTP) from the authentication service. The default value for the SSH grace time period is 60 seconds. To set SSH grace time on the system in the command-line interface, use the **`chsecurity -sshgracetime`** command.

The following prerequisite steps on Duo Security must be completed before you can configure multifactor authentication on the system:

1. Create a subscription for Duo Security. For more information, see Editions & Pricing page on Duo Security website. During subscription creation, you specify a tenant that is used to create a URL to access the Duo Security dashboard.
2. To create a Duo Security application, enter the following:

   ```
   admin-XXXXXXXX.duosecurity.com
   ```

   where XXXXXXXX is unique to your subscription that you specified when you created your subscription.
3. On **Applications** page, click **Protect an Application** and select **Web SDK**. The credentials on this page are the OpenID Client ID and Secret required for GUI access.
4. On **Applications** page, click **Protect an Application** and select **UNIX Application**. The credentials on this page are the API Client Integration Key and Secret Key required for CLI access.

**Note:** Duo Security does not support the use of IP addresses when accessing your system by using a web browser. Once multifactor authentication is configured, make sure that a hostname is used when accessing your system by using a web browser.

## Using the management GUI

To configure multifactor authentication on Duo Security, complete these steps:

1. Select **Settings** > **Security** > **Multifactor Authentication**.
2. On **Multifactor Authentication** page, select **Duo Security** authentication provider.
3. Enter the hostname and port of the authentication server. For Duo Security, enter the following:

   ```
   api-XXXXXXXX.duosecurity.com
   ```

   Where XXXXXXXX is unique to your subscription that you specified when you created your subscription.
4. For the **OpenID Credentials**, complete the following steps in the Duo Security interface to add the **Client ID** and **Client Secret**:

   a. Select **Dashboard** > **Applications** > **Protect an Application**.

   b. Search **Web SDK**, click **Protect**.

c. Copy **Client ID** and **Client Secret**.

5. For the **API Client Credentials**, complete the following steps in the Duo Security interface to add the **Integration Key** and **Secret Key**:
   a. Select **Dashboard** > **Applications** > **Protect an Application**.
   b. Search **UNIX Application**, click **Protect**.
   c. Copy **Integration Key** and **Secret Key**.
6. Click **Save**.
7. On the confirmation page, click **Confirm** to enable multifactor authentication for the system.

Multifactor authentication is enabled for the system. You can configure user groups to use multifactor authentication. Click **Navigate** to launch the **User Groups** page.

## Using the CLI

You can enable multifactor authentication on the system by using the following command:

- To enable multifactor authentication with Duo Security, enter the following command:

```
chauthmultifactorduo -enable -hostname api-XXXXXXXX.duosecurity.com -openidclientid
openid_client_id -openidclientsecret openid_client_secret -integrationkey integration_key
-secretkey secret_key
```

  where parameter **-enable** enables multifactor authentication service and **-hostname** is the hostname of the Duo Security tenant. The **-openidclientid** and the **-openidclientsecret** are the Open ID Client and Open ID Secret for the system, required to enable multifactor authentication for login to the storage system GUI. The **-integrationkey** and the **-secretkey** are the integration key and secret key for the system, required to enable multifactor authentication for login to the Storage Virtualize CLI.

*Configuring user groups for multifactor authentication with Duo Security*
After you configure multifactor authentication on the system, you must enable multifactor authentication for user groups and add users to those groups in Duo Security.

As part of multifactor authentication configuration, you must enable the function per user group. The system supports enabling multifactor authentication for local and remote user groups.

The security administrator must define all local users manually in Duo Security. When users log into the system with multifactor authentication, their username is used to look up the required second factor in Duo Security. Users can set up multiple second factors to avoid getting locked out of the system.

**Note:** When you configure multifactor authentication for the first time, follow these guidelines to avoid getting locked out of the system unintentionally:

- Ensure at least one user with Security Administrator role does not have multifactor authentication enabled. If the security administrator gets locked out of the system because of errors in the multifactor authentication set up, the additional user can still access the system.
- Enable multifactor authentication on a user group without logged in users or leave an SSH session active to avoid locking out users on the system.

For information on how to configure directory synchronization for remote users with Duo Security, see Duo Directory Synchronization.

## Using the management GUI

To enable multifactor authentication on user groups, complete these steps:

**For existing user groups:**

1. In the management GUI, select **Access** > **Users by Group**.
2. Select the user group from the left navigation and select **User Group Actions** > **Properties**.

3. On the **User Group Properties** page, select **On** under **Multifactor Authentication** to enable second-factor authentication for all users with the user group. These users authenticate with the first factors that are stored on the local system and then are required to provide a second factor to access the system through a supported authentication service.
4. Click **OK**.

**For new user groups for local users:**

1. In the management GUI, select **Access** > **Users by Group** > **Create User Group**.
2. On the **Create User Group** page, enter the following information:

   **Group Name**
   Enter a name of the user group.

   **Ownership Group**
   If ownership groups are configured on your system, you can select an ownership group for the user group.

   **Multifactor Authentication**
   Select **On** to enable second-factor authentication for local users on the system. These users authenticate with the first factors that are stored on the local system and then are required to provide a second factor to access the system through a supported authentication service.

   **Role**
   Select a role that for the user group.
3. Click **Create**.
4. Select **Access** > **Users by Group** > **Create User**.
5. On the **Create User** page, enter the following information:

   **Name**
   Enter a user name for the user. This user name must match the user name that is added on Duo Security.

   **Authentication mode**
   Select **Local**.

   **User Group**
   Select the name of the user group that the local user belongs to.

   **Password**
   Enter a password that is used as the first factor for management GUI access.

   **SSH key**
   For CLI users, include a public SSH key that is used as the first factor for CLI access.
6. Click **Create**. Repeat these steps for all local users.

**For new user groups for remote users:**

1. Select **Access** > **Users** > **Create User Groups**.
2. On the **Create User Group** page, enter the following information:

   **Group Name**
   Enter the name of the group that is on the remote LDAP server. The name of the group on the system must match.

   **Ownership Group**
   If ownership groups are configured on your system, you can select an ownership group for the user group.

   **Remote Authentication**
   Select **LDAP**.

   **Multifactor authentication**
   Select **On** to enable second-factor authentication for remote users on the system. These users authenticate with the first factors that are stored on the remote LDAP server and then are

required to provide a second factor to access the system through a supported authentication service.

3. Click **Create**.
4. Remote users are defined on the remote authentication service. Only remote users who require access to the command-line interface using a Secure Shell (SSH) key, need to be created. To create remote users, who require access to the command-line interface, select **Access** > **Users by Groups** > **Create User**.
5. On the **Create User** page, enter the following information:

   **Name**
   Enter a name for the user.

   **Authentication mode**
   Select **Remote**.

   **SSH key**
   For remote users who require to access the system, select the public SSH key which stored locally on the system.
6. Click **Create**. Repeat these steps for all remote users.

## Using the CLI

To enable multifactor authentication on user groups, enter the following commands:

**For existing user groups**

```
chusergrp -multifactor yes <id or name of group>
```

**For new user groups for local users**

```
mkusergrp -name name -role role -multifactor yes
```

**For new user groups for remote users**

```
mkusergrp -name name -role role -multifactor yes -remote
```

*Add local users to Duo Security*
After user groups are enabled for multifactor authentication on the system, add local users in Duo Security.

To add local users in Duo Security complete these steps:

1. Log into Duo Security with your Duo Administrator account.
2. Complete the steps described in the Duo Security documentation. For more information, see Enrolling users.

*Add remote users to Duo Security*
After user groups are enabled for multifactor authentication on the system, add remote users in Duo Security.

For remote users that authenticate with LDAP servers, Duo Security supports automatic enrollment. Duo Security's Directory Sync feature duplicates any users and groups that are defined on the source LDAP server. Any subsequent changes that are made to the source LDAP server are copied automatically to Duo Security. To add remote users in Duo Security, complete these steps:

1. Log in to Duo Security with your Duo Administrator account.
2. Configure the Directory Sync feature on Duo Security.

   For information on how to configure directory synchronization for remote users with Duo Security, see Duo Directory Synchronization.

# Single Sign-on

Single Sign-on (SSO) authentication requires users to register their credentials only once when the user signs on to the application for the first time. The user information is stored at the Identity Provider (IdP) that manages the user credentials and determines whether the user is required to authenticate again or not.

SSO is a feature that delegates all authentication to a trusted Identity Provider (IdP). The applications are configured with the IdP. When the user logs in to the application, the IdP validates the registered information and enables the user to log in to the application. In this method, both the first and second factor authentications are delegated to a trusted Identity Provider, which is useful if you do not want IBM Storage Virtualize to do any of the authentications. Also, additional authentication factors can be configured to provide multifactor authentication for users.

IBM Storage Virtualize uses the OpenID Connect (OIDC) protocol to communicate with the SSO service when a user logs in to the GUI. SSO is supported only for web browsers. SSO can be used only to protect user login to the management GUI. The SSO feature is not supported for the command-line interface. To provide multifactor authentication for CLI access, the system must be integrated natively with IBM Security Verify. For more information, refer to "Multifactor authentication" on page 92. The SSO service is only supported for remote users that are defined on a remote LDAP server, and does not support local users that are defined locally within the IBM Storage Virtualize system, including the superuser.

Remote users are users who are defined on a remote LDAP server. For remote users that authenticate with LDAP servers, install and configure IBM Security Verify Bridge for Directory Sync on your LDAP server, such as Windows Active Directory. IBM Security Verify Bridge for Directory Sync duplicates any users and groups that are defined on the source LDAP server into the Cloud Directory in IBM Security Verify. Any subsequent changes that are made to the source LDAP server are copied automatically to the Cloud Directory in IBM Security Verify. For more information, see IBM Security Verify Bridge for Directory Sync in the IBM Security Verify documentation.

To protect CLI access for local users or remote users with SSO, local and remote user groups can be configured with the password and SSH key authentication method, which requires users to authenticate with both a password and an SSH key when a user logs in through the CLI.

When SSO is configured and enabled for the user groups, the GUI displays a new **Sign In with SSO** option on the login window. The user is redirected to do the authentication with the configured SSO service. When the authentication is successful using SSO, the user does not have to reauthenticate on subsequent login attempts.

## MFA versus SSO

Multifactor authentication (MFA) and Single Sign-on (SSO) can both be configured to provide multiple authentication factors when a user logs in to the system. However, it's important to understand the key differences and limitations of each authentication method. The following table summarizes the difference and limitations of both the authentication methods and enables a user to select a method based on their needs.

*Table 25. Authentication method summary*

| | **Multifactor Authentication** | **Single Sign-on** |
|---|---|---|
| First factor handled by | IBM Storage Virtualize | Identity Provider (IdP) |
| Second factor handled by | Multifactor authentication service (IBM Security Verify) | Identity Provider (IdP) |
| CLI login support | Yes | No |
| GUI login support | Yes | Yes |
| Local user support | Yes | No |
| Remote user support | Yes | Yes |

| *Table 25. Authentication method summary (continued)* | | |
|---|---|---|
| | **Multifactor Authentication** | **Single Sign-on** |
| Supported services | IBM Security Verify | Microsoft Active Directory Federation Services |
| Protocol | OpenID Connect (OIDC) | OpenID Connect (OIDC) |

### LDAP versus SSO

LDAP (Lightweight Directory Access Protocol) or Single Sign-on (SSO) can be configured for authenticating remote users when a user logs in to the system. However, it's important to understand the key differences and limitations of each authentication method. The following table summarizes the difference and limitations of both the authentication methods and enables a user to select a method based on their needs.

| *Table 26. Difference and limitations between LDAP and SSO* | | |
|---|---|---|
| | **LDAP authentication** | **Single Sign-on** |
| First factor handled by | LDAP server | Identity Provider (IdP) |
| Second factor handled by | NA | Identity Provider (IdP) |
| CLI login support | Yes | No |
| GUI login support | Yes | Yes |
| Local user support | No | No |
| Remote user support | Yes | Yes |
| Supported services | • Microsoft Active Directory<br>• IBM Security Directory Server<br>• Others (for example, OpenLDAP) | Microsoft Active Directory Federation Services |
| Protocol | LDAP | OpenID Connect (OIDC) |

## Configuring single sign-on

Single sign-on delegates all authentication to a trusted Identity Provider (IdP).

With single sign-on, users need to provide their credentials once when they log in to an application or system, rather than repeatedly providing the credentials for every individual application or system. Each individual IBM Storage Virtualize system is considered a separate application and must be added to the IdP. When single sign-on is enabled on the system, both first-factor and second-factor authentication are delegated to the IdP. For more information on supported single sign-on providers, see Supported authentication providers.

In the management GUI, users select **Sign In with SSO** on the log in prompt and are redirected to complete authentication through the configured IdP. You can configure the IdP to provide first factor authentication to your users. Several multifactor authentication cloud-based providers can be added to the single sign-on configuration to require additional user authentication if necessary. When authentication is completed successfully, the uses are redirected to management GUI.

### Prerequisites

Ensure that the following prerequisite tasks are completed on the system before you configure single sign-on:

1. Ensure that the system is updated to 8.5.0 or later release.

2. Configure a DNS server. To create a DNS server, select **Settings** > **Network** > **DNS**.
3. If your authentication server is outside of your private network, you must configure an HTTP proxy server or configure your firewall to access your authentication server.

   **Note:** If you plan to use a proxy server to access the Identity Provider, ensure that you select **Yes** on the **Use proxy** option when you configure single sign-on in the management GUI.

   To create an HTTP proxy server, **Settings** > **Network** > **Internal Proxy Server**. If your authentication server is within your private network, proxy or firewall changes are not required.
4. For the management GUI, ensure that the inactivity logout is equal to or greater than the time it takes for a user to receive a one-time passcode (OTP) from the authentication service. The default value for the inactivity timeout is 30 minutes for the management GUI. To set the inactivity timeout in the management GUI, select **Settings** > **Security** > **Inactivity Logout**. To set the GUI inactivity timeout on the command-line interface, use the **`chsecurity -guitimeout`** command.

## Choose an authentication provider

Configure single sign-on using a supported authentication provider. Once single sign-on has been configured, you must then decide which user groups should be enabled for single sign-on.

### *Configuring single sign-on with PingOne*

PingOne can be configured as the authentication provider for the system.

**Note:** Ensure that the prerequisite tasks are completed on the system before you configure single sign-on. For more information, see the prerequisites in “Configuring single sign-on” on page 114.

## Prerequisites

The following prerequisite steps on PingOne must be completed before you can configure single sign-on on the system:

1. Create a subscription for PingOne. A free trial subscription is also available. For more information, see Ping Single Sign-On. During subscription creation, you specify a tenant that is used to create a URL to access the PingOne dashboard.
2. Access the PingOne administrator dashboard by entering the following URL in a web browser (this is usually received in your welcome email). The sign-on URL is specific to your admin console and includes your environment ID. For example, the URL might look like this:

   ```
   http://console.pingone.asia/index.html?env=cde34209-9366-45aa-9e5f-a19d3ba3eada
   ```
3. In the PingOne administrator console, from the menu select **Connections** > **Applications**.
4. Select **Application +** to add a new application.

   **Note:** Each system must be added as a separate application.

   The following table shows the required fields and actions for the **Register an application** page in the PingOne interface.

   *Table 27. Register an application page*

   | **Field** | **Action** |
   |---|---|
   | Application Name | Enter a name to identify the system on PingOne. If you are adding multiple systems, enter a unique name. |
   | Application Type | Select OIDC Web App. |

5. Click **Save**. After the system is saved as a new app registration, the application reloads with the **Overview** page selected.

6. To find the OpenID Configuration Endpoints URL for the tenant, on the Configuration page expand URLs and copy the OIDC Discovery Endpoint URL. This value must be specified as the OpenID Configuration Endpoint URL on the Single sign-on page in the management GUI.
7. On the Configuration page, expand **General** and copy the Client ID. This value must be specified as the Client ID on the Single sign-on page in the management GUI.
8. On the Configuration page, expand **General** and copy the Client Secret. This value must be specified as the Client ID on the Single sign-on page in the management GUI.

   **Note:** The system does not support client secret values that begin with a hyphen character. Use the reveal button to inspect the Client Secret automatically generated by PingOne. If the Client Secret begins with a hyphen character, click the **Generate New Secret** button.
9. On the Configuration page, click the **Edit** button. For **Response Type**, ensure that **Code** is selected. For **Grant Type**, ensure that **Authorization Code** is selected. For **Redirect URLs**, multiple redirect URIs can be specified for the management GUI. For management GUI access, the redirect URI consists of the management IP address or hostname followed by /sso. For example, https://hostname.com/sso. For **Token EndPoint Authentication Method**, select **Client Secret Post**. Click the **Save** button.
10. On the Resources page, click **Edit** on the Allowed Scopes and ensure that **openid** and **profile** are selected. Click the **Save** button.
11. Close the Application window to return to the Applications page. Click the toggle to enable user access to this application.
12. In the PingOne administrator console, from the menu select **Connections** > **Resources** and then select **OpenID Connect**.
13. Ensure the authentication provider is configured to send back the group claim to the system. The group claim identifies which groups the authenticating user belongs to. Some authentication providers do not send the group claim by default, so this typically requires some configuration on the authentication provider. On the **Attributes** tab, click the **Edit** button. Click **+ Add** to add a new attribute for the group claim. Ensure that the **Attributes** field contains **group** and the **PingOne Mappings** field contains a value of **Group Names**. Configure the **ID Token** and **UserInfo** values according to your organization's requirements; this determines how the **group** attribute is sent back to the application. Click the **Save** button.
14. If you require the group claim to be sent back to the application in the ID Token, open the **Scopes** page and **Edit** the **profile** scope. Click inside the **Mapped Attributes** area and type **group** to add the group attribute. Click the **Save** button.
15. On the Attributes page, ensure that the existing `preferred_username` attribute has a **PingOne Mappings** value of **Username** and that **ID Token** is selected.
16. In the PingOne administrator console, from the menu select **Connections** > **Applications**.
17. Select your application. On the Attribute Mappings page, click **Edit** on **Custom Attributes**. Click **+ Add** to add a new attribute. For the **Attributes** field, enter groups. For the **PingOneMappings** field, select **GroupNames**. Click **Save**.

## Using the management GUI

To configure single sign-on with PingOne, complete these steps:

1. Select **Settings** > **Security** > **Single Sign-on**.
2. Enter the OpenID Configuration Endpoint URL of the authentication server that you obtained from the application details earlier. For PingOne, the format is:

   ```
   https://auth.pingone.region/environment-id/as/.well-known/openid-configuration
   ```

   Where region is the region that the PingOne tenant is hosted in, and environment-id is the Environment ID specific to this application. This can be found by selecting from the menu **Connections** > **Applications** > **Click your application** > **Configuration** > **OIDC Discovery Endpoint**.

3. For the **OpenID Credentials**, add the **Client ID** and **Client Secret** that you copied on from the Configuration page in the PingOne interface.
4. For the User claim, the value to enter depends on how your authentication provider is configured. The User claim must match the name that the authentication service uses to specify the username attribute in the ID Token it sends to the system. Typically this value is preferred_username, but can be customized on the authentication provider.
5. For the Group claim, the value to enter depends on how your authentication provider is configured. The Group claim must match the name that the authentication service uses to specify the group attribute, either in the ID Token it sends to the system, or in the UserInfo endpoint. Typically this value is groups, but can be customized on the authentication provider.
6. For Proxy server, consider how the system accesses the authentication provider. For an authentication provider within your network, a proxy server usually isn't needed. If you connect to the authentication provider through the Internet, check the box and ensure a proxy server is defined on the system.
7. Click **Save**. On the confirmation page, click **Confirm** to enable single sign-on for the system.

Single sign-on is enabled for the system. You can configure user groups to use single sign-on. Click **Navigate** to launch the **User Groups** page.

## Using the command-line interface

To enable single sign-on, enter the following command:

```
chauthsinglesignon -oidcconfigurationendpoint https://auth.pingone.region/environment-
id/as/.well-known/openid-configuration -clientid xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx
-clientsecret xxxxxxxx -userclaim preferred_username -groupclaim groups -enable
```

In the example, region is the tenant region that is associated with your subscription, and environment-id is the Environment ID specific to this application. The values for the **-clientid** and the **-clientsecret** are the Client ID and Client Secret that are automatically generated when you created your system as an application in PingOne, and can be obtained from the Configuration page of your application in the PingOne interface. The values for the **-userclaim** and the **-groupclaim** should match the name of the claims configured for the ID Token on the authentication provider.

### *Configuring single sign-on with IBM Security Verify*

IBM Security Verify can be configured as the authentication provider for the system.

**Note:** Ensure that the prerequisite tasks are completed on the system before you configure single sign-on. For more information, see the Prerequisites section in "Configuring single sign-on" on page 114.

## Prerequisites

The following prerequisite steps on IBM Security Verify must be completed before you can enable single-sign on the system:

1. Create a subscription for IBM Security Verify. You need an IBMid to create a subscription. A 90-day free trial subscription is also available. For more information, see Cloud identity and access management (IAM) solutions. During subscription creation, you specify a tenant that is used to create a URL to access the IBM Security Verify dashboard.
2. Access the IBM Security Verify administrator dashboard by entering the following URL in a web browser:

   ```
   https://tenant.verify.ibm.com/ui/admin
   ```

   Where `tenant` is the name of the tenant that you specified when you created your subscription. Usually this tenant name is associated with your company or organization.
3. In the IBM Security Verify interface, select **Applications** > **Applications** > **Add application**.
4. Select > **Add Application**.

   **Note:** Each system must be added as a separate application.

The following table shows the required fields and actions for the **General** tab in the IBM Security Verify interface.

*Table 28. General tab*

| Field | Action |
|---|---|
| Name | Enter a name to identify the system on IBM Security Verify. If you are adding multiple systems, enter a unique name. |
| Description | Enter a brief description of the system. |
| Company name | Name of organization or company. |

The following table shows the required fields and actions for the **Sign-on** tab in the IBM Security Verify interface. The Sign-on tab is used to add the management GUI as an API-based client.

*Table 29. Sign-on tab*

| Field | Action | Details |
|---|---|---|
| Application URL | Enter the URL for your system. | Enter the URL that is used to access the management GUI. |
| Grant type | Select **Authorization code** and **JWT bearer**. | Two grant types are required for setting up SSO for the system. **Authorization code** indicates that the client can request access to protected resources on behalf of users. |
| Client ID | This value is automatically generated when the system is saved as an application. | This value must be entered to the Single Sign-on page in the management GUI under **OpenID Credentials**. |
| Client secret | This value is automatically generated when the system is saved as an application. | This value must be entered to the Single Sign-on page in the management GUI under **OpenID Credentials**. |
| User consent | Select **Do not ask for consent**. | |
| Redirect URIs | Enter the locations where the authorization server sends users after they are successfully authorized and granted an authorization code or access token. | Multiple redirect URIs can be specified for the management GUI. The redirect URI is comprised of the management IP address or hostname followed by `/sso`. For example,<br><br>https://hostname.com/sso |
| JWT bearer user identification | Select **Username**. | Indicates that the username field in the JWT bearer is used to find users in the Cloud Directory and determines what second factors IBM Security Verify presents to users when they log into the system. |

| *Table 29. Sign-on tab (continued)* | | |
|---|---|---|
| **Field** | **Action** | **Details** |
| JWT bearer default identity source | Ensure **Cloud Directory** is selected. | Indicates that the IBM Security Verify Cloud Directory is used to look up the second factor for the username. After you configure multifactor authentication on the system, users and user groups must be added to the Cloud Directory. |
| Generate refresh token | Ensure that this option is unchecked. | |
| Send all known user attributes in the ID token | Ensure that this option is checked. | |
| Access policies | Complete these steps:<br>a. Deselect **Use default policy**.<br>b. Click the **Edit** icon.<br>c. Select **Always require 2FA in all devices**.<br>d. Click **OK**. | This action creates an access policy which controls the authentication steps for system access. Access policies can specify different authentication requirements based on properties of the user or connection. In this case, all users must complete a second factor authentication every time they access the system from all devices. |
| Restrict custom scopes | Ensure this option is unchecked. | |

There is no action required for the **API access** tab in IBM Security Verify.

The following table shows the required fields and actions for the **API access** tab in the IBM Security Verify. Click **Save**. After the system is saved as a new application, the **Custom Application** reloads with the **Entitlements** tab selected.

5. On the **Entitlements** tab, select **Automatic access for all users and groups**.
6. Click **Save**.
7. Select **Applications** and select the application name that represents the system.
8. On the **Sign-on** tab, copy the **Client ID** and the **Client secret**. These values must be specified as the OpenID credentials on the Single sign-on page in the management GUI
9. Ensure the authentication provider is configured to send back the group claim in the ID Token sent to the system. The group claim identifies which groups the authenticating user belongs to. Some authentication providers do not send the group claim by default, so this typically requires some configuration on the authentication provider.

## Using the management GUI

1. Select **Settings** > **Security** > **Single Sign-on**.
2. Enter the OpenID Configuration Endpoint URL of the authentication server. For IBM Security Verify, enter the following:

```
https://tenant.verify.ibm.com/oidc/endpoint/default/.well-known/openid-configuration
```

where `tenant` is the name that is associated with your subscription.

3. For the **OpenID Credentials**, add the **Client ID** and **Client Secret** that you copied on from the Sign-on tab in the IBM Security Verify interface.
4. For the User claim, the value to enter depends on how your authentication provider is configured. The User claim must match the name that the authentication service uses to specify the username attribute in the ID Token it sends to the system. Typically this value is preferred_username, but can be customized on the authentication provider.
5. For the Group claim, the value to enter depends on how your authentication provider is configured. The Group claim must match the name that the authentication service uses to specify the group attribute in the ID Token it sends to the system. Typically, this value is groupIds, but can be customized on the authentication provider.
6. For Proxy server, consider how the system accesses the authentication provider. For an authentication provider within your network, a proxy server usually is not needed. If you connect to the authentication provider through the internet, check the box and ensure a proxy server is defined on the system.
7. Click **Save**. On the confirmation page, click **Confirm** to enable single sign-on for the system.

Single sign-on is enabled for the system. You can configure user groups to use single sign-on. Click **Navigate** to launch the **User Groups** page.

## Using the command-line interface

To enable single sign-on, enter the following command:

```
chauthsinglesignon -oidcconfigurationendpoint https://tenant.verify.ibm.com/oidc/endpoint/
default/.well-known/openid-configuration -clientid xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx
-clientsecret xxxxxxxx -userclaim preferred_username -groupclaim groupIds -enable
```

In the example, tenant is the tenant name that is associated with your subscription. The values for the **-clientid** and the **-clientsecret** are the Open ID Client and Open ID Secret that are automatically generated when you created your system as a custom application in IBM Security Verify on the Sign-on tab in the IBM Security Verify interface. The values for the **-userclaim** and the **-groupclaim** should match the name of the claims configured for the ID Token on the authentication provider.

### *Configuring single sign-on with Microsoft Active Directory Federation Services (AD FS)*

Microsoft Active Directory Federation Services (AD FS) can be configured as the authentication provider for the system.

**Note:** Ensure that the prerequisite tasks are completed on the system before you configure single sign-on. For more information, see the prerequisites in “Configuring single sign-on” on page 114.

## Prerequisites

The following prerequisite steps on AD FS must be completed before you can enable single-sign on the system:

1. Ensure that you are running at least AD FS 5.0 on a Microsoft Windows 2019 server.
2. Access the Server Manager dashboard.
3. From right menu, select **Tools** > **AD FS Management**.
4. From left menu, choose **AD FS** > **Application Groups**. Select **Add Application Group...** from the **Actions** menu.
5. The **Add Application Group Wizard** displays. The following table shows the required fields and actions for the **Welcome** step:

| *Table 30. Welcome* | |
|---|---|
| **Field** | **Action** |
| Name | Enter a name to identify the system on the AD FS instance. If you are adding multiple systems, each name must be unique. For example, specify the actual name of the IBM Storage Virtualize system. |
| Description | Enter a brief description of the system. |
| Template | Select **Server application accessing a web API**. |

6. Click **Next**.
7. The following table shows the required fields and actions for the **Server application** step:

| *Table 31. Server application* | | |
|---|---|---|
| **Field** | **Action** | **Details** |
| Name | The wizard automatically adds "Server application" to the name that you specified on the **Welcome** page. If multiple systems need to be added, you can change this to include other identifying information. | |
| Client Identifier | This value is automatically generated when the system is saved as an application. Copy this value and save it. You need to add this value on the **Configure Web API** page in the wizard and in the management GUI. | This value must be entered to the **Single Sign-on** page in the management GUI under **OpenID Credentials** as the **Client ID**. |
| Redirect URI | Enter the redirect URI for the management GUI. For management GUI access, the redirect URI is consisted of the host name, fully qualified domain name, or management IP address followed by `/sso`. Depending on how you access the system, multiple redirect URIs can be specified. For example:<br>**Hostname**<br>`https://hostname/sso`<br>**Fully-qualified domain name**<br>`https://hostname.com/sso`<br>**IP address**<br>`https://1.2.3.4/sso`<br>**Note:** Single sign-on is not supported on the service assistant interface. | The redirect URI is the location where AD FS sends users after they are successfully authenticated. |

8. Click **Next**.
9. The following table shows the required fields and actions for the **Configure Application Credentials** step.

| *Table 32. Configure Application Credentials* | | |
|---|---|---|
| **Field** | **Action** | **Details** |
| Generate a shared secret | a. Select **Generate a shared secret**.<br>b. After the shared secret is automatically generated, click **Copy to clipboard**.<br>**Note:** This shared secret must be added to the management GUI when you enable single sign-on for the system. After it is generated, you cannot display this secret in the application properties after the system is added to AD FS. Ensure that you retain the copied version of the client secret. The client secret is sensitive so keep it safe. | This value must be entered on the **Single Sign-on** page in the management GUI under **OpenID Credentials** as the **Client secret**. |

10. Click **Next**.
11. The following table shows the required fields and actions for the **Configure Web API** step.

| *Table 33. Configure Web API* | | |
|---|---|---|
| **Field** | **Action** | **Details** |
| Name | The wizard automatically adds "Web API" to the name that you specified on the **Welcome** page. If multiple systems need to be added, you can change this to include other identifying information. | |
| Identifier | a. Copy the automatically generated value in the **Client Identifier** field from the **Server Application** page.<br>b. Click **Add** | |

12. Click **Next**.
13. The following table shows the required fields and actions for the **Choose Access Control Policy** step.

| *Table 34. Choose Access Control Policy* | | |
|---|---|---|
| **Field** | **Action** | **Details** |
| Choose an access control policy | From the options, select one of the following supported access policies:<br>**Permit everyone**<br>Select this option if you are only configuring single sign-on without any additional factors to authenticate users.<br>**Permit everyone and require MFA**<br>Select this option if you are configuring single sign-on with additional factors to authenticate users. | These two options are not the only valid access control policies available to select for the system. These options are suggested since they cover most use cases; however, custom access policies can also be created to suit your needs and security policies. |

14. Click **Next**.
15. The following table shows the required fields and actions for the **Configure Application Permissions** step.

| *Table 35. Configure Application Permissions* | | |
|---|---|---|
| **Field** | **Action** | **Details** |
| Permitted scopes | Select the following scopes:<br>• Select **allatclaims**.<br>• Verify that **openid** is selected. | |

16. Click **Next**.
17. On the **Summary** page, verify the settings and click **Next**.
18. Click **Close** on the confirmation message. The **Application Groups** page displays. The system is now configured as an application group in AD FS.
19. Right click the new application group that you created and select **Properties**.
20. On the **Properties** page, select the name of the system under the Web API section and click **Edit...**.
21. On the **Web API Properties** page, select **Issuance Transform Rules** tab.
22. On the **Issuance Transform Rules** tab, click **Add Rule...**.
23. On the **Choose Rule Type** page, select **Send LDAP Attributes as Claims** as the claim rule template and click **Next**.
24. On the **Choose Claim Rule** page, enter the following values on the page:

    **Claim rule name:**
    Enter a name for the rule that you are setting.

    **Attribute store:**
    Select **Active Directory**.

    **Mapping of LDAP attributes to outgoing claim types:**
    In the **LDAP Attributes** column, select **Token-Groups - Unqualified Names**.

    In the **Outgoing Claim Type** column, select **Group**.
25. Click **Finish**. The **Web API Properties** page displays.
26. On the **Web API Properties** page, click **Apply**.

Before you enable single sign-on on the system, ensure that you have copied and saved the following information:

**Client Identifier**
This value is automatically generated on the **Server Application** task of the **Add Application Group** wizard. This value must be entered on the **Single Sign-on** page in the management GUI under **OpenID Credentials** as the **Client ID**.

**Shared secret**
This value is automatically generated on the **Configure Application Credentials** task of the **Add Application Group** wizard. This value must be entered on the **Single Sign-on** page in the management GUI under **OpenID Credentials** as the **Client secret**.

**OpenID configuration endpoint**
The administrator of your single sign-on server can provide this information. However you can display the OpenID configuration endpoint by selecting **Server Manager** > **AD FS Management** > **AD FS** > **Service** > **Endpoints**

## Using the management GUI

To configure single sign-on on the system, complete these steps:

1. Select **Settings** > **Security** > **Single Sign-on**.
2. For the **Authentication server**, enter the OpenID Connect Discovery endpoint that is associated with the Identity Provider. For example, for an AD FS endpoint, enter the following:

   ```
   https://<AD FS server URL>/adfs/.well-known/openid-configuration
   ```

   where <AD FS server URL> is the fully-qualified domain name for the single sign-on server.
3. Under **OpenID Credentials**, paste the **Client Identifier** that you copied from the **Server application** task from the **Add Application Group** wizard in the **Client ID** field. Paste the client secret that was generated on the **Configure Application Credentials** task in the **Add Application Group** wizard in the **Client secret** field.
4. Under **Claims**, the system automatically sets the **User claim** to upn and **Group claim** to group.
5. Click **Save**. On the confirmation page, click **Confirm** to enable single sign-on for the system.

Single sign-on is enabled for the system. You can configure user groups to use single sign-on. Click **Navigate** to launch the **User Groups** page.

## Using the command-line interface

To enable single sign-on, enter the following command:

```
chauthsinglesignon -oidcconfigurationendpoint https://authentication-server-url/adfs/.well-
known/openid-configuration -clientid xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx -clientsecret
xxxxxxxx -groupclaim group -userclaim upn -enable
```

In the example, authentication-server-url is the URL to your authentication server. The values for the **-clientid** and the **-clientsecret** are **Client Identifier** and the client secret that was generated in the **Add Application Group** wizard in the AD FS management dashboard.

### *Configuring single sign-on with Microsoft Azure AD*

Microsoft Azure AD can be configured as the authentication provider for the system.

**Note:** Ensure that the prerequisite tasks are completed on the system before you configure single sign-on. For more information, see the prerequisites in “Configuring single sign-on” on page 114.

## Prerequisites

The following prerequisite steps on Microsoft Azure AD must be completed before you can configure single sign-on on the system:

1. Create a subscription for Microsoft Azure AD. A free trial subscription is also available. For more information, see Azure Active Directory (Azure AD). During subscription creation, you specify a tenant that is used to create a URL to access the Microsoft Azure AD dashboard.
2. Access the Microsoft Azure AD administrator dashboard by entering the following URL in a web browser:

   ```
   https://portal.azure.com
   ```

3. In the Microsoft Azure AD interface, from the menu select **Azure Active Directory** > **App registrations** > **New registration**.
4. Select **IBM Storage Virtualize** > **Add application**.

   **Note:** Each system must be added as a separate application.

   The following table shows the required fields and actions for the **Register an application** page in the Microsoft Azure AD interface.

   *Table 36. Register an application page*

   | Field | Action |
   |---|---|
   | Name | Name Enter a name to identify the system on Microsoft Azure AD. If you are adding multiple systems, enter a unique name. |
   | Supported account types | Select which user accounts should be able to use this application. |
   | Redirect URI | Select Web. Multiple redirect URIs can be specified for the management GUI. The redirect URI is consisted of the management IP address or hostname followed by `/sso`. For example, `https://hostname.com/sso` |

5. Click **Register**. After the system is saved as a new app registration, the application reloads with the **Overview** page selected.
6. To find the OpenID Configuration Endpoints URL for the tenant, select **Overview** > **Endpoints** and copy the OpenID Connect metadata document URL. This value must be specified as the OpenID Configuration Endpoint URL on the Single sign-on page in the management GUI.
7. On the **Overview** page, copy the **Application (client) ID**. This value must be specified as the Client ID on the Single sign-on page in the management GUI.
8. Select **Manage** > **Certificates & secrets** > **New client secret**. Enter a description and expiry value. Click Add. After the client secret has been added, the Certificates & secrets page appears. Copy the **Secret ID**. This value must be specified as the Client secret on the Single sign-on page in the management GUI.
9. Ensure the authentication provider is configured to send back the group claim in the ID Token sent to the system. The group claim identifies which groups the authenticating user belongs to. Some authentication providers do not send the group claim by default, so this typically requires some configuration on the authentication provider. Select **App registrations** > select the application > **Add groups claim**. Select which group types to include. Select **Customize token properties by type** for the ID token, then choose your preferred method of identifying group names. For example, choosing Group ID means that user groups must be referred to using their numeric ID. Click **Add**.
10. Enable ID Tokens for the application. Select **Manage > Authentication**. For **Implicit grant and hybrid flows**, ensure **ID tokens** (used for implicit and hybrid flows) is selected. Click **Save**.

### Using the management GUI

To configure single sign-on with Microsoft Azure AD, complete these steps:

1. Select **Settings** > **Security** > **Single Sign-on**.
2. Enter the OpenID Configuration Endpoint URL of the authentication server. For Microsoft Azure AD, enter the following:

   ```
   https://login.microsoftonline.com/tenant/v2.0/.well-known/openid-configuration
   ```

   Where `tenant` is the ID of the tenant. This can be found by selecting from the menu **Azure Active Directory** > **Click your application** > **Overview** > **Endpoints** > **OpenID Connect metadata document**.
3. For the **OpenID Credentials**, add the **Client ID** and **Client Secret** that you copied on from the Overview page in the Microsoft Azure AD interface.
4. For the User claim, the value to enter depends on how your authentication provider is configured. The User claim must match the name that the authentication service uses to specify the username attribute in the ID Token it sends to the system. Typically this value is preferred_username, but can be customized on the authentication provider.
5. For the Group claim, the value to enter depends on how your authentication provider is configured. The Group claim must match the name that the authentication service uses to specify the group attribute in the ID Token it sends to the system. Typically this value is groups, but can be customized on the authentication provider.
6. For Proxy server, consider how the system accesses the authentication provider. For an authentication provider within your network, a proxy server usually isn't needed. If you connect to the authentication provider through the Internet, check the box and ensure a proxy server is defined on the system.
7. Click **Save**. On the confirmation page, click **Confirm** to enable single sign-on for the system.

Single sign-on is enabled for the system. You can configure user groups to use single sign-on. Click **Navigate** to launch the **User Groups** page.

### Using the command-line interface

To enable single sign-on, enter the following command:

```
chauthsinglesignon -oidcconfigurationendpoint https://login.microsoftonline.com/tenant/
v2.0/.well-known/openid-configuration-clientid xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx
-clientsecret xxxxxxxx -userclaim preferred_username -groupclaim groups -enable
```

In the example, `tenant` is the tenant name that is associated with your subscription. The values for the **`-clientid`** and the **`-clientsecret`** are the Open ID Client and Open ID Secret that are automatically generated when you created your system as an app registration in Microsoft Azure AD can be obtained from the Overview page in the Microsoft Azure AD interface. The values for the **`-userclaim`** and the **`-groupclaim`** should match the name of the claims configured for the ID Token on the authentication provider.

## *Configuring single sign-on with Okta*

Okta can be configured as the authentication provider for the system.

**Note:** To configure single sign-on with Okta, you must upgrade the system to 8.5.3 or later release.

**Note:** Ensure that the prerequisite tasks are completed on the system before you configure single sign-on. For more information, see the Prerequisites section in "Configuring single sign-on" on page 114.

### Prerequisites

The following prerequisite steps on Okta must be completed before you can configure single sign-on on the system:

1. Create a subscription for Okta. A free trial subscription is also available. For more information, see Okta. During subscription creation, you specify a tenant that is used to create a URL to access the Okta dashboard.
2. Access the Okta administrator dashboard by entering the following URL in a web browser:

   ```
   https://tenant.okta.com/admin/dashboard
   ```

   Where tenant is the name of the tenant that you specified when you created your subscription. Usually this tenant name is associated with your company or organization.
3. In the Okta interface, from the menu select **Applications** > **Applications** > **Create App Integration**.
4. For the Sign-in method, select **OIDC** – **OpenID Connect**. For the Application type, select **Web Application**. Click **Next**.

   **Note:** Each system must be added as a separate application.

   The following table shows the required fields and actions for the **General Settings** page in the Okta interface

| *Table 37. General settings on the app integration page* | |
|---|---|
| **Field** | **Action** |
| App integration name | Enter a name to identify the system on Okta. If you are adding multiple systems, enter a unique. |
| Grant type | Ensure **Authorization Code** is selected. |
| Sign-in redirect URIs | Multiple redirect URIs can be specified for the management GUI. The redirect URI is consisted of the management IP address or hostname followed by `/sso`. For example, `https://hostname.com/sso`. |

   The following table shows the required fields and actions for the **Assignments** page in the Okta interface.

| *Table 38. Assignments settings on the app integration page* | |
|---|---|
| **Field** | **Action** |
| Controlled access | Select which users or groups can access the app. If you are unsure, select **Allow everyone** . |
| Enable immediate access | Ensure **Enable immediate access with Federation Broker Mode** is disabled |

   Click **Save**. After the system is saved as a new app integration, the application reloads with the General page selected.
5. On the **General** page, copy the **Client ID**. This value must be specified as the Client ID on the Single sign-on page in the management GUI.
6. On the **General** page, ensure **Client authentication** is set to **Client secret**.
7. On the **General** page, ensure **Proof Key for Code Exchange (PKCE)** is disabled. The system does not support PKCE.
8. On the **General** page, copy the **Client Secret** or create a new one. This value must be specified as the Client ID on the Single sign-on page in the management GUI.
9. Ensure the authentication provider is configured to send back the group claim in the ID Token sent to the system. The group claim identifies which groups the authenticating user belongs to. Some authentication providers do not send the group claim by default, so this typically requires some configuration on the authentication provider. From the menu, select **Security** > **API** > **Authorization**

**Servers**. Edit the default authorization server. Under the Claims tab, edit the claim type **ID**. Add a new claim called **groups** with **Include in token type** set to **ID Token / Always**, a **Value type** of Groups, **a Filter** of **Matches regex** with value .* and an **Include in** of **Any scope**.

10. Click **Save**.

## Using the management GUI

To configure single sign-on with Okta, complete these steps:

1. Select **Settings** > **Security** > **Single Sign-on**.
2. 2. Enter the OpenID Configuration Endpoint URL of the authentication server. For Okta, enter the following:

   ```
   https://tenant.okta.com/oauth2/default/.well-known/openid-configuration
   ```

   where `tenant` is the hostname of the tenant.
3. For the **OpenID Credentials**, add the **Client ID** and **Client Secret** that you copied on from the Overview page in the Okta interface.
4. For the User claim, the value to enter depends on how your authentication provider is configured. The User claim must match the name that the authentication service uses to specify the username attribute in the ID Token it sends to the system. Typically this value is preferred_username, but can be customized on the authentication provider.
5. For the Group claim, the value to enter depends on how your authentication provider is configured. The Group claim must match the name that the authentication service uses to specify the group attribute in the ID Token it sends to the system. Typically this value is groups, but can be customized on the authentication provider.
6. For Proxy server, consider how the system accesses the authentication provider. For an authentication provider within your network, a proxy server usually is not needed. If you connect to the authentication provider through the internet, check the box and ensure a proxy server is defined on the system.
7. Click **Save**. On the confirmation page, click **Confirm** to enable single sign-on for the system.

Single sign-on is enabled for the system. You can configure user groups to use single sign-on. Click **Navigate** to launch the **User Groups** page.

## Using the command line interface

To enable single sign-on, enter the following command:

```
chauthsinglesignon -oidcconfigurationendpoint https://tenant.okta.com/oauth2/default/.well-
known/openid-configuration-clientid xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx -clientsecret xxxxxxxx
-userclaim preferred_username -groupclaim groups -enable
```

In the example, `tenant` is the tenant name that is associated with your subscription. The values for the **`-clientid`** and the **`-clientsecret`** are the Open ID Client and Open ID Secret that are automatically generated when you created your system as an app integration in Okta can be obtained from the General page in the Okta. The values for the **`-userclaim`** and the **`-groupclaim`** should match the name of the claims configured for the ID Token on the authentication provider.

### *Configuring single sign-on with Duo Security*

Duo Security can be configured as the authentication provider for the system.

Duo Security does not support acting as an Identity Provider (IdP) for single sign-on. To use Duo for single sign-on, an alternative authentication source must be configured. Refer to the Duo Security documentation for Configure Your Authentication Source. You might also need to refer to the documentation for your chosen authentication source in order to integrate it with Duo Security.

**Note:** Ensure that the prerequisite tasks are completed on the system before you configure single sign-on. For more information, see the Prerequisites section in “Configuring single sign-on” on page 114.

## Prerequisites

The following prerequisite steps on Duo Security must be completed before you can configure single sign-on on the system:

1. Create a subscription for Duo Security. A 30-day free trial subscription is also available. For more information, see Sign Up for a Free Trial | Duo Security. During subscription creation, you specify a tenant that is used to create a URL to access the Duo Security dashboard.
2. Access the Duo Security administrator dashboard by entering the following URL in a web browser:

   ```
   https://admin-tenant.duosecurity.com
   ```

   where `tenant` is the name of the tenant that you specified when you created your subscription. Usually, the tenant name is associated with your company or organization.
3. In the Duo Security interface, select **Applications** > **Protect an Application**.
4. Select **Generic OIDC Relying Party**.

   **Note:** Each system must be added as a separate application.

   The new application loads. The following table shows the required fields and actions for the **Metadata** section in the Duo Security interface.

*Table 39. Metadata section*

| Field | Action | Details |
|---|---|---|
| Client ID | This value is automatically generated when the system is saved as an application. | This value must be entered on the Single Sign-on page in the management GUI under **OpenID Credentials**. |
| Client Secret | This value is automatically generated when the system is saved as an application. | This value must be entered on the Single Sign-on page in the management GUI under **OpenID Credentials**. |
| Discovery URL | | Indicates the OpenID Connect configuration end point URL of the authentication server. This value must be entered on the Single Sign-on page in the management GUI under **Authentication server**. |

The following table shows the required fields and actions for the **Relying Party** section in the Duo Security interface.

*Table 40. Relying Party section*

| Field | Action | Details |
|---|---|---|
| Grant Type | Select **Authorization Code**. | **Authorization code** indicates that the client can request access to protected resources on behalf of users. |

| *Table 40. Relying Party section (continued)* | | |
|---|---|---|
| **Field** | **Action** | **Details** |
| Sign-in Redirect URLs | Enter the locations where the authorization server sends users after they are successfully authorized and granted an authorization code or access token. | Multiple redirect URIs can be specified for the management GUI. For management GUI access, the redirect URI consists of the management IP address or hostname followed by `/sso`. For example,<br>`https://hostname.com/sso` |

**Note:** Duo Security does not support the use of IP addresses when accessing your system by using a web browser. Once single sign-on is configured, make sure that a hostname is used when accessing your system by using a web browser.

The following table shows the required fields and actions for the **OIDC Response** section in the Duo Security interface. In this section, you configure the OpenID Connect response from Duo Security to the system during authentication. Depending on the name of the IdP attributes on your authentication source, you create mappings between the IdP attributes (for example a SAML attribute sent back from your authentication source) and the claims that Duo Security sends back in the OIDC response.

| *Table 41. OIDC Response section* | | |
|---|---|---|
| **Field** | **Action** | **Details** |
| Scopes | Select **openid** and **profile**. | After selecting profile, a list of IdP Attribute and Claim mappings will appear.<br>Ensure at least two attributes (username and group) are mapped correctly to claims in the OIDC response. For example, to map an IdP Attribute called **Username** from your authentication source, enter **Username** into the IdP Attribute field.<br>You can choose the name of the claim that gets sent back in the OIDC response to the system. For example, to map the IdP Attribute called **Username** to an OIDC claim called **groups**, enter groups in the Claim field. |
| Additional Scopes | No action is required. | Additional scopes are not required. |

Optionally, configure the policy settings for this application in the **Policy** section.

The following table shows the required fields and actions for the **Settings** section in the Duo Security interface.

| *Table 42. Settings section* | | |
|---|---|---|
| **Field** | **Action** | **Details** |
| Name | Enter a name. | Enter a name to identify the system on Duo Security. If you are adding multiple systems, enter a unique name. |
| Additional Scopes | No action is required. | Additional scopes are not required. |

5. Click **Save**. After the system is saved as a new application, the application reloads with the chosen settings.

## Using the management GUI

To configure single sign-on with Duo Security, complete these steps:

1. Select **Settings** > **Security** > **Single Sign-on**.
2. Enter the OpenID Configuration Endpoint URL of the authentication server. This is the Discovery URL from the Metadata settings of the application you created earlier in Duo Security. For Duo Security, it is in the following format:

   ```
   https://sso-tenant.sso.duosecurity.com/oidc/clientid/.well-known/openid-configuration
   ```

   where `tenant` is the name that is associated with your subscription, and `clientid` is the client ID that is associated with your application.
3. For the **OpenID Credentials**, add the **Client ID** and **Client Secret** that you copied on from the Sign-on tab in the Duo Security interface.
4. For the User claim, the value to enter depends on how your authentication provider is configured. The User claim must match the name that the authentication service uses to specify the username attribute in the ID Token of the OIDC response it sends to the system. Use the **Claim** name that maps to your username **IdP Attribute** from the OIDC Response section of your application.
5. For the Group claim, the value to enter depends on how your authentication provider is configured. The Group claim must match the name that the authentication service uses to specify the group attribute in the ID Token of the OIDC response it sends to the system. Use the **Claim** name that maps to your group **IdP Attribute** from the OIDC Response section of your application.
6. For Proxy server, consider how the system accesses the authentication provider. For an authentication provider within your network, a proxy server usually is not needed. If you connect to the authentication provider through the Internet, check the box, and ensure that a proxy server is defined on the system.
7. Click **Save**. On the confirmation page, click **Confirm** to enable single sign-on for the system.

Single sign-on is enabled for the system. You can configure user groups to use single sign-on. Click **Navigate** to launch the **User Groups** page.

## Using the command-line interface

To enable single sign-on, enter the following command:

```
https://sso-tenant.sso.duosecurity.com/oidc/clientid/.well-known/openid-configuration
```

In the example, `tenant` is the tenant name that is associated with your subscription, and `clientid` is the client ID that is associated with your application. The values for the **`-clientid`** and the **`-clientsecret`** are the Open ID Client and Open ID Secret that are automatically generated when you created your system as an application in Duo Security, and are displayed in the Metadata section of the application on the Duo Security interface. The values for the **`-userclaim`** and the **`-groupclaim`** must match the name of the claims that are configured for the ID Token in the **`OIDC Response`** section of the application on the authentication provider.

### *Configuring user groups for single sign-on*

After you have configured the system to use single sign-on, you can configure user groups to use single sign-on. Like LDAP, you must create the remote groups with names which match the name of a group on the Identity provider.

#### Using the management GUI

To enable single sign-on for user groups, complete these steps:

**For existing user groups:**

1. In the management GUI, select **Access** > **Users by Group**.
2. Select the user group from the left navigation and select **User Group Actions** > **Properties**. For each user group on the authentication service, a corresponding user group must be created with the same name.
3. On the **User Group Properties** page, select **Single sign-on** under **Remote Authentication**.
4. Click **OK**.

**For new user groups**

1. Select **Access** > **Users by Group** > **Create User Group**.
2. On the **Create User Group** page, enter the following information:

   **Group Name**
   Enter the name of the group that is from AD FS. The name of the group on the system must match.

   **Ownership Group**
   If ownership groups are configured on your system, you can select an ownership group for the user group.

   **Remote Authentication**
   Select **Single Sign-on**.

   **Role**
   Select a role for the user group. The role determines privileges that users are granted when they are assigned to the user group.
3. Click **Create**.

#### Using the CLI

To configure user groups for single sign-on, complete these steps:

**For existing user groups**

```
chusergrp -remote yes <id or name of group>
```

**For new user groups for remote users**

```
mkusergrp -name <name> -role <role> -remote
```

## Encryption

To use encryption on the system you must purchase and activate encryption licenses, set up your method of key management, and then create encrypted objects.

### Encryption overview

Encryption is a technology that uses cryptography to ensure confidentiality of sensitive information. Encryption uses keys to encode information so that it cannot be understood by unauthorized parties.

Depending on your model, the system supports both encryption of data-at-rest and encryption of data-in-flight.

To use encryption of data-at-rest on the system, you need to:

- Purchase and activate licenses for the encryption feature.
- Decide which methods to use for managing the main keys.
- Configure encryption with the chosen key management method(s).
- Create encrypted objects such as arrays, pools, or cloud accounts.

To use encryption of data-in-flight on the system, you need to:

- Purchase and activate licenses for the encryption feature.
- Create secure IP partnerships between systems.

Configuring encryption is a nondisruptive procedure. During the procedure, the system continues to process I/O operations normally and the existing storage objects are not impacted.

## Encryption of Data at Rest (EDaR)

EDaR is the encryption of static data that is being stored on an internal drive or external storage. Data is automatically encrypted as it is written to the storage, and automatically decrypted as it is read from the storage. This feature protects against the potential exposure of sensitive user data and user metadata that is stored on discarded, lost, or stolen storage devices. Depending on your model, the system supports data encryption that uses encryption-capable hardware and software.

To use EDaR the encryption feature must be licensed and configured on the system.

EDaR is performed that uses the symmetric Advanced Encryption Standard (AES) algorithm with 256-bit encryption keys in XTS mode (XTS-AES-256), as defined in the IEEE 1619-2007 standard and NIST Special Publication 800-38E. The data encryption keys are protected by use of NIST's AES key wrap that uses an intermediate 256-bit wrapping key. This intermediate key is in turn protected that uses NIST's AES key wrap that uses an intermediate 256-bit wrapping key. The access keys are managed that uses whichever key management methods are enabled on the system. The wrapped keys are stored securely on the system in nonvolatile memory so that they can be accessed when the system starts up. During normal system operation, all unwrapped keys are stored securely in volatile memory, the contents of which are securely discarded on system shutdown or power loss. All data encryption keys (DEKs) and key encryption keys (KEKs) on the system are AES-256 bit keys, and all keys are protected by using an AES wrap-key operation.

All NVMe drives that are supported by the system, including IBM FlashCore Modules (FCMs) and a range of other third-party drives, are self-encrypting drives (SEDs) that encrypt data within the electrical circuit of each individual drive, with no performance penalty. Depending on the model, IBM FlashCore Modules are either FIPS 140-2 Level 2 or FIPS 140-3 Level 3 validated. For drives that support FIPS 140-3, the system must be running software version 8.6.2 or later to enable this.

When drives are connected through a Serial Attached SCSI (SAS) network, the SAS protocol chip provides data encryption capabilities, with no performance penalty. The SAS chip uses data encryption algorithms that are FIPS 140-2 Level 1 compliant.

The system software also can apply encryption to data on external storage devices that do not support built-in encryption. In this scenario, the software offloads the job of data encryption to the Advanced Encryption Standard New Instructions (AES-NI)-capable CPU within the hardware, with a small performance penalty, which depends on your configuration. The software uses algorithms that are FIPS 140-2 Level 1 compliant.

EDaR is always applied after data reduction technologies, such as compression and deduplication.

## Encryption of Data in Flight (EDiF)

EDiF is the encryption of data that is being transmitted from one location to another. Data is encrypted before it is sent over the link and is decrypted when it is received on the other side. This feature protects

against threats such as eavesdropping and Man In The Middle (MITM) attacks. Depending on your model, the system supports data encryption that uses encryption-capable hardware and software.

To use EDiF the encryption feature must be licensed, but there is no need to configure encryption on the system.

The system supports EDiF for data that is being replicated between two systems that are connected by Ethernet and configured in a secure Internet Protocol (IP) partnership. When a secured IP partnership is created, for example, between a production system and a recovery system, the data is secured as it travels through the network between the production system and the recovery system. Secured IP partnerships use a combination of IPsec and IKEv2 to secure data in flight. IKEv2 is an IPsec-based tunnelling protocol that uses secure key exchange algorithms to establish a secure connection to the partner system. IPsec is a suite of security protocols that ensures packets that are transmitted over the network are authenticated and encrypted.

In secured IP partnerships, the partner systems authenticate with each other, negotiate the security parameters, exchange encryption keys, and establish secured network tunnels through which encrypted data travels. Partner systems are authenticated by certificates that are issued by either the system's internal root CA or a trusted third-party root CA or intermediate CA.

The system does not support EDiF between host servers and the system, the system and back-end storage devices, or systems that are connected by FC that are configured in an FC replication partnership.

## Licensing encryption

To use encryption, an encryption feature license must be purchased and activated for every machine in the system. A machine is defined to be either a control enclosure or a node, depending on the platform and model used. As an example, an eight-node system consisting of four IBM Storage FlashSystem 9200 control enclosures would require four feature licenses. Alternatively, an eight-node system composing of eight IBM SAN Volume Controller nodes would require eight feature licenses. Licenses for the encryption feature are only made available in countries that permit the use of encryption technologies. Encryption can be configured once all required feature licenses have been activated. See for more details.

## Configuring key management

To configure encryption on the system, the user must have the SecurityAdmin role. When configuring encryption on the system, one or more key management methods must be used to manage the main keys for the system.

The system supports the following key management methods:

- 
- 

The system also supports an encryption recovery key, which can be used as a backup method alongside any other key management methods, to ensure that encrypted data is available when there is a problem accessing the main keys.

For organizations with strict security policies regarding USB flash drives, the system supports disabling the USB ports to prevent unauthorized transfer of system data to portable media devices. If you have such security requirements, consider that uses key servers to manage encryption keys instead.

It is possible to simultaneously configure USB flash drives and key servers to ensure that access to encrypted data is retained if either method is inaccessible, or if the keys are permanently lost for one of the methods.

**Note:** To protect against permanent key loss for one of the methods, a simultaneous configuration must be planned. It is not permitted to enable another key method when the keys for an existing method have already been lost.

If your system contains existing volumes in nonencrypted pools, you can migrate these volumes to encrypted pools after encryption has been configured.

### Using encryption

Once the encryption feature has been configured, the system allows logical configuration objects to be created as encrypted, meaning that the system will automatically encrypt and decrypt the data being stored for that object.

Depending on your configuration, the system will automatically apply the right method of encryption (hardware-based or software-based). Encrypted arrays consisting of internal drives will always use hardware-based encryption and will only use the encryption key belonging to the array, even if the array is part of an encrypted pool. Encrypted pools consisting of external storage will always use software-based encryption and will only use the encryption key belonging to the pool.

### Unlocking an encrypted system

When the system is unlocked and running normally, each node in the system keeps a copy of the main key in secure volatile memory. When a node restarts, the in-memory copy of the key is discarded and another copy is fetched from a partner node, without needing to retrieve the main key from elsewhere. If all the nodes in the system have lost the in-memory copy of the main key simultaneously (in scenarios such as a full system shut down or unexpected power loss), the system tries to retrieve the main key from one of the configured key management methods.

If key servers are configured, the key is retrieved automatically from any key servers that are online and available to the system. If USB flash drives are configured, the key is automatically retrieved from any USB flash drives that are locally installed. Should any of these methods be unavailable at the time the system requires a master key. The system is locked and all encrypted arrays, pools, and cloud accounts are held offline.

**Note:** If all copies of the main key are unavailable when the system is locked, and your system contains encrypted arrays with SAS drives, then all SAS drives connected to the enclosure go offline, even drives belonging to unencrypted arrays.

If the encryption recovery key is configured, the recovery key can be supplied to unlock the system to bring encrypted data online.

The system requires a main encryption key to be available during the following operations:

- System power-on
- System restart
- User-initiated rekey operations
- System recovery

## Licensing encryption

Before you can configure encryption on the system, you must purchase and activate encryption licenses. If you intend to use encryption of data-in-flight to secure IP connections between partnered systems, you also require an encryption license. If you have not already purchased a license, contact a customer representative.

### Before you begin

For systems that support more than one control enclosure, a licensed key for the encryption function must be added to all the control enclosures in the system. To obtain license keys, you need the machine type and model (MTM), serial number (S/N), and machine signature to manually activate the keys. Before you can obtain MTM, S/N, and machine signature, ensure that the control enclosure has been added to the system. These values are required if you are activating keys manually on the system.

Complete the following steps to find machine type and model (MTM), serial number (S/N), and machine signature:

1. In the management GUI, select **Monitoring** > **System Hardware**. The system automatically detects if there is a second control enclosure candidate available.

Click **Add Enclosure** and complete the wizard to add the second control enclosure to the system.

2. After the second control enclosure is added to the system, both control enclosures display. For each control enclosure, select the expand icon to open the System Hardware - Enclosure Details page.
3. On the System Hardware - Enclosure Details page, select **Enclosure Actions** > **Properties**. The machine type and model (MTM), serial number (S/N), and machine signature display on the Properties page. Complete this for both control enclosures.

## Using the management GUI

Within the management GUI, there are two ways to activate an encryption license on the system. During system setup, you are prompted to either manually or automatically activate the license on the system. Automatic activation requires that the notebook that is being used to activate the license is connected to an external network.

If you purchased a license after system setup is completed, go to **Settings** > **Systems** > **Licensed Functions** and click to expand **Encryption Licenses**. These instructions assume that system setup is completed.

If you completed system setup and want to activate an encryption license, complete these steps:

1. In the management GUI, select **Settings** > **Systems** > **Licensed Functions**.
2. Click to expand **Encryption Licenses** and select the control enclosure on which to activate the license. You can select manual activation of encryption.
3. To activate encryption manually, complete these steps:
   a. Select the control enclosures on which encryption will be activated, and select **Actions** > **Activate License Manually**.
   b. On the Activate License Manually page, you must retrieve license keys by completing the form at https://www.ibm.com/storage/dsfa/. To complete the form by selecting your product and entering machine type and model, serial number, machine signature, and authorization code that was sent in your license agreement. Copy or download the keys.
   c. Click **Activate**.

## Using the command-line interface

You can use the CLI command to either activate the key directly or provide a path to the file where the key resides. As with activation through the management GUI, you need to use the authorization code that you received with your purchase agreement to obtain the key. If you purchased a license after system setup is completed, use the `Licensed Functions` option to activate the license. You can also use either of the following commands to activate an encryption license on the system:

1. To activate the license by using the key directly, enter the **`activatefeature -licensekey`** command, where key is the license key to activate a feature. The key consists of 16 hexadecimal characters that are organized in four groups of four characters with each group separated by a hyphen (such as 0123-4567-89AB-CDEF).
2. To activate the license with a file path that stores the key, complete these steps:
   a. Use scp to copy the license key file (2076_XXXXXXX.xml) to the /tmp directory.
   b. Using the command-line interface, enter the **`activatefeature -licensekeyfile filepath`**, where filepath is full path-to-file that contains all required license information (such as /tmp/keyfile.xml).

   For more information, see **`Licensing and featurization commands`**.

## Encryption with key servers

A key server is a centralized system that generates and manages encryption keys that are used by the system. Key servers are ideal in environments with many systems, since key servers send keys to the system automatically over the network without requiring physical access to the systems.

Some key servers support replication of keys among multiple key servers. If multiple key servers are supported, you can specify up to four key servers that connect to the system over both a public network or a separate private network.

The system supports the Key Management Interoperability Protocol (KMIP), which is a standard for encryption of stored data and management of cryptographic keys.

The system supports IBM Security Guardium® Key Lifecycle Manager, Thales CipherTrust Manager, or Gemalto SafeNet Key Secure key servers to handle key management on the system. These supported key server management applications create and manage cryptographic keys for the system and provide access to these keys through a certificate. Only one type of key server management application can be enabled on the system at a time. Mutual TLS authentication takes place when certificates are exchanged between the system and the key server. Certificates must be managed closely because expired certificates can cause system outages. Key servers must be installed and configured before they are defined on the system.

One of the key servers defined on the system must be designated as the primary key server. The primary key server is used by the system to create new encryption keys during a rekey operation. All key servers defined on the system are used to fetch the current encryption key when required. When using key servers to manage the master key for the system, a copy of the master key is stored on each defined key server. The master key is a 256-bit AES key, generated by the key server.

When key management is configured using key servers, the system automatically fetches the key from each key server when required. Additionally, every 30 minutes the system will automatically validate that each configured key server is accessible and able to provide the current key to the system. The validation happens on every node in the system, to every key server defined on the system. A key server can have one of three statuses:

- **Online**: the key server is accessible and able to provide the current encryption key to all nodes in the system
- **Degraded**: the key server is accessible and able to provide the current encryption key to only some nodes in the system
- **Offline**: the key server is not accessible and cannot provide the current encryption key to any node in the system

For the supported list of key servers, refer to Supported Key servers - IBM Storage Virtualize.

## Using key servers with IBM Security Guardium Key Lifecycle Manager

The system supports different types of key server configurations on IBM Security Guardium Key Lifecycle Manager. The following configurations are supported:

- IBM Security Guardium Key Lifecycle Manager key servers designate one primary key server, which can have up to three secondary key servers (also known as clones) defined. These additional key servers support more paths when it delivers keys to the system. However, during rekey operations, only the path to the primary key server is used. When the system is rekeyed, secondary key servers are not used until the primary key server replicates the new keys to these secondary key servers. Replication must be complete before keys can be used on the system. You can either schedule automatic replication or complete it manually with IBM Security Guardium Key Lifecycle Manager. During replication, key servers are not available to distribute keys or accept new keys. The total time that it takes for a replication to complete on the IBM Security Guardium Key Lifecycle Manager depends on the number of key servers that are configured as clones. If replication is triggered manually, the IBM Security Guardium Key Lifecycle Manager issues a completion message when the replication completes. Verify that all key servers contain replicated key and certificate information before keys are used on the system.

- Key servers can also be configured with multiple primary key servers where each key server can create new encryption keys. In this instance, any server can be set as the primary key server. The primary key server is the key server that the system uses when you create any new key server encryption keys. If multiple primary servers are enabled on the IBM Security Guardium Key Lifecycle Manager, the key is immediately replicated to the other key servers in the configuration.

For more information about the supported versions, see the IBM Documentation for IBM Security Guardium Key Lifecycle Manager.

When you create key server objects on the system for IBM Security Guardium Key Lifecycle Manager key servers, you must create a device group, in addition to name, IP address, port, and certificate information. The *device group* is a collection of security credentials (including keys and groups of keys) that allows for restricted management of subsets of devices within a larger pool. The system must be defined on the key server to the **SPECTRUM_VIRT** device group if you are using the default settings. If the **SPECTRUM_VIRT** device group does not exist on the key server, it must be created based on the GPFS device family. If you are configuring multiple key servers, the **SPECTRUM_VIRT** device group must be defined on the primary and all additional key servers.

If you are using IBM Security Guardium Key Lifecycle Manager to create and manage keys, ensure that you are using version 2.7.0 or later. If you are using version 2.7, the system supports one master (primary) key server and secondary key servers. Keys are not available until replication is completed between the key servers. If you use version 3.0 or higher, the system supports multiple master key servers, which automatically replicates keys to all configured key servers.

## Using key servers with Thales CipherTrust Manager and Gemalto SafeNet KeySecure

Thales CipherTrust Manager and Gemalto SafeNet KeySecure key servers also supports KMIP and creates keys on demand, sharing with the other clustered servers, providing redundant access. The system supports different types of configurations on key servers. The following configurations are supported:

- Thales CipherTrust Manager and KeySecure key servers use an active-active model, where multiple key servers are used to provide redundancy. In these configurations one key server must be specified as the primary key server. The primary key server is the key server that the system uses when you create any new encryption keys. The key is immediately replicated to the other key servers in the cluster. All of the key servers that are defined on the system can be used to retrieve keys. Although it is possible to configure a single key server instance, two key servers are recommended to ensure availability of keys, if one key server experiences an outage.
- The system supports up to four key servers. If the system is accessing multiple key servers, they need to belong to the same cluster of key servers.

If you are using Gemalto SafeNet KeySecure key servers to create and manage keys, determine whether the system needs a user name and password to authenticate to the KeySecure key servers. If you plan to use a username and password to authenticate the system to these key servers, you must configure user credentials for authentication in the key server management interface. For KeySecure versions of 8.10 and up, administrators can configure a username and password to authenticate the system when it connects. Before version KeySecure 8.10, the use of a password is optional.

If you are using Thales CipherTrust Manager key server to create and manage keys, determine whether the system needs a username and password to authenticate to the CipherTrust Manager key servers. If you plan to use a username and password to authenticate the system to these key servers, you must configure user credentials for authentication in the key server management interface. After configuration, you can select your username from the interface. For more information on migrating key servers, refer to .

### *Enabling encryption with key servers*

You can use the management GUI or the command-line interface (CLI) to enable encryption with key servers.

## Configuring encryption with IBM Security Guardium Key Lifecycle Manager key servers

Ensure that you complete the following tasks on the IBM Security Guardium Key Lifecycle Manager before you enable encryption:

1. The system supports only TLS versions 1.2 and 1.3. In the IBM Security Guardium Key Lifecycle Manager, specify **TLSv1.2** to use the TLS1.2 protocol, or TLSv1.3 to use the TLS1.3 protocol.
2. Ensure that the DB2 database service is started automatically on startup.
3. Ensure that a valid SSL certificate from IBM Security Guardium Key Lifecycle Manager is installed on the system and in use. If automatic replication is configured on IBM Security Guardium Key Lifecycle Manager, then this certificate needs to be uploaded to the system one time. However, if automatic replication is not configured on the IBM Security Guardium Key Lifecycle Manager, a certificate for each stand-alone key server must be uploaded to the system.
4. Ensure that a device group exists on IBM Security Guardium Key Lifecycle Manager called **SPECTRUM_VIRT**, which is based on the GPFS family. If you are configuring multiple key servers, the **SPECTRUM_VIRT** device group must be defined on the primary and all secondary key servers.
5. If encryption is enabled with USB flash drives, insert at least one of the USB flash drives into the system before key servers can be configured for managing keys.

For more information about completing these tasks, see the IBM Documentation for IBM Security Guardium Key Lifecycle Manager.

To enable encryption with a IBM Security Guardium Key Lifecycle Manager key server in the management GUI, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Encryption**.
2. Click **Enable Encryption**.
3. On the **Welcome** page, select **Key Servers**. Click **Next**.

   **Note:** You can also select both **Key Servers** and **USB flash drives** to configure both methods to manage encryption keys. If either method becomes disabled, you can use the other method to access encrypted data on your system.
4. Select IBM Security Guardium Key Lifecycle Manager **(with KMIP)** for the key server type.
5. Enter the name, IP address or domain name, and port for each key server. If you are configuring multiple key servers, the first key server that you specify is the primary key server. If you specify a fully qualified domain name, a DNS server must be configured on your system. To configure a DNS server for the system, select **Settings** > **Network** > **DNS**.
6. Select **SPECTRUM_VIRT** for the device group for the key servers. This device group must also be configured on each of the key servers for the system.
7. On the **Key Server Certificate** page, you must upload all the necessary key server certificates to the system. The key server certificate can be the key server endpoint certificate, the root CA certificate, or a file that contains all CA certificates within that chain. This file does not need to include the key server certificate, it should only have the intermediate and root CA certificates. In case both the endpoint certificate and the CA certificate of key servers are installed on the system, the endpoint certificate takes priority over the CA certificate. If the key servers are configured for automatic replication, the certificate is copied from the primary key server to all secondary key servers. All the IBM Security Guardium Key Lifecycle Manager instances are connected over secure connections with the same key server certificate. If replication is used on the IBM Security Guardium Key Lifecycle Manager, only one key server certificate needs to be installed on the system. The IBM Security Guardium Key Lifecycle Manager uses this single certificate to replicate keys with each other. If only one certificate is used and automatically replicated to all configured key servers, select the certificate

that you downloaded to the system in the **Certificate** field. If automatic replication is not configured, select all the valid certificates that you downloaded to the system for each of the configured key servers. Click **Next**.

8. On the **System Certificate** page, click **Export system certificate** and download the system certificate. Copy the system certificate to the truststore for the **SPECTRUM_VIRT** device group on each configured key server. You must not upload the root certificate to the key servers, as IBM Security Guardium Key Lifecycle Manager does not support chain of trust checking for the SPECTRUM_VIRT device group. For more information, see online documentation of IBM Security Guardium Key Lifecycle Manager.
9. If you have USB flash drives configured as your encryption method, the **Disable USB Encryption** page displays. If you want to migrate to key servers and disable USB flash drives, select **Yes**. If you want keep both encryption methods, click **No**.
10. Click **Next**.
11. On the **Summary** page, verify the configuration for the key servers and click **Finish**.

To enable encryption with an IBM Security Guardium Key Lifecycle Manager key server in the command-line interface, complete the following steps:

1. To enable encryption on your system, see **chencryption** command.
2. To enable the key server type and supply the certificates of the key server, see **chkeyserverisklm** command.
3. To use an internally signed certificate or an externally signed certificate, see **satask exportrootcertificate** command and **chsystemcert** command.
4. To create the primary key server and up to three more secondary key servers and specify the key server certificate, see **mkkeyserver** command.
5. To verify that the system is prepared, see **lsencryption** command.

## Configuring encryption with Thales CipherTrust Manager or Gemalto SafeNet KeySecure key servers

For SafeNet KeySecure key servers, ensure that you complete the following tasks before you enable encryption:

1. Each key server must be configured to allow TLS 1.2 for secure communications.
2. Ensure that a valid SSL certificate from each KeySecure key server is installed on the system and in use. Either add the server certificate for each KeySecure key server, or add the root CA certificate that was used to sign each server certificate.
3. If you plan to use a username and password to authenticate the system to these key servers, you must configure user credentials for authentication in the key server management interface. For KeySecure versions of 8.10 and up, administrators can configure a username and password to authenticate the system when it connects. Before version KeySecure 8.10, the use of a password is optional. To set up authentication with a username and password between the system and KeySecure key servers, disable global keys on the **High Security** menu in the SafeNet KeySecure interface. When global keys are disabled, key servers cannot authenticate clients to create or access keys without valid credentials.
4. The Storage Virtualize certificate must be trusted by the SafeNet KeySecure key servers. If the system's root CA is used to sign the certificate, then the system's root certificate must be installed as an External CA in SafeNet KeySecure and added to the list of CAs that can be used for KMIP. Alternatively, the system certificate can be signed by a trusted third-party CA. The third-party root certificate must be installed as an External CA in SafeNet KeySecure and added to the list of CAs that can be used for KMIP. If Storage Virtualize certificate is self-signed, then the self-signed certificate must be installed as an External CA in SafeNet KeySecure and added to the list of CAs that can be used for KMIP. It is recommended to not use a self-signed certificate, as the connection between Storage Virtualize and the key servers is interrupted when the certificate is renewed, until the new certificate is added to the key servers.

5. If you currently have encryption that is enabled with USB flash drives, at least one of the USB flash drives must be inserted into the system before key servers can be configured for managing keys.

To enable encryption with a Thales CipherTrust Manager or KeySecure key server using the management GUI, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Encryption**.
2. Click **Enable Encryption**.
3. On the **Welcome** page, select **Key Servers**. Click **Next**.

   **Note:** You can also select both **Key Servers** and **USB Flash Drives** to configure both methods to manage encryption keys. If either method becomes unavailable, you can use the other method to access encrypted data on your system.
4. Select Thales CipherTrust Manager or **Gemalto SafeNet KeySecure** for the key server type.
5. Enter the name, IP address or domain name, and port for each key server. If you are configuring multiple key servers, the first key server that you specify is the primary key server. If you specify a fully qualified domain name, a DNS server must be configured on your system. To configure a DNS server for the system, select **Settings** > **Network** > **DNS**.
6. On the **Key Server Credentials** page, enter a user name and password that is used to authenticate the system to the key servers.
7. On the **Key Server Certificate** page, you must upload all the necessary key server certificates to the system. The key servers can use either a certificate from a trusted third party, a self-signed certificate, or a combination of these certificates. All instances are connected over secure connections with the same key server certificate. Either the server certificate for each key server, or the root CA certificate or a file that contains all CA certificates within that chain. This file does not need to include the key server certificate, only the intermediate and root CA certificates. Any server certificates take priority over any CA certificate that is installed on the system for the key servers. Click **Next**.
8. If you are using Thales CipherTrust Manager, the Storage Virtualize certificate must be signed by a CA.
   - If the Storage Virtualize certificate is signed by the system's root CA, click **Export Root Certificate**. Install the root certificate as an external CA on the Thales CipherTrust Manager key servers and add it to the list of external CAs that can be used for KMIP.
   - If the Storage Virtualize certificate is signed by a trusted third-party CA, install the third-party CA's root certificate as an external CA on the Thales CipherTrust Manager key servers and add it to the list of external CAs that can be used for KMIP.

   If you are using SafeNet KeySecure:

   - If the Storage Virtualize certificate is signed by the system's root CA, click **Export Root Certificate**. Install the root certificate as an external CA on the SafeNet KeySecure key servers and add it to the list of external CAs that can be used for KMIP.
   - If the Storage Virtualize certificate is signed by a trusted third-party CA, install the third-party CA's root certificate as an external CA on the SafeNet KeySecure key servers and add it to the list of external CAs that can be used for KMIP.
   - If the Storage Virtualize certificate is self-signed, click **Export System Certificate**. Install the Storage Virtualize certificate as an external CA on the SafeNet KeySecure key servers and add it to the list of external CAs that can be used for KMIP.
9. Select **The system's public key certificate has been transferred to each configured key server**.
10. If you have USB flash drives configured as your encryption method, the **Disable USB Encryption** page displays. If you want to migrate to key servers and disable USB flash drives, select **Yes**. If you want both encryption methods that are configured simultaneously, click **No**.
11. Click **Next**.
12. On the **Summary** page, verify the configuration for the key servers and click **Finish**.

To enable encryption with a Thales CipherTrust Manager or a KeySecure key server in the command-line interface, complete the following steps:

1. To enable encryption on your system, see **chencryption** command.
2. To enable the key server type and supply the root certificate authority (CA) certificate, see **chkeyserverciphertrustmanager** command.
3. To use an internally signed certificate or an externally signed certificate, see **satask exportrootcertificate** command and **chsystemcert** command.
4. To create the primary key server and up to three more secondary key servers and specify the key server certificate, see **mkkeyserver** command.
5. To verify that the system is prepared, see **lsencryption** command.

### *Rekeying a system with key servers*

The key server master key for the system can be changed by performing a rekey operation. During the rekey process, the key server generates a new master key and the existing master key becomes obsolete.

Consider rekeying the key server master key periodically according to your organization's security policy, or when you think the existing key has been compromised or is known by an unauthorized party.

Only a single encryption method can be rekeyed at once. If you have multiple methods of encryption configured on your system, ensure that the current rekey operation is completed before starting another rekey operation.

Ensure that all key servers are online before starting, as the current encryption key will need to be supplied to the system during the rekey process. The key server master key cannot be rekeyed unless the current key is available.

**Note:** If all key servers have lost the current key, you must disable and then re-enable encryption using key servers. For more information, see

If the system is using an encrypted cloud account for Transparent Cloud Tiering, the cloud account must be online during the rekey operation.

If you are using multiple master or active-active key servers, new keys are automatically replicated to all configured key servers. In configurations with a single primary key server and multiple secondary key servers, only the primary key server is updated during the rekey operation. Any additional key servers go offline, and the system reports an error against those key servers until the new key is replicated from the primary to the secondary key servers.

**Note:** To avoid data loss, back up your data on the key server management application every time that you rekey.

## Using the management GUI

Before creating a new key on all configured key servers, the key servers must be online and connected to the system. In the management GUI, select **Settings** > **Security** > **Encryption**. Expand **Key Servers** to display details on all the configured key servers on the system. Verify that the status of the key servers is online and available to the system.

To rekey the system that uses key server encryption, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Encryption**.
2. Expand **Key Servers** to display all the configured key servers on the system and select **Rekey**.
3. Click **OK** on the message dialog. The encryption key is generated by the primary key server and is copied to the primary key server. If errors occur during the rekey process, status messages display problems with the copy or creation of a new key. To determine and fix other possible errors, select **Monitoring** > **Events**.

After the rekey operation completes, the new keys are replicated instantly if you have multiple primary or active-active key server configurations.

### Using the command-line interface

Before creating a new key on all configured key servers, the key servers must be online and connected to the system. In the command-line interface, run the **lskeyserver** command to verify whether the key servers are online and available to the system.

To rekey a system with encryption using key servers, complete these steps:

1. To start the rekey process, enter the following command to prepare a new master key:

   ```
   chencryption -keyserver newkey -key prepare
   ```

   The new key is created successfully when the **keyserver_pmk_rekey_uid** field shows the name of the new key.
2. Commit the new key by running the following command:

   ```
   chencryption -keyserver newkey -key commit
   ```

   For more information, see **chencryption** command.

   The master key has been rekeyed successfully when the **keyserver_pmk_uid** field shows the name of the new key, and the **keyserver_pmk_rekey_uid** field is blank. For more information, see **lsencryption** command.

## *Disabling encryption with key servers*

Encryption with key servers can be disabled using either the management GUI or the command-line interface.

**Note:** For security, encryption methods (including the recovery key) can only be disabled when physically connected to the technician port on the configuration node.

### Using the management GUI

When disabling encryption using the management GUI, encryption using key servers is automatically disabled in the process. See for instructions on disabling encryption using the management GUI.

To disable only key servers, refer to the instructions described in "Using the command-line interface".

### Using the command-line interface

Follow these steps to disable encryption using key servers:

1. Identify the configuration node of the system. For more information, see “Configuration node” on page 5.
2. Connect your computer to the technician port of the configuration node. For more information, see Technician port.
3. In a terminal window, use Secure Shell (SSH) software to connect to the cluster IP address of the system and authenticate using the credentials of any user with the SecurityAdmin role:

   ```
   ssh username@cluster_ip
   ```

   For more information, see **Connecting to the CLI with OpenSSH**.
4. To disable the encryption using key server, enter the following command:

   ```
   chencryption -keyserver disable
   ```

   For more information, see **chencryption** command.

   Encryption using key server has been disabled successfully when the **keyserver_pmk_uid** field is blank and the **keyserver_status** field is licensed. For more information, see **lsencryption** command.

### *Migrating from Gemalto SafeNet KeySecure to Thales CipherTrust Manager key servers*

You can migrate from Gemalto SafeNet KeySecure key servers to Thales CipherTrust Manager key servers non-disruptively with the management GUI. The command-line interface must be used to view the unique ID of the encryption key.

## Prerequisites

Before you migrate key servers, ensure that you complete the following tasks.

**Update system to supported level**
The system must be updated to a level that includes support for Thales CipherTrust Manager key servers.

**Migrate encryption keys from Gemalto SafeNet KeySecure to Thales CipherTrust Manager key servers**
The encryption key that the system uses must be migrated from the SafeNet KeySecure servers to the Thales CipherTrust Manager key servers. Use the **`svcinfo lsencryption`** to view the **`keyserver_pmk_uid`** that identifies the encryption key. The Thales CipherTrust Manager documentation covers migrating encryption keys from SafeNet KeySecure servers toThales CipherTrust Manager servers. You must use these instructions to migrate these encryption keys before you can complete the rest of the migration.

**Configure system certificates on the system**

IBM Storage Virtualize uses certificates to establish a secure connection to encryption key servers. Certificates must be configured before migrating to Thales CipherTrust Manager.

SafeNet KeySecure supports self-signed IBM Storage Virtualize certificates. Thales CipherTrust Manager does not support self-signed certificates, so the IBM Storage Virtualize certificate must be signed by a certificate authority (CA).

The IBM Storage Virtualize certificate can be signed by the system's root CA or by a trusted third-party CA. If the certificate is signed by the system's root CA, the root certificate must be added as an external certificate authority in Thales CipherTrust Manager. If the certificate is signed by a trusted third-party CA, the third-party root certificate must be added as an external certificate authority.

The root certificate must also be installed on the SafeNet KeySecure servers to ensure that the system can communicate with the SafeNet KeySecure servers when the new signed system certificate is installed.

If Thales CipherTrust Manager is configured to require a username in KMIP client certificates, then the username should be included in the IBM Storage Virtualize certificate.

**Download the key server certificate**
Download the key server certificate that is used with the KMIP interface in Thales CipherTrust Manager to your local workstation. You need to upload this certificate to the system during migration.

**Create a Thales CipherTrust Manager username**
By default, Thales CipherTrust Manager is configured to require a username in the 'common name' field of the client's SSL certificate. On Thales CipherTrust Manager, ensure that the following tasks are completed for this username.

1. Create a user with this username.
2. Ensure that this user owns the encryption key that was migrated to the Thales CipherTrust Manager key servers.
3. Ensure that this user is added to the **Key Users** group.

## Using the management GUI

To migrate key servers, complete the following steps.

**Note:** Do not regenerate new encryption keys on the system until after the migration is completed.

1. In the management GUI, select **Settings** > **Security** > **Encryption**.

2. On the **Encryption** page, verify that all the SafeNet KeySecure key servers are online.
3. On the **Certificate** page, under **Key Server certificate authority** , click **Update Certificate**, and select the key server certificate. This key server certificate was downloaded to your local system from the KMIP interface in Thales CipherTrust Manager as part of the prerequisites.
4. Right-click one of the non-primary SafeNet KeySecure key servers and select **Remove**.
5. Select **Add Key Server** and add the key server details for the first CipherTrust Manager key server. Ensure that **Make Primary Key Server** is selected.

   **Note:** A certificate does not need to be added if the certificate authority was updated in Step 3.
6. Repeat steps 4 and 5 until all of the SafeNet KeySecure key servers are replaced byThales CipherTrust Manager key servers.

   **Note:** Do not select **Make Primary Key Server** for the remaining server.
7. The migration is complete. If necessary, you can rekey encryption keys on the system.

### Using the CLI

## Encryption with USB flash drives

USB flash drives are low-cost storage devices that can be used to manage the master encryption key for the system. You can configure encryption and use USB flash drives store local copies of the master encryption key for the system, which can be provided to the system when required by installing them in USB ports.

**Note:** According to your organization's security policy, the USB ports on the system can be disabled to prevent them from being used to store encryption key files.

During configuration of encryption, the system must store a minimum of three copies of the master key on USB flash drives. It is possible to create more than three copies of the master key during the preparation phase of configuration, by physically swapping out the current USB flash drives with new USB flash drives until the wanted number of key copies have been written successfully. It is also possible to make more copies of the master key after configuration, by making a copy of the encryption key file from one of the USB flash drives. It can then be stored on other storage media, or managed using secret management software.

**Note:** The encryption key file written to USB flash drives is highly sensitive and should be managed securely. If copying encryption key files, take care not to make the file accessible to unauthorized parties.

Two options are available for accessing key information on USB flash drives:

**USB flash drives are left inserted in the system at all times**
: If you want the system to bring encrypted data online automatically after a system restart, a USB flash drive must be left installed in one or more canisters in the system. When the system powers on, the encryption key is read from any available USB flash drive installed in the system. This method requires that the physical environment where the system is located is secure. If the location is secure, it prevents an unauthorized person from stealing the USB flash drives to make copies of the encryption keys, as well as stealing system components such as drives. The risk with this approach is that if the location isn't secure and the master encryption keys are attached to the system, then stealing both the system and the USB flash drives means the encrypted data at rest can be accessed.

**USB flash drives are not left inserted in the system**
: For the most secure operation, do not keep the USB flash drives inserted into the canisters in the system. However, this method requires that you manually install the USB flash drives containing the encryption master key in the canisters during certain operations when the system requires an encryption key to be present. USB flash drives that contain the current master key must be stored securely to prevent theft or loss. During operations where the system requires an encryption key to be present, one or more USB flash drives must be installed manually into any canister so data can be accessed. After the system reads the encryption key from USB flash drives, the system will be unlocked and encrypted data will be accessible. After the system is online, the USB flash drives must be removed and stored securely to prevent theft or loss. The advantage of this approach is that even

if the system or drives are stolen, the encrypted data at rest cannot be accessed because the master keys are not present.

### *Enabling encryption with USB flash drives*

The system supports enabling encryption that uses USB flash drives to store encryption keys. USB flash drive-based encryption requires physical access to the systems and is effective in environments with a minimal number of systems. For organizations that require strict security policies regarding USB flash drives, the system supports disabling these ports to prevent unauthorized transfer of system data to portable media devices.

If you have such security requirements, use key servers to manage encryption keys. In addition, if you are using USB flash drives to manage encryption keys but want to disable access to these ports for security reasons, you can migrate to encryption that uses key servers. For more information, see .

If the location is not secure, all USB flash drives with the key can be removed from the system and be stored securely. Extra copies of the key must be created and stored securely to ensure access to the system if the USB flash drives become damaged or stolen. During these operations, you are responsible for ensuring the security of the system. Use these general guidelines when you enable encryption and manage flash drives that contain an encryption key.

1. In addition to the copies that are generated on the USB flash drives when encryption is enabled on the system, make at least one more copy on another USB flash drive and store it in a secure location.
2. In addition, copy the encryption key to other forms of storage to provide resiliency and to mitigate risk, if, for example, the USB flash drives are from a faulty batch of drives.
3. Ensure that each copy of the encryption key is valid before writing any user data to the system. The system validates any key material on a USB flash drive when it is inserted into the canister. If the key material is not valid, the system logs an error. If the USB flash drive is not usable or failed, the system does not display the drive as an active USB flash drive.
4. Securely store all copies of the encryption key. As an example, any USB flash drives that are not left inserted into the system might be locked in a safe. Comparable precautions should be taken to securely protect any other copies of the encryption key stored on other forms of storage.

## Using the management GUI

To enable USB flash drive encryption that uses the management GUI, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Encryption**.
2. Click **Enable Encryption**.
3. On the **Welcome** window, select **USB Flash Drives**, then click **Next**.

   **Note:** You can also select both **Key Servers** and **USB Flash Drives** to configure both methods to manage encryption keys. If either method becomes unavailable, you can use the other method to access encrypted data on your system.
4. In the wizard, you are prompted to insert the required number of USB flash drives into the system. If the system has three or more USB ports, at least three USB flash drives must be entered. If the system has less than three USB ports, a USB flash drive must be entered into every port on the system.
5. When the system detects the USB flash drives, the encryption key is automatically copied to the USB flash drives. Ensure that you create any required extra copies for backups.
6. If the system has less than three USB ports, the USB flash drives containing the new key must be removed from the system and new drives added until the number of key copies is three or more. When a new USB flash drive is added while the rekey is in progress, the key is copied onto it automatically.
7. After all copies are completed, click **Confirm** to complete the enablement process.

   **Note:** You can leave the USB flash drives inserted into the system. However, the area where the system is located must be secure to prevent the USB flash drives from being lost or stolen. If the area where the system is located is not secure, remove all the USB flash drives from the system and store them in a secure location.

## Using the command-line interface

1. To enable the encryption that uses USB flash drives, enter the following command:

   ```
   chencryption -usb enable
   ```

   Encryption that uses USB flash drives has been enabled successfully when the status field shows enabled upon executing **lsencryption** command.
2. Insert the required number of USB flash drives into the system. If the system has three or more USB ports, at least three USB flash drives must be entered. If the system has less than three USB ports, a USB flash drive must be entered into every port on the system.
3. Enter the following command to prepare a new USB master key:

   ```
   chencryption -usb newkey -key prepare
   ```

   The new key has been created successfully when the `usb_rekey_filename` field shows the name of the new key upon executing **lsencryption** command.

   If the system has less than three USB ports, the USB flash drives containing the new key must be removed from the system and new drives added until the number of key copies is three or more. When a new USB flash drive is added while the rekey is in progress, the key is copied onto it automatically.
4. After all copies are completed, run the following command to complete the enablement process:

   ```
   chencryption -usb newkey -key commit
   ```

   The master key has been rekeyed successfully when the `usb_key_filename` field shows the name of the new key and the `usb_rekey_filename` is blank, upon executing **lsencryption** command. For more information, see **lsencryption** command.
5. If an error occurs during the creation of the initial master key creation process, it can be cancelled by running the following command:

   ```
   chencryption -usb newkey -key cancel
   ```

   If the key creation is cancelled, the process must be re-attempted before encryption with USB flash drives is fully enabled on the system. For more information, see **chencryption**.

### *Rekeying a system with USB flash drives*

Rekeying is the process of creating a new key for the system. To create a new key, encryption must be enabled on the system; however, the rekey operation works whether there are encrypted objects or not. Only a single encryption method can be rekeyed at once.

If you have multiple methods of encryption configured on your system, ensure that the current rekey operation is completed before starting another rekey operation. If you are generating new keys for cloud storage, the cloud account must be online during the rekeying operation. Rekeying is a non-disruptive process and can be completed without any impact to availability of encrypted or non-encrypted storage.

Before creating a new key, ensure that at least one USB port contains a USB flash drive that contains the current key. During the rekey process, a new key is generated and copied to the USB flash drives. The new key is then used instead of the current key. The rekey operation fails unless at least one USB flash drive contains the current key. To complete a rekey of the system, you need at least three USB flash drives to store the copied key material.

Once a rekey of the system is complete, the old key material will no longer be of use and cannot be used to unlock the encrypted storage. If any USB flash drives which contained a copy of the old key were not plugged in during the rekey scenario, they will not contain a copy of the new master key and the key will have to be copied to the drive manually.

### Using the management GUI

To rekey the system with encryption using USB flash drives in the management GUI, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Encryption**.
2. Expand **USB Flash Drives** to display all the detected USB flash drives on the system and select **Rekey**.
3. In the wizard, you are prompted to insert the required number of USB flash drives into the system. If the system has three or more USB ports, at least three USB flash drives must be entered. If the system has less than three USB ports, a USB flash drive must be entered into every port on the system.
4. When the system detects the USB flash drives, the encryption key is automatically copied to the USB flash drives. Ensure that you create any required extra copies for backups.
5. If the system has less than three USB ports, the USB flash drives containing the new key must be removed from the system and new drives added until the number of key copies is three or more. When a new USB flash drive is added while the rekey is in progress, the key is copied onto it automatically.
6. After all copies are completed, click **Confirm** to complete the rekey process.
7. If errors occur during the rekey process, status messages display problems with the copy or creation of a new key. For example, if the minimum number of USB drives are inserted but none of them have an existing encryption key, the rekey operation fails. To determine and fix other possible errors, select **Monitoring** > **Events**.

### Using the command-line interface

To rekey the system with encryption using USB flash drives on the command-line interface, complete these steps:

1. Insert the required number of USB flash drives into the system. If the system has three or more USB ports, at least three USB flash drives must be entered. If the system has less than three USB ports, a USB flash drive must be entered into every port on the system.
2. To start the rekey process, run the following command to prepare a new master key:

   ```
   chencryption -usb newkey -key prepare
   ```

   The new key has been created successfully when the **usb_rekey_filename** field shows enabled upon executing **lsencryption** command.
3. If the system has less than three USB ports, the USB flash drives containing the new key must be removed from the system and new drives added until the number of key copies is three or more. When a new USB flash drive is added while the rekey is in progress, the key is copied onto it automatically.
4. After all copies are completed, run the following command to complete the rekey process:

   ```
   chencryption -usb newkey -key commit
   ```

   The master key has been rekeyed successfully when the **usb_key_filename** field shows the name of the new key and the **usb_rekey_filename** is blank upon executing the **lsencryption** command. For more information, see **lsencryption** command.
5. If an error occurs during the rekey process, it can be cancelled by running the following command:

   ```
   chencryption -usb newkey -key cancel
   ```

   If the rekey is cancelled, the system will continue to use the existing master key. For more information, see **chencryption**.

### *Disabling encryption with USB flash drives*

Encryption with USB flash drives can be disabled using the management GUI.

**Note:** For security, encryption methods (including the recovery key) can only be disabled when physically connected to the technician port on the configuration node.

### Using the management GUI

When disabling encryption using the management GUI, encryption using USB flash drives is automatically disabled in the process. See for instructions on disabling encryption using the management GUI.

### Using the command-line interface

Follow these steps to disable encryption using USB flash drives:

1. Identify the configuration node of the system. For more information, see “Configuration node” on page 5.
2. Connect your computer to the technician port of the configuration node. For more information, see Technician port.
3. In a terminal window, use Secure Shell (SSH) software to connect to the cluster IP address of the system and authenticate using the credentials of any user with the SecurityAdmin role:

   ```
   ssh username@cluster_ip
   ```

   For more information, see **Connecting to the CLI with OpenSSH**.
4. To disable the encryption using USB flash drives, enter the following command:

   ```
   chencryption -usb disable
   ```

   For more information, see **chencryption** command.

   Encryption using USB flash drives has been disabled successfully when the **usb_key_filename** field is blank and the **status** field is licensed. For more information, see **lsencryption** command.
5. The encryption key files remain on the USB flash drives but are not used again. The security administrator of the system is responsible for removing unused or expired encryption key files from the USB flash drives.

## Encryption recovery key

The system supports enablement of an encryption recovery key to supplement an existing encryption method (such as USB flash drives or key servers). The encryption recovery key can be used to bring the system’s encrypted storage back online after an outage.

The encryption recovery key is an ASCII string, which should be copied down and stored in a safe location, such as a password manager. Once enabled, the encryption recovery key can be supplied in place of a USB flash drives encryption key or key server encryption key to bring the system’s encrypted storage back online after an outage. The encryption recovery key is intended to supplement an existing encryption method (such as USB flash drives encryption or key server encryption) and should not be used ‘stand-alone’.

### *Enabling the encryption recovery key*

You can use the management GUI or the command-line interface (CLI) to enable the encryption recovery key.

### Using the management GUI

To enable the encryption recovery key while enabling encryption on the system for the first time, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Encryption** and click **Enable Encryption**.
2. Once enabled, the encryption recovery key page will be displayed. Click **Generate recovery key** to generate key and begin the enablement process.

   **Note:** The encryption recovery key will not be redisplayed after completing the wizard.
3. Enter the encryption recovery key into the input field and click **Finish**.

4. If entered correctly, an information message should appear to state the recovery key has been entered correctly. Click **Close**. The encryption recovery key is enabled and can be used to unlock the system.

To enable the encryption recovery key while also enabling a second encryption method (such as USB flash drives encryption or key server encryption), toggle the corresponding drop-down menu on the **Settings** > **Security** > **Encryption** page and click **Configure**. Follow the wizard to enable the selected encryption method.

### Using the command-line interface

To enable the encryption recovery key while both USB flash drives encryption and key server encryption are already enabled, you must use the command-line interface to configure the recovery key. Refer to the **chencryption** command.

Follow these steps to the enable encryption recovery key:

1. Enter the following command to enable the recovery key:

```
chencryption -recoverykey enable
```

2. Enter the following command to prepare a new recovery key:

```
chencryption -recoverykey newkey -key prepare
```

   The new recovery key will be displayed on screen.

   **Note:** The recovery key is sensitive and must be stored in a safe location.

3. To confirm that the recovery key has been stored correctly, the system requires the recovery key to be confirmed. To confirm the recovery key, run the following command and enter the new recovery key when prompted:

```
chencryption -recoverykey newkey -key confirm
Enter the new recovery key for the system:
```

   A confirmation message is displayed when the recovery key has been entered correctly.

4. Commit the recovery key by running the following command:

```
chencryption -recoverykey newkey -key commit
```

### *Rekeying the encryption recovery key*

The encryption recovery key for the system can be changed by performing a rekey operation. During the rekey process, the system generates a new encryption recovery key and the existing recovery key becomes obsolete.

Consider rekeying the encryption recovery key periodically according to your organization's security policy, or when you think the existing recovery key has been compromised or is known by an unauthorized party.

Only a single encryption method can be rekeyed at once. If you have multiple methods of encryption configured on your system, ensure that the current rekey operation is completed before starting another rekey operation.

Ensure that the existing recovery key is available, as it needs to be supplied to the system during the rekey process. The recovery key cannot be rekeyed unless the current recovery key is known and available.

**Note:** If you have lost the recovery key, you must disable and then re-enable the recovery key.

The new encryption recovery key for the system is sensitive, so make sure that the new recovery key is stored securely in a safe location. It is suggested that the rekey procedure is used in a private location, where your browser or terminal window cannot be seen by others.

### Using the management GUI

The system does not currently support rekeying the encryption recovery key by using the management GUI. Refer to the instructions described in "Using the command-line interface" section.

### Using the command-line interface

Follow these steps to rekey the encryption recovery key:

1. The current recovery key must be supplied to the system within 30 minutes of starting a rekey. To validate the current recovery key, enter the following command:

   ```
   chencryption -recoverykey validate
   Enter the recovery key for the system:
   ```

   An interactive prompt is displayed on the screen. When the correct recovery key has been entered, a success message appears.
2. To start the rekey process, enter the following command to prepare a new recovery key:

   ```
   chencryption -recoverykey newkey -key prepare
   ```

   The new recovery key will be displayed on screen.

   **Note:** The recovery key is sensitive and must be stored in a safe location.

   The system creates an identifier for the new recovery key, which can be used as a label when storing the key in a safe location (for example, in a password manager).
3. To confirm that the recovery key has been stored correctly, the system requires the recovery key to be confirmed. To confirm the recovery key, run the following command and enter the new recovery key when prompted:

   ```
   chencryption -recoverykey newkey -key confirm
   Enter the new recovery key for the system:
   ```

   A confirmation message is displayed when the recovery key has been entered correctly.
4. Commit the recovery key by running the following command:

   ```
   chencryption -recoverykey newkey -key commit
   ```

   For more information, see **chencryption** command.

   The recovery key has been rekeyed successfully when the **recovery_key_name** field shows the name of the new recovery key, and the **recovery_key_rekey_name** field will be blank. For more information, see **lsencryption** command.

## *Unlocking the system using the encryption recovery key*

Systems with encryption configured must be unlocked by using an encryption key to bring encrypted storage online.

During startup, the system automatically attempts to fetch an encryption key to unlock the system. Depending on the configuration, the system automatically looks for encryption keys that are stored on locally attached USB flash drives or network-attached key servers. When certain events or outages occur, such as a full system power down, the system must retrieve an encryption key that can unlock the system. If the system is not able to automatically retrieve an encryption key then the encrypted arrays, pools, and cloud accounts are held offline and data are inaccessible.

If the system has configured an encryption recovery key, this key can be manually supplied to the system to bring the encrypted storage back online. After entering the recovery key, monitor the system to confirm that the encrypted storage is now online and that encrypted data is accessible to host applications.

### Using the management GUI

The system does not currently support unlocking the system with the encryption recovery key by using the management GUI. Refer to the instructions described in "Using the command-line interface" section.

### Using the command-line interface

If the system is online, but encrypted storage is held offline, the recovery key can be supplied by running the following command on the configuration node:

```
chencryption -recoverykey validate
```

For more information see **chencryption** command.

In rare scenarios, the system is not able to come online because all quorum disks are offline. This can happen when all quorum disks are configured on Serially Attached SCSI (SAS) drives, and the SAS adapter is locked by a missing encryption key. In this scenario, nodes are held in service state (reporting node error 550) and cannot join the system. The recovery key can be supplied by running the following command on any node in service state:

```
satask -recoverykey
```

Enter the recovery key that uses the interactive prompt. A confirmation message is displayed when the recovery key has been entered correctly.

## *Disabling the encryption recovery key*

The encryption recovery key can be disabled using either the management GUI or the command-line interface.

**Note:** For security, encryption methods (including the recovery key) can only be disabled when physically connected to the technician port on the configuration node.

### Using the management GUI

When disabling encryption using the management GUI, encryption recovery key is automatically disabled in the process. See for instructions on disabling encryption using the management GUI.

To disable only the encryption recovery key, refer to the instructions described in "Using the command-line interface".

### Using the command-line interface

Follow these steps to disable the encryption recovery key:

1. Identify the configuration node of the system. For more information, see “Configuration node” on page 5.
2. Connect your computer to the technician port of the configuration node. For more information, see Technician port.
3. In a terminal window, use Secure Shell (SSH) software to connect to the cluster IP address of the system and authenticate using the credentials of any user with the SecurityAdmin role:

   ```
   ssh username@cluster_ip
   ```

   For more information, see **Connecting to the CLI with OpenSSH**.
4. To disable the encryption recovery key, enter the following command:

   ```
   chencryption -recoverykey disable
   ```

   For more information, see **chencryption** command.

5. The recovery key has been disabled successfully when the **recovery_key_name** field is blank and the **recovery_key_status** field is licensed. For more information, see **lsencryption** command.

## Decommissioning encryption

If encryption is no longer needed, you can disable the function on the system. This might be required when all encryption master keys for the system have been lost, when returning a loan system, or when repurposing a storage system for a different use. After decommissioning encryption on the system, all encrypted data and encryption configuration is removed.

During the disablement process, if the system is using key servers to manage keys, the master encryption keys remain on the key server but are not used again. The administrator of the key server is responsible for removing unused or expired keys from the key server. If the system is using USB flash drives to manage keys, the encryption key files remain on the USB flash drives but are not used again. The security administrator of the system is responsible for removing unused or expired encryption key files from the USB flash drives. If the system is using an encryption recovery key, any stored recovery keys are no longer used and can be discarded.

**Note:** All encrypted arrays, pools, and cloud accounts must be removed before disabling encryption.

You can use the management GUI to fully disable encryption on the system.

**Note:** For security, encryption methods (including the recovery key) can only be disabled when physically connected to the technician port on the configuration node.

### Before you begin

Follow these steps to connect to the technician port on the configuration node:

1. Identify the configuration node of the system. See Configuration node for more details.
2. Connect your computer to the technician port of the configuration node. For more information, see Technician port for more details.

### Using the management GUI

To fully disable encryption in the management GUI, complete the following steps:

1. In the management GUI, select **Settings** > **Security** > **Encryption**.
2. On the Encryption page, change the State to **Disabled**.

### Using the command-line interface

To fully disable encryption using the command-line interface, complete the following steps:

1. In a terminal window, use Secure Shell (SSH) software to connect to the configuration node and authenticate using the credentials of any user with the SecurityAdmin role:

   ```
   ssh username@192.168.0.1
   ```

   For more information, see **Connecting to the CLI with OpenSSH**.
2. To disable encryption on a system that uses an encryption recovery key, enter the following command:

   ```
   chencryption -recoverykey disable
   ```

   The recovery key has been disabled successfully. For more information, see **chencryption**.
3. To verify whether recovery key has been disabled successfully, run **lsencryption** command and check that the `recovery_key_name` field is blank and the `recovery_key_status` is licensed. For more information, see **lsencryption**.
4. To disable encryption on a system that uses key servers to manage encryption keys, enter the following command:

```
chencryption -keyserver disable
```

The recovery key that uses the key server has been disabled successfully. For more information, see **chencryption**.

5. To verify whether recovery key that uses the key server has been disabled successfully, run **lsencryption** command and check that the `keyserver_pmk_uid` field is blank and the `keyserver_status` is licensed. For more information, see **lsencryption**.
6. To disable encryption on a system that uses USB flash drives to manage encryption keys, enter the following command:

```
chencryption -usb disable
```

The recovery key that uses the key server has been disabled successfully. For more information, see **chencryption**.

7. To verify whether recovery key that uses the key server has been disabled successfully, run **lsencryption** command and check that the `usb_key_filename` field is blank and the `status` is licensed. For more information, see **lsencryption**.

## Password policy

With password policy support, system administrators can set security requirements that are related to password creation and expiration, timeout for inactivity, and actions after failed logon attempts.

Password policy support allows administrators to set security rules that are based on their organization's security guidelines and restrictions. The system supports the following password and security-related rules with this support.

**Password creation rules**

Administrator can set and manage the following rules for all passwords that are created on the system:

- Specify password length requirements for all users.
- Require passwords to use uppercase and lowercase characters.
- Require passwords to contain special characters.
- Prevent users from reusing recent passwords. This parameter is not supported on FlashSystem 5015, FlashSystem 5035 and FlashSystem 5045.
- Require users to change password on next login under any of these conditions:
  - Their password expired.
  - An administrator created new accounts with temporary passwords.

**Password expiration and rules for locking accounts**

The administrator can create the following rules for password expiration:

- Set password expiration limit.
- Set a password to expire immediately.
- Set number of failed login attempts before the account is locked.
- Set time for locked accounts.
- Automatic log out for inactivity.
- Locking superuser account access.

  **Note:** The superuser account is the default user that can complete installation, initial configuration, and other service-related actions on the system. If the superuser account is locked, service tasks cannot be completed.

# Setting up an SSH client

Secure Shell (SSH) is a client/server network application. It is used as a communication vehicle between the host system (for example, a laptop computer) and the system command-line interface (CLI).

## Overview

The system acts as the SSH server in this relationship. If you require command-line access without entering a password, it uses the principles of public and private keys for authentication.

## Authenticating SSH logins

Generate a Secure Shell (SSH) key pair to use the command-line interface (CLI). Additionally, when you use the SSH to log in to the system, you must use the RSA-based private key authentication.

When you are using AIX® hosts, SSH logins are authenticated on the system by using the RSA-based authentication that is supported in the OpenSSH client that is available for AIX . This scheme is based on the supplied password (or if you require command-line access without entering a password, then public-key cryptography is used) by using an algorithm that is known commonly as RSA.

**Note:** The authentication process for host systems that are not AIX is similar.

With this scheme (as in similar OpenSSH systems on other host types), the encryption, and decryption is done by using separate keys. This scheme means that it is not possible to derive the decryption key from the encryption key.

Because physical possession of the private key allows access to the system, the private key must be kept in a protected place, such as the .ssh directory on the AIX host, with restricted access permissions.

When SSH client (A) attempts to connect to SSH server (B), the SSH password (if you require command-line access without entering a password, the key pair) authenticates the connection. The key consists of two halves: the public keys and private keys. The SSH client public key is put onto SSH Server (B) using some means outside of the SSH session. When SSH client (A) tries to connect, the private key on SSH client (A) is able to authenticate with its public half on SSH server (B).

The system supports up to 32 interactive SSH sessions on the management IP address simultaneously.

**Note:** After an SSH interactive session times out, session gets automatically closed. Session timeout limit is set to 15 minutes, by default. The limit value can be changed by using the CLI command. See **`chsecurity`**.

To connect to the system, the SSH client requires a user login name and an SSH password (or if you require command-line access without entering a password, the key pair). Authenticate to the system by using a management username and password. When you use an SSH client to access a system, you must use your username and password. The system uses the password (and if not a password, the SSH key pair) to authorize the user who is accessing the system.

For multifactor authentication, IBM Security Verify communicates with the system and uses a PAM module to handle second factor authentication for SSH logins.

For Microsoft Windows hosts, PuTTY can be downloaded from the internet and used at no charge to provide an SSH client.

Microsoft Windows 10 includes the OpenSSH client. The **`ssh`** and **`scp`** commands work on the Microsoft Windows 10 command line. Use your existing private key or generate a new SSH keypair, and then place the private key in the `.ssh` folder in your user folder.

You can connect to the system by using the same username with which you log in to the system.

## Changing security protocol levels

Security administrators can change the security protocol level for either SSL or SSH protocols. When you change the security level for either of these security protocols, you can control which encryption algorithms, ciphers, and version of the protocol are permitted on the system.

Depending on your security requirements for your organization or geography, you can change the level for both SSL and SSH protocols.

The system supports OpenSSL and Java SSL ciphers to provide strong encryption for secure connections using the SSL or TLS protocols. On a new system, the default SSL protocol level is 5, and the default SSH protocol level is 3. If you want to allow the use of more cipher suites for compatibility with some applications, you can select a lower value. Selecting a higher value further restricts the list of supported cipher suites.

By default, the system uses the suggested SSL protocol and SSH protocol levels. If the suggested SSL and SSH protocol levels change on a future system upgrade, the system applies the new levels automatically. However, if you manually select a new level, then the system no longer uses the suggested levels and does not modify the level on future system upgrades. To use automatic suggestions, reset the SSL and SSH protocol levels using the management GUI or CLI.

**Note:**

The suggested SSL and SSH protocol levels might be increased in future code upgrades as security requirements change. To automatically update the protocol level to the new suggested level whenever you upgrade the system, select **Automatic** for the protocol level.

For servers or services that do not support TLS 1.3, do not set the security level to a level that supports only TLS 1.3, such as levels 6 and 7. Currently, the following servers or services do not support TLS 1.3:

- KeySecure key servers
- Duo Security for Multifactor Authentication on the Command Line Interface (CLI)
- Transparent Cloud Tiering (TCT)

If you use KeySecure key servers and want to use TLS 1.3 for secure communication, migrate to CipherTrust Manager key servers.

**Note:** For services or features involving two systems (system A and system B) that communicate with each other (for example, policy-based replication or secured IP partnerships), using mutual TLS authentication, you can configure different security levels on each system. Do not configure system A to only support TLS 1.2 (SSL protocol levels 2-4) and system B to support only TLS 1.3 (levels 6-7), or vice versa.

The following table describes each security level, minimum version of SSL or TLS allowed and the supported ciphers for each level:

| *Table 43. New supported Ciphers with certificate keytype* | | |
|---|---|---|
| **Security Level** | **Certificate Keytype** | **Supported Ciphers** |
| 2 | rsa2048, rsa4096 | ECDHE-RSA-AES256-GCM-SHA384<br>DHE-RSA-AES256-GCM-SHA384<br>ECDHE-RSA-CHACHA20-POLY1305<br>DHE-RSA-CHACHA20-POLY1305<br>DHE-RSA-AES256-CCM8<br>DHE-RSA-AES256-CCM<br>ECDHE-ARIA256-GCM-SHA384<br>DHE-RSA-ARIA256-GCM-SHA384<br>ECDHE-RSA-AES256-SHA384<br>DHE-RSA-AES256-SHA256<br>ECDHE-RSA-CAMELLIA256-SHA384<br>DHE-RSA-CAMELLIA256-SHA256<br>ECDHE-RSA-AES256-SHA<br>AES256-GCM-SHA384<br>AES256-CCM8<br>AES256-CCM<br>ARIA256-GCM-SHA384<br>AES256-SHA256<br>CAMELLIA256-SHA256<br>AES256-SHA<br>ECDHE-RSA-AES128-GCM-SHA256<br>DHE-RSA-AES128-GCM-SHA256<br>DHE-RSA-AES128-CCM8<br>DHE-RSA-AES128-CCM<br>ECDHE-ARIA128-GCM-SHA256<br>DHE-RSA-ARIA128-GCM-SHA256<br>ECDHE-RSA-AES128-SHA256<br>DHE-RSA-AES128-SHA256<br>ECDHE-RSA-CAMELLIA128-SHA256<br>DHE-RSA-CAMELLIA128-SHA256<br>ECDHE-RSA-AES128-SHA<br>AES128-GCM-SHA256<br>AES128-CCM8 |

| *Table 43. New supported Ciphers with certificate keytype (continued)* | | |
|---|---|---|
| **Security Level** | **Certificate Keytype** | **Supported Ciphers** |
| | ecdsa384, ecdsa521 | ECDHE-ECDSA-AES256-GCM-SHA384<br>ECDHE-ECDSA-CHACHA20-POLY1305<br>ECDHE-ECDSA-AES256-CCM8<br>ECDHE-ECDSA-AES256-CCM<br>ECDHE-ECDSA-ARIA256-GCM-SHA384<br>ECDHE-ECDSA-AES256-SHA384<br>ECDHE-ECDSA-CAMELLIA256-SHA384<br>ECDHE-ECDSA-AES256-SHA<br>ECDHE-ECDSA-AES128-GCM-SHA256<br>ECDHE-ECDSA-AES128-CCM8<br>ECDHE-ECDSA-AES128-CCM<br>ECDHE-ECDSA-ARIA128-GCM-SHA256<br>ECDHE-ECDSA-AES128-SHA256<br>ECDHE-ECDSA-CAMELLIA128-SHA256<br>ECDHE-ECDSA-AES28-SHA |

*Table 43. New supported Ciphers with certificate keytype (continued)*

| Security Level | Certificate Keytype | Supported Ciphers |
|---|---|---|
| 3 | rsa2048, rsa4096 | ECDHE-RSA-AES256-GCM-SHA384<br>DHE-RSA-AES256-GCM-SHA384<br>ECDHE-RSA-CHACHA20-POLY1305<br>DHE-RSA-CHACHA20-POLY1305<br>DHE-RSA-AES256-CCM8<br>DHE-RSA-AES256-CCM<br>ECDHE-ARIA256-GCM-SHA384<br>DHE-RSA-ARIA256-GCM-SHA384<br>ECDHE-RSA-AES256-SHA384<br>DHE-RSA-AES256-SHA256<br>ECDHE-RSA-CAMELLIA256-SHA384<br>DHE-RSA-CAMELLIA256-SHA256<br>AES256-GCM-SHA384<br>AES256-CCM8<br>AES256-CCM<br>ARIA256-GCM-SHA384<br>AES256-SHA256<br>CAMELLIA256-SHA256<br>ECDHE-RSA-AES128-GCM-SHA256<br>DHE-RSA-AES128-GCM-SHA256<br>DHE-RSA-AES128-CCM8<br>DHE-RSA-AES128-CCM<br>ECDHE-ARIA128-GCM-SHA256<br>DHE-RSA-ARIA128-GCM-SHA256<br>ECDHE-RSA-AES128-SHA256<br>DHE-RSA-AES128-SHA256<br>ECDHE-RSA-CAMELLIA128-SHA256<br>DHE-RSA-CAMELLIA128-SHA256<br>AES128-GCM-SHA256<br>AES128-CCM8<br>AES128-CCM<br>ARIA128-GCM-SHA256<br>AES128-SHA256 |

| *Table 43. New supported Ciphers with certificate keytype (continued)* | | |
|---|---|---|
| **Security Level** | **Certificate Keytype** | **Supported Ciphers** |
| | ecdsa384, ecdsa521 | ECDHE-ECDSA-AES256-GCM-SHA384<br>ECDHE-ECDSA-CHACHA20-POLY1305<br>ECDHE-ECDSA-AES256-CCM8<br>ECDHE-ECDSA-AES256-CCM<br>ECDHE-ECDSA-ARIA256-GCM-SHA384<br>ECDHE-ECDSA-AES256-SHA384<br>ECDHE-ECDSA-CAMELLIA256-SHA384<br>ECDHE-ECDSA-AES128-GCM-SHA256<br>ECDHE-ECDSA-AES128-CCM8<br>ECDHE-ECDSA-AES128-CCM<br>ECDHE-ECDSA-ARIA128-GCM-SHA256<br>ECDHE-ECDSA-AES128-SHA256<br>ECDHE-ECDSA-CAMELLIA128-SHA256 |
| 4 | ecdsa384, ecdsa521 | ECDHE-ECDSA-AES256-GCM-SHA384<br>ECDHE-ECDSA-AES128-GCM-SHA256 |

*Table 43. New supported Ciphers with certificate keytype (continued)*

| Security Level | Certificate Keytype | Supported Ciphers |
|---|---|---|
| 5 | rsa2048, rsa4096 | TLS_AES_256_GCM_SHA384<br>TLS_CHACHA20_POLY1305_SHA256<br>TLS_AES_128_GCM_SHA256<br>TLS_AES_128_CCM_8_SHA256<br>TLS_AES_128_CCM_SHA256<br>TLS_AES_128_CCM_8_SHA256<br>ECDHE-RSA-AES256-GCM-SHA384<br>ECDHE-RSA-CHACHA20-POLY1305<br>ECDHE-ARIA256-GCM-SHA384<br>ECDHE-RSA-AES256-SHA384<br>ECDHE-RSA-CAMELLIA256-SHA384<br>DHE-RSA-AES256-GCM-SHA384<br>DHE-RSA-CHACHA20-POLY1305<br>DHE-RSA-AES256-CCM8<br>DHE-RSA-AES256-CCM<br>DHE-RSA-ARIA256-GCM-SHA384<br>DHE-RSA-AES256-SHA256<br>DHE-RSA-CAMELLIA256-SHA256<br>ECDHE-RSA-AES128-GCM-SHA256<br>ECDHE-ARIA128-GCM-SHA256<br>ECDHE-RSA-AES128-SHA256<br>ECDHE-RSA-CAMELLIA128-SHA256<br>DHE-RSA-AES128-GCM-SHA256<br>DHE-RSA-AES128-CCM8<br>DHE-RSA-AES128-CCM<br>DHE-RSA-ARIA128-GCM-SHA256<br>DHE-RSA-AES128-SHA256<br>DHE-RSA-CAMELLIA128-SHA256 |

*Table 43. New supported Ciphers with certificate keytype (continued)*

| Security Level | Certificate Keytype | Supported Ciphers |
|---|---|---|
| | ecdsa384, ecdsa521 | TLS_AES_256_GCM_SHA384<br>TLS_CHACHA20_POLY1305_SHA256<br>TLS_AES_128_GCM_SHA256<br>TLS_AES_128_CCM_8_SHA256<br>TLS_AES_128_CCM_SHA256<br>TLS_AES_128_CCM_8_SHA256<br>ECDHE-ECDSA-AES256-GCM-SHA384<br>ECDHE-ECDSA-CHACHA20-POLY1305<br>ECDHE-ECDSA-AES256-CCM8<br>ECDHE-ECDSA-AES256-CCM<br>ECDHE-ECDSA-ARIA256-GCM-SHA384<br>ECDHE-ECDSA-AES256-SHA384<br>ECDHE-ECDSA-CAMELLIA256-SHA384<br>ECDHE-ECDSA-AES128-GCM-SHA256<br>ECDHE-ECDSA-AES128-CCM8<br>ECDHE-ECDSA-AES128-CCM<br>ECDHE-ECDSA-ARIA128-GCM-SHA256<br>ECDHE-ECDSA-AES128-SHA256<br>ECDHE-ECDSA-CAMELLIA128-SHA256 |
| 6 | rsa2048, ecdsa384, ecdsa521, rsa4096 | TLS_AES_256_GCM_SHA384<br>TLS_CHACHA20_POLY1305_SHA256<br>TLS_AES_128_GCM_SHA256<br>TLS_AES_128_CCM_8_SHA256<br>TLS_AES_128_CCM_SHA256 |
| 7 | rsa2048, ecdsa384, ecdsa521, rsa4096 | TLS_AES_256_GCM_SHA384 |

*Table 44. Supported SSL/TLS security levels*

| Security level | Description | Minimum security allowed | Supported Java SSL ciphers |
|---|---|---|---|
| 2 | Sets the system to disallow SSL version 3.0, TLS version 1.0, and TLS version 1.1. | TLS 1.2 | • SSL_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384<br>• SSL_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384<br>• SSL_ECDHE_RSA_WITH_AES_256_CBC_SHA384<br>• SSL_ECDHE_RSA_WITH_AES_256_GCM_SHA384<br>• SSL_RSA_WITH_AES_256_CBC_SHA256<br>• SSL_RSA_WITH_AES_256_GCM_SHA384<br>• SSL_ECDH_ECDSA_WITH_AES_256_CBC_SHA384<br>• SSL_ECDH_RSA_WITH_AES_256_GCM_SHA384<br>• SSL_DHE_RSA_WITH_AES_256_CBC_SHA256<br>• SSL_DHE_RSA_WITH_AES_256_CBC_SHA256<br>• SSL_DHE_RSA_WITH_AES_256_GCM_SHA384<br>• SSL_DHE_DSS_WITH_AES_256_CBC_SHA256<br>• SSL_DHE_DSS_WITH_AES_256_GCM_SHA384<br>• SSL_ECDHE_ECDSA_WITH_AES_256_CBC_SHA<br>• SSL_ECDHE_RSA_WITH_AES_256_CBC_SHA<br>• SSL_RSA_WITH_AES_256_CBC_SHA<br>• SSL_ECDH_ECDSA_WITH_AES_256_CBC_SHA<br>• SSL_ECDH_RSA_WITH_AES_256_CBC_SHA<br>• SSL_DHE_RSA_WITH_AES_256_CBC_SHA<br>• SSL_DHE_DSS_WITH_AES_256_CBC_SHA<br>• SSL_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256<br>• SSL_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256<br>• SSL_ECDHE_RSA_WITH_AES_128_CBC_SHA256<br>• SSL_ECDHE_RSA_WITH_AES_128_GCM_SHA256<br>• SSL_RSA_WITH_AES_128_CBC_SHA256<br>• SSL_RSA_WITH_AES_128_GCM_SHA256<br>• SSL_ECDH_ECDSA_WITH_AES_128_CBC_SHA256<br>• SSL_ECDH_ECDSA_WITH_AES_128_GCM_SHA256<br>• SSL_ECDH_RSA_WITH_AES_128_CBC_SHA256<br>• SSL_ECDH_RSA_WITH_AES_128_GCM_SHA256<br>• SSL_DHE_RSA_WITH_AES_128_CBC_SHA256<br>• SSL_DHE_RSA_WITH_AES_128_GCM_SHA256<br>• SSL_DHE_DSS_WITH_AES_128_CBC_SHA256<br>• SSL_DHE_DSS_WITH_AES_128_GCM_SHA256<br>• SSL_ECDHE_ECDSA_WITH_AES_128_CBC_SHA<br>• SSL_ECDHE_RSA_WITH_AES_128_CBC_SHA<br>• SSL_RSA_WITH_AES_128_CBC_SHA<br>• SSL_ECDH_ECDSA_WITH_AES_128_CBC_SHA<br>• SSL_ECDH_RSA_WITH_AES_128_CBC_SHA<br>• SSL_DHE_RSA_WITH_AES_128_CBC_SHA |

*Table 44. Supported SSL/TLS security levels (continued)*

| Security level | Description | Minimum security allowed | Supported Java SSL ciphers |
|---|---|---|---|
| 3 | Sets the system to disallow SSL version 3.0, TLS version 1.0, and TLS version 1.1 and to allow cipher suites that are exclusive to TLS version 1.2. | TLS 1.2 | • SSL_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384<br>• SSL_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384<br>• SSL_ECDHE_RSA_WITH_AES_256_CBC_SHA384<br>• SSL_ECDHE_RSA_WITH_AES_256_GCM_SHA384<br>• SSL_RSA_WITH_AES_256_CBC_SHA256<br>• SSL_RSA_WITH_AES_256_GCM_SHA384<br>• SSL_ECDH_ECDSA_WITH_AES_256_CBC_SHA384<br>• SSL_ECDH_RSA_WITH_AES_256_GCM_SHA384<br>• SSL_DHE_RSA_WITH_AES_256_CBC_SHA256<br>• SSL_DHE_RSA_WITH_AES_256_CBC_SHA256<br>• SSL_DHE_RSA_WITH_AES_256_GCM_SHA384<br>• SSL_DHE_DSS_WITH_AES_256_CBC_SHA256<br>• SSL_DHE_DSS_WITH_AES_256_GCM_SHA384<br>• SSL_ECDHE_ECDSA_WITH_AES_256_CBC_SHA<br>• SSL_ECDHE_RSA_WITH_AES_256_CBC_SHA<br>• SSL_RSA_WITH_AES_256_CBC_SHA<br>• SSL_ECDH_ECDSA_WITH_AES_256_CBC_SHA<br>• SSL_ECDH_RSA_WITH_AES_256_CBC_SHA<br>• SSL_DHE_RSA_WITH_AES_256_CBC_SHA<br>• SSL_DHE_DSS_WITH_AES_256_CBC_SHA<br>• SSL_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256<br>• SSL_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256<br>• SSL_ECDHE_RSA_WITH_AES_128_CBC_SHA256<br>• SSL_ECDHE_RSA_WITH_AES_128_GCM_SHA256<br>• SSL_RSA_WITH_AES_128_CBC_SHA256<br>• SSL_RSA_WITH_AES_128_GCM_SHA256<br>• SSL_ECDH_ECDSA_WITH_AES_128_CBC_SHA256<br>• SSL_ECDH_ECDSA_WITH_AES_128_GCM_SHA256<br>• SSL_ECDH_RSA_WITH_AES_128_CBC_SHA256<br>• SSL_ECDH_RSA_WITH_AES_128_GCM_SHA256<br>• SSL_DHE_RSA_WITH_AES_128_CBC_SHA256<br>• SSL_DHE_RSA_WITH_AES_128_GCM_SHA256<br>• SSL_DHE_DSS_WITH_AES_128_CBC_SHA256<br>• SSL_DHE_DSS_WITH_AES_128_GCM_SHA256 |

*Table 44. Supported SSL/TLS security levels (continued)*

| Security level | Description | Minimum security allowed | Supported Java SSL ciphers |
|---|---|---|---|
| 4 | Sets the system to disallow SSL version 3.0, TLS version 1.0, and TLS version 1.1, and to allow cipher suites that are exclusive to TLS version 1.2. Sets the system to disallow RSA ciphers and static key exchange ciphers. | TLS 1.2 | • SSL_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384<br>• SSL_DHE_DSS_WITH_AES_256_GCM_SHA384<br>• SSL_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256<br>• SSL_DHE_DSS_WITH_AES_128_GCM_SHA256 |

*Table 44. Supported SSL/TLS security levels (continued)*

| Security level | Description | Minimum security allowed | Supported Java SSL ciphers |
|---|---|---|---|
| 5 | Sets the system to disallow SSL version 3.0, TLS version 1.0, and TLS version 1.1 and to allow cipher suites that are exclusive to TLS version 1.2 and 1.3. | TLS 1.2, TLS 1.3 | **For TLS 1.3**<br>• TLS_AES_256_GCM_SHA384<br>• TLS_CHACHA20_POLY1305_SHA256<br>• TLS_AES_128_GCM_SHA256<br>• TLS_AES_128_CCM_8_SHA256<br>• TLS_AES_128_CCM_SHA256<br>**For TLS 1.2**<br>The security level 5 supports all the Java SSL ciphers supported at the security level 3. A few additional TLS 1.2 specific Java SSL ciphers supported at level 5 are:<br>• TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384<br>• TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256<br>• TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384<br>• TLS_RSA_WITH_AES_256_GCM_SHA384<br>• TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384<br>• TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384<br>• TLS_DHE_DSS_WITH_AES_256_GCM_SHA384<br>• TLS_DHE_RSA_WITH_AES_256_GCM_SHA384<br>• TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256<br>• TLS_RSA_WITH_AES_128_GCM_SHA256<br>• TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256<br>• TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256<br>• TLS_DHE_RSA_WITH_AES_128_GCM_SHA256<br>• TLS_DHE_DSS_WITH_AES_128_GCM_SHA256 |
| 6 | Sets the system to disallow SSL version 3.0, TLS version 1.0, TLS version 1.1, and TLS version 1.2 and to allow cipher suites that are exclusive to TLS version 1.3. | TLS 1.3 | • TLS_AES_256_GCM_SHA384<br>• TLS_CHACHA20_POLY1305_SHA256<br>• TLS_AES_128_GCM_SHA256<br>• TLS_AES_128_CCM_8_SHA256<br>• TLS_AES_128_CCM_SHA256 |

| Table 44. Supported SSL/TLS security levels (continued) | | | |
|---|---|---|---|
| **Security level** | **Description** | **Minimum security allowed** | **Supported Java SSL ciphers** |
| 7 | Sets the system to disallow SSL version 3.0, TLS version 1.0, TLS version 1.1, and TLS version 1.2 and to allow the TLS 1.3 cipher suites that are FIPS mode compliant. | TLS 1.3 | TLS_AES_256_GCM_SHA384 |

The following table describes the SSH security levels supported by the system:

| Table 45. SSH algorithms supported at each security level | | | | |
|---|---|---|---|---|
| **Security level** | **Key Exchange** | **Cipher Suite** | **MAC Algorithm** | **Host Key Algorithms** |
| 1 | • curve25519-sha256<br>• curve25519-sha256@libssh.org<br>• ecdh-sha2-nistp256<br>• ecdh-sha2-nistp384<br>• ecdh-sha2-nistp521<br>• diffie-hellman-group-exchange-sha256<br>• diffie-hellman-group16-sha512<br>• diffie-hellman-group18-sha512<br>• diffie-hellman-group14-sha256<br>• diffie-hellman-group14-sha1<br>• diffie-hellman-group1-sha1<br>• diffie-hellman-group-exchange-sha1 | aes256-ctr<br>aes192-ctr<br>aes128-ctr<br>chacha20-poly1305@openssh.com<br>aes256-gcm@openssh.com<br>aes128-gcm@openssh.com<br>aes256-cbc<br>aes192-cbc<br>aes128-cbc | hmac-sha2-256<br>hmac-sha2-512<br>hmac-sha2-256-etm@openssh.com<br>hmac-sha2-512-etm@openssh.com<br>hmac-sha1 | rsa-sha2-256<br>rsa-sha2-512<br>ecdsa-sha2-nistp521<br>ecdsa-sha2-nistp521-cert-v01@openssh.com<br>ssh-rsa<br>ssh-rsa-cert-v01@openssh.com |

*Table 45. SSH algorithms supported at each security level (continued)*

| Security level | Key Exchange | Cipher Suite | MAC Algorithm | Host Key Algorithms |
|---|---|---|---|---|
| 2 | • curve25519-sha256<br>• curve25519-sha256@libssh.org<br>• ecdh-sha2-nistp256<br>• ecdh-sha2-nistp384<br>• ecdh-sha2-nistp521<br>• diffie-hellman-group-exchange-sha256<br>• diffie-hellman-group16-sha512<br>• diffie-hellman-group18-sha512<br>• diffie-hellman-group14-sha256<br>• diffie-hellman-group14-sha1 | aes256-ctr<br>aes192-ctr<br>aes128-ctr<br>chacha20-poly1305@openssh.com<br>aes256-gcm@openssh.com<br>aes128-gcm@openssh.com | hmac-sha2-256<br>hmac-sha2-512<br>hmac-sha2-256-etm@openssh.com<br>hmac-sha2-512-etm@openssh.com<br>hmac-sha1 | rsa-sha2-256<br>rsa-sha2-512<br>ecdsa-sha2-nistp521<br>ecdsa-sha2-nistp521-cert-v01@openssh.com<br>ssh-rsa<br>ssh-rsa-cert-v01@openssh.com |
| 3 | • curve25519-sha256<br>• curve25519-sha256@libssh.org<br>• ecdh-sha2-nistp256<br>• ecdh-sha2-nistp384<br>• ecdh-sha2-nistp521<br>• diffie-hellman-group-exchange-sha256<br>• diffie-hellman-group16-sha512<br>• diffie-hellman-group18-sha512<br>• diffie-hellman-group14-sha256 | aes256-ctr<br>aes192-ctr<br>aes128-ctr<br>chacha20-poly1305@openssh.com<br>aes256-gcm@openssh.com<br>aes128-gcm@openssh.com | hmac-sha2-256<br>hmac-sha2-512<br>hmac-sha2-256-etm@openssh.com<br>hmac-sha2-512-etm@openssh.com | rsa-sha2-256<br>rsa-sha2-512<br>ecdsa-sha2-nistp521<br>ecdsa-sha2-nistp521-cert-v01@openssh.com<br>ssh-rsa<br>ssh-rsa-cert-v01@openssh.com |

| Table 45. SSH algorithms supported at each security level (continued) | | | | |
|---|---|---|---|---|
| **Security level** | **Key Exchange** | **Cipher Suite** | **MAC Algorithm** | **Host Key Algorithms** |
| 4 | • curve25519-sha256<br>• curve25519-sha256@libssh.org<br>• ecdh-sha2-nistp256<br>• ecdh-sha2-nistp384<br>• ecdh-sha2-nistp521 | aes256-ctr<br>aes192-ctr<br>aes128-ctr<br>aes256-gcm@openssh.com<br>aes128-gcm@openssh.com | hmac-sha2-256<br>hmac-sha2-512<br>hmac-sha2-256-etm@openssh.com<br>hmac-sha2-512-etm@openssh.com | rsa-sha2-256<br>rsa-sha2-512<br>ecdsa-sha2-nistp521<br>ecdsa-sha2-nistp521-cert-v01@openssh.com |

**Restriction:** The 3-site-orchestrator does not support SSH protocol level 4.

When you change the SSL security protocol level, you must restart any service using SSL/TLS. All current session are ended to ensure no sessions are open using the old security level. It can take a few minutes for the service to be available.

## Using the management GUI

You can use the management GUI to update protocol levels for SSL and SSH connections:

**SSL/TLS security protocol level**

By default, the SSL protocol level is set to 5, and the SSH protocol level is set to 3 . To change the SSL or SSH security protocol levels, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Security protocol levels**.
2. You can update any of the following details:

**SSL protocol level**

**Note:** Changing the **SSL protocol level** causes the GUI to restart.

SSL ensures that the data is securely transferred. By default, security level 5 is set to allow both TLS 1.2 and TLS 1.3. You can select the required **SSL protocol level** from the following options:

- Automatic - Use suggested level of 5.
- 2 - TLS 1.2, allow TLS 1.0, 1.1, and 1.2 ciphers.
- 3 - TLS 1.2, allow TLS 1.2 ciphers.
- 4 - TLS 1.2, allow TLS 1.2 ciphers but disallow RSA and static key exchange ciphers.
- 5 - TLS 1.2 and TLS 1.3, disallow static key exchange ciphers.
- 6 - TLS 1.3, allow only TLS 1.3 ciphers.
- 7 - TLS 1.3, allow only ciphers that support FIPS.

**SSH protocol level**

Select the SSH protocol level that is used for connections to the command-line interface. Each level supports different algorithms for key exchange. The range is 1 - 4, where 3 is the default value. Select the required **SSH protocol level** from the following options:

- Automatic - Use suggested level of 3.
- 1 - Allow block ciphers.

- 2 - Disallow block ciphers.
- 3 - Disallow SHA1.
- 4 - Disallow Diffie-Hellman.

3. Click **Save**.

**Note:** The suggested SSL and SSH protocol levels might be increased in future code upgrades as security requirements change. To automatically update the protocol level to the new suggested level whenever you upgrade the system, select **Automatic** for the protocol level.

**SSH rules**

To update the SSH rules settings, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **SSH Rules**.
2. You can update any of the following details:

   **SSH login grace period (seconds)**
   Indicates the amount of time in seconds to log in before SSH times out. The range is 15 - 1800.

   **Maximum login attempts (SSH)**
   Indicates the total number of login attempts allowed per single SSH connection. The range is 1 - 10.

3. Click **Save**.

## Using the command-line interface (CLI)

The `chsecurity` command allows you to set the ciphers and protocols that are allowed by secure interfaces to reduce the vulnerability to attack. However, changing the security level might break the connection to external systems such as web browsers and anything that is connected through CIM such as VMWare provisioning utilities or IBM Spectrum Control software.

1. To display your current system SSL, TLS, and SSH security settings, enter the following command:

```
lssecurity
```

The results show the current setting as shown in the following example:

```
sslprotocol 5
sshprotocol 3
gui_timeout_mins 30
cli_timeout_mins 15
restapi_timeout_mins 60
min_password_length 8
password_special_chars 0
password_upper_case 0
password_lower_case 0
password_digits 0
check_password_history no
max_password_history 6
min_password_age_days 1
password_expiry_days 0
expiry_warning_days 14
superuser_locking disabled
max_failed_login_attempts 0
lockout_period_mins 10
superuser_multi_factor no
ssh_grace_time_seconds 60
ssh_max_tries 6
superuser_password_sshkey_required no
superuser_gui_disabled no
superuser_rest_disabled no
superuser_cim_disabled yes
two_person_integrity_enabled no
two_person_integrity_superuser_locked no
ssl_protocols_enabled TLSv1.2:TLSv1.3
ssl_protocol_suggested yes
ssh_protocol_suggested yes
```

2. To change SSL/TLS settings, enter **chsecurity -sslprotocol security_level**, where **security_level** is 2, 3, 4, 5, 6, or 7.

   **Note:** You might lose the connection to the management GUI when the security level is changed. If you lose the connection, use the CLI to decrease the security level to a lower setting.
3. To change SSH settings, enter **chsecurity -sshprotocol security_level**, where **security_level** is 1, 2, 3, or 4.

# Security levels and supported security ciphers

You can use connections based on the Transport Layer Security (TLS), which is the successor of the Secure Sockets Layer (SSL) protocol, to ensure safer communications.

## Version

This information about security settings applies to the current release only.

## SSL certificates

The system uses a certificate to authenticate SSL connections. For more information about managing certificates, see Creating and managing certificate authority store by using CLI .

## TLS or SSL connections and security levels

**Note:** The terms TLS and SSL are often used interchangeably in the industry.

The system uses TLS or SSL connections to control access to interfaces such as the management GUI, the service assistant GUI, the key server, and RESTful API. TLS or SSL connections use security ciphers to help control access.

You can use security ciphers that are supported by different levels of TLS or SSL. Each level supports ciphers that provide differing strengths of encryption. You can set the security level to 5, 6, or 7 to be compliant with the NIST 800-52 standard. Security level 7 allows only the cipher suite TLS_AES_256_GCM_SHA384, which the NIST recommends for Federal Information Processing Standards (FIPS) mode.

SSL protocol levels 2 to 4 do not support the TLS 1.3 protocol and the cipher suites that are approved by the NIST 800-52 standard. Currently, the TLS or SSL security level 7 is the maximum level that is supported and TLS or SSL security level 2 is the lowest security level supported.

## SSL protocols and ciphers supported at each security level

Table 46 on page 171 displays the protocols that are supported at security level 7.

*Table 46. Protocols supported at level 7*

| **Protocol level** | **Is it supported?** |
|---|---|
| TLS 1.3 | Yes |
| TLS 1.2 | No |
| TLS 1.1 | No |
| TLS 1.0 | No |
| SSL 3 and earlier | No |

Table 47 on page 172 displays Java™ SSL ciphers that are supported at security level 7.

| *Table 47. Java SSL ciphers supported at security level 7* |
|---|
| **Java SSL ciphers** |
| TLS_AES_256_GCM_SHA384 |

Table 48 on page 172 displays OpenSSL security ciphers that are supported by security level 7.

| *Table 48. OpenSSL ciphers supported at level 7 (`chsecurity -sslprotocol 7`)* | | |
|---|---|---|
| **Cipher** | **Bulk encryption algorithm** | **Hashing algorithm** |
| AES-256-GCM-SHA384 | AES-256-GCM | SHA384 |

Table 49 on page 172 displays the protocols that are supported at security level 6.

| *Table 49. Protocols supported at level 6* | |
|---|---|
| **Protocol level** | **Is it supported?** |
| TLS 1.3 | Yes |
| TLS 1.2 | No |
| TLS 1.1 | No |
| TLS 1.0 | No |
| SSL 3 and earlier | No |

Table 50 on page 172 displays Java SSL ciphers that are supported at security level 6.

| *Table 50. Java SSL ciphers supported at security level 6* |
|---|
| **Java SSL ciphers** |
| TLS_AES_256_GCM_SHA384 |
| TLS_CHACHA20_POLY1305_SHA256 |
| TLS_AES_128_GCM_SHA256 |
| TLS_AES_128_CCM_8_SHA256 |
| TLS_AES_128_CCM_SHA256 |

Table 51 on page 172 displays OpenSSL security ciphers that are supported by security level 6.

| *Table 51. OpenSSL ciphers supported at level 6 (`chsecurity -sslprotocol 6`)* | | |
|---|---|---|
| **Cipher** | **Bulk encryption algorithm** | **Hashing algorithm** |
| AES-256-GCM-SHA384 | AES-256-GCM | SHA384 |
| CHACHA20-POLY1305-SHA256 | CHACHA20-POLY1305 | SHA256 |
| AES-128-GCM-SHA256 | AES-128-GCM | SHA256 |
| AES-128-CCM-8-SHA256 | AES-128-CCM-8 | SHA256 |
| AES-128-CCM-SHA256 | AES-128-CCM | SHA256 |

Table 52 on page 173 displays the protocols that are supported at security level 5.

| *Table 52. Protocols supported at level 5* | |
|---|---|
| **Protocol level** | **Is it supported?** |
| TLS 1.3 | Yes |
| TLS 1.2 | Yes |
| TLS 1.1 | No |
| TLS 1.0 | No |
| SSL 3 and earlier | No |

Table 53 on page 173 displays Java SSL ciphers that are supported at security level 5.

| *Table 53. Java SSL ciphers supported at security level 5* |
|---|
| **Java SSL ciphers** |
| **For TLS 1.3** |
| TLS_AES_256_GCM_SHA384 |
| TLS_CHACHA20_POLY1305_SHA256 |
| TLS_AES_128_GCM_SHA256 |
| TLS_AES_128_CCM_8_SHA256 |
| TLS_AES_128_CCM_SHA256 |
| **For TLS 1.2**<br>The security level 5 supports all the Java SSL ciphers that are supported at the security level 3. A few more TLS 1.2 specific Java SSL ciphers that are supported at level 5 are: |
| TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 |
| TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 |
| TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| TLS_RSA_WITH_AES_256_GCM_SHA384 |
| TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384 |
| TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384 |
| TLS_DHE_DSS_WITH_AES_256_GCM_SHA384 |
| TLS_DHE_RSA_WITH_AES_256_GCM_SHA384 |
| TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 |
| TLS_RSA_WITH_AES_128_GCM_SHA256 |
| TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256 |
| TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256 |
| TLS_DHE_RSA_WITH_AES_128_GCM_SHA256 |
| TLS_DHE_DSS_WITH_AES_128_GCM_SHA256 |

Table 54 on page 174 and Table 55 on page 174 displays OpenSSL security ciphers that are supported by security level 5.

| *Table 54. OpenSSL ciphers supported at level 5 for TLS 1.3 (`chsecurity -sslprotocol 5`)* | | |
|---|---|---|
| **Cipher** | **Bulk encryption algorithm** | **Hashing algorithm** |
| AES-256-GCM-SHA384 | AES-256-GCM | SHA384 |
| CHACHA20-POLY1305-SHA256 | CHACHA20-POLY1305 | SHA256 |
| AES-128-GCM-SHA256 | AES-128-GCM | SHA256 |
| AES-128-CCM-8-SHA256 | AES-128-CCM-8 | SHA256 |
| AES-128-CCM-SHA256 | AES-128-CCM | SHA256 |

| *Table 55. OpenSSL ciphers supported at level 5 for TLS 1.2 (`chsecurity -sslprotocol 5`)* | | | | |
|---|---|---|---|---|
| **Cipher** | **Kx** | **Au** | **Enc** | **Mac** |
| ECDHE-RSA-AES256-GCM-SHA384 | ECDH | RSA | AESGCM(256) | AEAD |
| ECDHE-ECDSA-AES256-GCM-SHA384 | ECDH | ECDSA | AESGCM(256) | AEAD |
| ECDHE-RSA-AES256-SHA384 | ECDH | RSA | AES(256) | SHA384 |
| ECDHE-ECDSA-AES256-SHA384 | ECDH | ECDSA | AES(256) | SHA384 |
| DHE-RSA-AES256-GCM-SHA384 | DH | RSA | AESGCM(256) | AEAD |
| DHE-RSA-AES256-SHA256 | DH | RSA | AES(256) | SHA256 |
| AES256-GCM-SHA384 | RSA | RSA | AESGCM(256) | AEAD |
| AES256-SHA256 | RSA | RSA | AES(256) | SHA256 |
| ECDHE-RSA-AES128-GCM-SHA256 | ECDH | RSA | AESGCM(128) | AEAD |
| ECDHE-ECDSA-AES128-GCM-SHA256 | ECDH | ECDSA | AESGCM(128) | AEAD |
| ECDHE-RSA-AES128-SHA256 | ECDH | RSA | AES(128) | SHA256 |
| ECDHE-ECDSA-AES128-SHA256 | ECDH | ECDSA | AES(128) | SHA256 |
| DHE-RSA-AES128-GCM-SHA256 | DH | RSA | AESGCM(128) | AEAD |
| DHE-RSA-AES128-SHA256 | DH | RSA | AES(128) | SHA256 |
| AES128-GCM-SHA256 | RSA | RSA | AESGCM(128) | AEAD |

| *Table 55. OpenSSL ciphers supported at level 5 for TLS 1.2 (`chsecurity -sslprotocol 5`) (continued)* | | | | |
|---|---|---|---|---|
| **Cipher** | **Kx** | **Au** | **Enc** | **Mac** |
| AES128-SHA256 | RSA | RSA | AES(128) | SHA256 |

Table 56 on page 175 displays the protocols that are supported at security level 4.

| *Table 56. Protocols supported at level 4* | |
|---|---|
| **Protocol level** | **Is it supported?** |
| TLS 1.2 | Yes |
| TLS 1.1 | No |
| TLS 1.0 | No |
| SSL 3 and earlier | No |

Table 57 on page 175 displays Java SSL ciphers that are supported at security level 4.

| *Table 57. Java SSL ciphers supported at security level 4* |
|---|
| **Java SSL ciphers** |
| SSL_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 |
| SSL_DHE_DSS_WITH_AES_256_GCM_SHA384 |
| SSL_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 |
| SSL_DHE_DSS_WITH_AES_128_GCM_SHA256 |

Table 58 on page 175 displays OpenSSL security ciphers that are supported by security level 4.

| *Table 58. OpenSSL ciphers supported at level 4 (`chsecurity -sslprotocol 4`)* | | | | |
|---|---|---|---|---|
| **Cipher** | **Kx** | **Au** | **Enc** | **Mac** |
| ECDHE-ECDSA-AES256-GCM-SHA384 | ECDH | ECDSA | AESGCM(256) | AEAD |
| DHE-DSS-AES256-GCM-SHA384 | DH | DSS | AESGCM(256) | AEAD |
| ECDHE-ECDSA-AES128-GCM-SHA256 | ECDH | ECDSA | AESGCM(128) | AEAD |
| DHE-DSS-AES128-GCM-SHA256 | DH | DSS | AESGCM(128) | AEAD |

Table 59 on page 175 displays the protocols that are supported at security level 3.

| *Table 59. Protocols supported at level 3* | |
|---|---|
| **Protocol level** | **Is it supported?** |
| TLS 1.2 | Yes |
| TLS 1.1 | No |
| TLS 1.0 | No |
| SSL 3 and earlier | No |

Table 60 on page 176 displays Java SSL ciphers that are supported at security level 3.

*Table 60. Java SSL ciphers supported at security level 3*

| **Java SSL ciphers** |
|---|
| SSL_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384 |
| SSL_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 |
| SSL_ECDHE_RSA_WITH_AES_256_CBC_SHA384 |
| SSL_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| SSL_RSA_WITH_AES_256_CBC_SHA256 |
| SSL_RSA_WITH_AES_256_GCM_SHA384 |
| SSL_ECDH_ECDSA_WITH_AES_256_CBC_SHA384 |
| SSL_ECDH_RSA_WITH_AES_256_GCM_SHA384 |
| SSL_DHE_RSA_WITH_AES_256_CBC_SHA256 |
| SSL_DHE_RSA_WITH_AES_256_CBC_SHA256 |
| SSL_DHE_RSA_WITH_AES_256_GCM_SHA384 |
| SSL_DHE_DSS_WITH_AES_256_CBC_SHA256 |
| SSL_DHE_DSS_WITH_AES_256_GCM_SHA384 |
| SSL_ECDHE_ECDSA_WITH_AES_256_CBC_SHA |
| SSL_ECDHE_RSA_WITH_AES_256_CBC_SHA |
| SSL_RSA_WITH_AES_256_CBC_SHA |
| SSL_ECDH_ECDSA_WITH_AES_256_CBC_SHA |
| SSL_ECDH_RSA_WITH_AES_256_CBC_SHA |
| SSL_DHE_RSA_WITH_AES_256_CBC_SHA |
| SSL_DHE_DSS_WITH_AES_256_CBC_SHA |
| SSL_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256 |
| SSL_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 |
| SSL_ECDHE_RSA_WITH_AES_128_CBC_SHA256 |
| SSL_ECDHE_RSA_WITH_AES_128_GCM_SHA256 |
| SSL_RSA_WITH_AES_128_CBC_SHA256 |
| SSL_RSA_WITH_AES_128_GCM_SHA256 |
| SSL_ECDH_ECDSA_WITH_AES_128_CBC_SHA256 |
| SSL_ECDH_ECDSA_WITH_AES_128_GCM_SHA256 |
| SSL_ECDH_RSA_WITH_AES_128_CBC_SHA256 |
| SSL_ECDH_RSA_WITH_AES_128_GCM_SHA256 |
| SSL_DHE_RSA_WITH_AES_128_CBC_SHA256 |
| SSL_DHE_RSA_WITH_AES_128_GCM_SHA256 |
| SSL_DHE_DSS_WITH_AES_128_CBC_SHA256 |

| *Table 60. Java SSL ciphers supported at security level 3 (continued)* |
|---|
| **Java SSL ciphers** |
| SSL_DHE_DSS_WITH_AES_128_GCM_SHA256 |

Table 61 on page 177 displays OpenSSL security ciphers that are supported by security level 3.

*Table 61. OpenSSL ciphers supported at level 3 (`chsecurity -sslprotocol 3`)*

| Cipher | Kx | Au | Enc | Mac |
|---|---|---|---|---|
| ECDHE-RSA-AES256-GCM-SHA384 | ECDH | RSA | AESGCM(256) | AEAD |
| ECDHE-ECDSA-AES256-GCM-SHA384 | ECDH | ECDSA | AESGCM(256) | AEAD |
| ECDHE-RSA-AES256-SHA384 | ECDH | RSA | AES(256) | SHA384 |
| ECDHE-ECDSA-AES256-SHA384 | ECDH | ECDSA | AES(256) | SHA384 |
| DHE-DSS-AES256-GCM-SHA384 | DH | DSS | AESGCM(256) | AEAD |
| DHE-RSA-AES256-GCM-SHA384 | DH | RSA | AESGCM(256) | AEAD |
| DHE-RSA-AES256-SHA256 | DH | RSA | AES(256) | SHA256 |
| ECDH-RSA-AES256-GCM-SHA384 E | ECDH/RSA | ECDH | AESGCM(256) | AEAD |
| ECDH-ECDSA-AES256-GCM-SHA384 | ECDH/ECDSA | ECDH | AESGCM(256) | AEAD |
| ECDH-RSA-AES256-SHA384 | ECDH/RSA | ECDH | AES(256) | SHA384 |
| ECDH-ECDSA-AES256-SHA384 | ECDH/ECDSA | ECDH | AES(256) | SHA384 |
| AES256-GCM-SHA384 | RSA | RSA | AESGCM(256) | AEAD |
| AES256-SHA256 | RSA | RSA | AES(256) | SHA256 |
| ECDHE-RSA-AES128-GCM-SHA256 | ECDH | RSA | AESGCM(128) | AEAD |
| ECDHE-ECDSA-AES128-GCM-SHA256 | ECDH | ECDSA | AESGCM(128) | AEAD |
| ECDHE-RSA-AES128-SHA256 | ECDH | RSA | AES(128) | SHA256 |

| *Table 61. OpenSSL ciphers supported at level 3 (`chsecurity -sslprotocol 3`) (continued)* | | | | |
|---|---|---|---|---|
| **Cipher** | **Kx** | **Au** | **Enc** | **Mac** |
| ECDHE-ECDSA-AES128-SHA256 | ECDH | ECDSA | AES(128) | SHA256 |
| DHE-DSS-AES128-GCM-SHA256 | DH | DSS | AESGCM(128) | AEAD |
| DHE-RSA-AES128-GCM-SHA256 | DH | RSA | AESGCM(128) | AEAD |
| DHE-RSA-AES128-SHA256 | DH | RSA | AES(128) | SHA256 |
| DHE-DSS-AES128-SHA256 | DH | DSS | AES(128) | SHA256 |
| ECDH-RSA-AES128-GCM-SHA256 | ECDH/RSA | ECDH | AESGCM(128) | AEAD |
| ECDH-ECDSA-AES128-GCM-SHA256 | ECDH/ECDSA | ECDH | AESGCM(128) | AEAD |
| ECDH-RSA-AES128-SHA256 | ECDH/RSA | ECDH | AES(128) | SHA256 |
| ECDH-ECDSA-AES128-SHA256 | ECDH/ECDSA | ECDH | AES(128) | SHA256 |
| AES128-GCM-SHA256 | RSA | RSA | AESGCM(128) | AEAD |
| AES128-SHA256 | RSA | RSA | AES(128) | SHA256 |

Table 62 on page 178 displays the protocols that are supported at security level 2.

| *Table 62. Protocols supported at level 2* | |
|---|---|
| **Protocol level** | **Is it supported?** |
| TLS 1.2 | Yes |
| TLS 1.1 | No |
| TLS 1.0 | No |
| SSL 3 and earlier | No |

Table 63 on page 178 displays Java SSL ciphers that are supported at security level 2.

| *Table 63. Java SSL ciphers supported at level 2* |
|---|
| SSL_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384 |
| SSL_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 |
| SSL_ECDHE_RSA_WITH_AES_256_CBC_SHA384 |
| SSL_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| SSL_RSA_WITH_AES_256_CBC_SHA256 |
| SSL_RSA_WITH_AES_256_GCM_SHA384 |

| *Table 63. Java SSL ciphers supported at level 2 (continued)* |
| --- |
| SSL_ECDH_ECDSA_WITH_AES_256_CBC_SHA384 |
| SSL_ECDH_RSA_WITH_AES_256_GCM_SHA384 |
| SSL_DHE_RSA_WITH_AES_256_CBC_SHA256 |
| SSL_DHE_RSA_WITH_AES_256_CBC_SHA256 |
| SSL_DHE_RSA_WITH_AES_256_GCM_SHA384 |
| SSL_DHE_DSS_WITH_AES_256_CBC_SHA256 |
| SSL_DHE_DSS_WITH_AES_256_GCM_SHA384 |
| SSL_ECDHE_ECDSA_WITH_AES_256_CBC_SHA |
| SSL_ECDHE_RSA_WITH_AES_256_CBC_SHA |
| SSL_RSA_WITH_AES_256_CBC_SHA |
| SSL_ECDH_ECDSA_WITH_AES_256_CBC_SHA |
| SSL_ECDH_RSA_WITH_AES_256_CBC_SHA |
| SSL_DHE_RSA_WITH_AES_256_CBC_SHA |
| SSL_DHE_DSS_WITH_AES_256_CBC_SHA |
| SSL_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256 |
| SSL_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 |
| SSL_ECDHE_RSA_WITH_AES_128_CBC_SHA256 |
| SSL_ECDHE_RSA_WITH_AES_128_GCM_SHA256 |
| SSL_RSA_WITH_AES_128_CBC_SHA256 |
| SSL_RSA_WITH_AES_128_GCM_SHA256 |
| SSL_ECDH_ECDSA_WITH_AES_128_CBC_SHA256 |
| SSL_ECDH_ECDSA_WITH_AES_128_GCM_SHA256 |
| SSL_ECDH_RSA_WITH_AES_128_CBC_SHA256 |
| SSL_ECDH_RSA_WITH_AES_128_GCM_SHA256 |
| SSL_DHE_RSA_WITH_AES_128_CBC_SHA256 |
| SSL_DHE_RSA_WITH_AES_128_GCM_SHA256 |
| SSL_DHE_DSS_WITH_AES_128_CBC_SHA256 |
| SSL_DHE_DSS_WITH_AES_128_GCM_SHA256 |
| SSL_ECDHE_ECDSA_WITH_AES_128_CBC_SHA |
| SSL_ECDHE_RSA_WITH_AES_128_CBC_SHA |
| SSL_RSA_WITH_AES_128_CBC_SHA |
| SSL_ECDH_ECDSA_WITH_AES_128_CBC_SHA |
| SSL_ECDH_RSA_WITH_AES_128_CBC_SHA |
| SSL_DHE_RSA_WITH_AES_128_CBC_SHA |
| SSL_DHE_DSS_WITH_AES_128_CBC_SHA |

Table 64 on page 180 displays OpenSSL security ciphers that are supported by security level 2.

*Table 64. OpenSSL ciphers supported at level 2 (`chsecurity -sslprotocol 2`)*

| Cipher | Kx | Au | Enc | Mac |
|---|---|---|---|---|
| ECDHE-RSA-AES256-GCM-SHA384 | ECDH | RSA | AESGCM(256) | AEAD |
| ECDHE-ECDSA-AES256-GCM-SHA384 | ECDH | ECDSA | AESGCM(256) | AEAD |
| ECDHE-RSA-AES256-SHA384 | ECDH | RSA | AES(256) | SHA384 |
| ECDHE-ECDSA-AES256-SHA384 | ECDH | ECDSA | AES(256) | SHA384 |
| DHE-DSS-AES256-GCM-SHA384 | DH | DSS | AESGCM(256) | AEAD |
| DHE-RSA-AES256-GCM-SHA384 | DH | RSA | AESGCM(256) | AEAD |
| DHE-RSA-AES256-SHA256 | DH | RSA | AES(256) | SHA256 |
| ECDH-RSA-AES256-GCM-SHA384 E | ECDH/RSA | ECDH | AESGCM(256) | AEAD |
| ECDH-ECDSA-AES256-GCM-SHA384 | ECDH/ECDSA | ECDH | AESGCM(256) | AEAD |
| ECDH-RSA-AES256-SHA384 | ECDH/RSA | ECDH | AES(256) | SHA384 |
| ECDH-ECDSA-AES256-SHA384 | ECDH/ECDSA | ECDH | AES(256) | SHA384 |
| AES256-GCM-SHA384 | RSA | RSA | AESGCM(256) | AEAD |
| AES256-SHA256 | RSA | RSA | AES(256) | SHA256 |
| AES256-SHA | RSA | RSA | AES(256) | SHA1 |
| ECDHE-RSA-AES128-GCM-SHA256 | ECDH | RSA | AESGCM(128) | AEAD |
| ECDHE-ECDSA-AES128-GCM-SHA256 | ECDH | ECDSA | AESGCM(128) | AEAD |
| ECDHE-RSA-AES128-SHA256 | ECDH | RSA | AES(128) | SHA256 |
| ECDHE-ECDSA-AES128-SHA256 | ECDH | ECDSA | AES(128) | SHA256 |
| DHE-DSS-AES128-GCM-SHA256 | DH | DSS | AESGCM(128) | AEAD |

| *Table 64. OpenSSL ciphers supported at level 2 (`chsecurity -sslprotocol 2`) (continued)* | | | | |
|---|---|---|---|---|
| **Cipher** | **Kx** | **Au** | **Enc** | **Mac** |
| DHE-RSA-AES128-GCM-SHA256 | DH | RSA | AESGCM(128) | AEAD |
| DHE-RSA-AES128-SHA256 | DH | RSA | AES(128) | SHA256 |
| DHE-DSS-AES128-SHA256 | DH | DSS | AES(128) | SHA256 |
| ECDH-RSA-AES128-GCM-SHA256 | ECDH/RSA | ECDH | AESGCM(128) | AEAD |
| ECDH-ECDSA-AES128-GCM-SHA256 | ECDH/ECDSA | ECDH | AESGCM(128) | AEAD |
| ECDH-RSA-AES128-SHA256 | ECDH/RSA | ECDH | AES(128) | SHA256 |
| ECDH-ECDSA-AES128-SHA256 | ECDH/ECDSA | ECDH | AES(128) | SHA256 |
| AES128-GCM-SHA256 | RSA | RSA | AESGCM(128) | AEAD |
| AES128-SHA256 | RSA | RSA | AES(128) | SHA256 |
| AES128-SHA | RSA | RSA | AES(128) | SHA1 |
| DES-CBC3-SHA | RSA | RSA | 3DES(168) | SHA1 |

## TCP and UDP ports

You can use firewall protections that restrict Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) ports. You can use external network communications to connect to these ports. Table 65 on page 181 lists all supported ports and describes how they can be used.

| *Table 65. TCP and UDP ports that are supported* | | | | |
|---|---|---|---|---|
| **Service** | **Traffic direction** | **Protocol** | **Port** | **Service type** |
| Email (SMTP) notification and inventory reports | Outbound | TCP | 25 | Optional |
| SNMP event notification | Outbound | UDP | 162 | Optional |
| Syslog event notification | Outbound | TCP<br>UDP | 6514 (TCP)<br>514 (UDP) | Optional |
| IPv4 DHCP (Node service address) | Outbound | UDP | 68 | Optional |
| IPv6 DHCP (Node service address) | Outbound | UDP | 547 | Optional |
| Network time server (NTP) | Outbound | UDP | 123 | Optional |
| SSH for command-line interface (CLI) access | Inbound | TCP | 22 | Mandatory |
| HTTP to HTTPS redirect for GUI access | Inbound | TCP | 80 | Optional |

*Table 65. TCP and UDP ports that are supported (continued)*

| Service | Traffic direction | Protocol | Port | Service type |
|---|---|---|---|---|
| HTTPS redirect for GUI access | Inbound | TCP | 443 | Mandatory |
| HTTP to HTTPS redirect for GUI access | Inbound | TCP | 8080 | Optional |
| HTTPS for GUI access | Inbound | TCP | 8443 | Mandatory |
| Remote user authentication service - HTTP | Outbound | TCP | 16310 | Optional |
| Remote user authentication service - HTTPS | Outbound | TCP | 16311 | Optional |
| Remote user authentication service - Lightweight Directory Access Protocol (LDAP) | Outbound | TCP | 389 | Optional |
| iSCSI | Inbound | TCP | 3260 | Optional |
| iSCSI iSNS | Outbound | TCP | 3260 | Optional |
| IP Partnership management IP communication | Inbound | TCP | 3260 | Optional |
| IP Partnership management IP communication | Outbound | TCP | 3260 | Optional |
| Long-distance partnerships by using TCP data path connections | Inbound | TCP | 3265 | Optional |
| Long-distance partnerships by using TCP data path connections | Outbound | TCP | 3265 | Optional |
| Ethernet Clustering data path connections | Inbound | TCP | 21455 | Optional |
| Ethernet Clustering data path connections | Outbound | TCP | 21456 | Optional |
| Short-distance partnerships by using RDMA data path connections | Inbound | TCP | 3265 | Optional |
| Short-distance partnerships by using RDMA data path connections | Outbound | TCP | 3265 | Optional |
| VASA Provider | Inbound | TCP | 8440 | Optional |
| RESTful API (HTTPS) | Inbound | TCP | 7443 | Optional |

**Note:** The management GUI is accessed by using an HTTPS connection. For convenience, port 80 is left open but redirects all requests to use an HTTPS connection. The web server for the management GUI runs as a non-privileged process for more security, and requires these settings:

- Port 80 to be redirected to port 8080.
- Port 443 to be redirected to port 8443.

*Table 66. SSH algorithms supported at each security level*

| Security level | Key Exchange | Cipher Suite | MAC Algorithm | Host Key Algorithms |
|---|---|---|---|---|
| 1 | • curve25519-sha256<br>• curve25519-sha256@libssh.org<br>• ecdh-sha2-nistp256<br>• ecdh-sha2-nistp384<br>• ecdh-sha2-nistp521<br>• diffie-hellman-group-exchange-sha256<br>• diffie-hellman-group16-sha512<br>• diffie-hellman-group18-sha512<br>• diffie-hellman-group14-sha256<br>• diffie-hellman-group14-sha1<br>• diffie-hellman-group1-sha1<br>• diffie-hellman-group-exchange-sha1 | aes256-ctr<br>aes192-ctr<br>aes128-ctr<br>chacha20-poly1305@openssh.com<br>aes256-gcm@openssh.com<br>aes128-gcm@openssh.com<br>aes256-cbc<br>aes192-cbc<br>aes128-cbc | hmac-sha2-256<br>hmac-sha2-512<br>hmac-sha2-256-etm@openssh.com<br>hmac-sha2-512-etm@openssh.com<br>hmac-sha1 | rsa-sha2-256<br>rsa-sha2-512<br>ecdsa-sha2-nistp521<br>ecdsa-sha2-nistp521-cert-v01@openssh.com<br>ssh-rsa<br>ssh-rsa-cert-v01@openssh.com |

*Table 66. SSH algorithms supported at each security level (continued)*

| Security level | Key Exchange | Cipher Suite | MAC Algorithm | Host Key Algorithms |
|---|---|---|---|---|
| 2 | • curve25519-sha256<br>• curve25519-sha256@libssh.org<br>• ecdh-sha2-nistp256<br>• ecdh-sha2-nistp384<br>• ecdh-sha2-nistp521<br>• diffie-hellman-group-exchange-sha256<br>• diffie-hellman-group16-sha512<br>• diffie-hellman-group18-sha512<br>• diffie-hellman-group14-sha256<br>• diffie-hellman-group14-sha1 | aes256-ctr<br>aes192-ctr<br>aes128-ctr<br>chacha20-poly1305@openssh.com<br>aes256-gcm@openssh.com<br>aes128-gcm@openssh.com | hmac-sha2-256<br>hmac-sha2-512<br>hmac-sha2-256-etm@openssh.com<br>hmac-sha2-512-etm@openssh.com<br>hmac-sha1 | rsa-sha2-256<br>rsa-sha2-512<br>ecdsa-sha2-nistp521<br>ecdsa-sha2-nistp521-cert-v01@openssh.com<br>ssh-rsa<br>ssh-rsa-cert-v01@openssh.com |
| 3 | • curve25519-sha256<br>• curve25519-sha256@libssh.org<br>• ecdh-sha2-nistp256<br>• ecdh-sha2-nistp384<br>• ecdh-sha2-nistp521<br>• diffie-hellman-group-exchange-sha256<br>• diffie-hellman-group16-sha512<br>• diffie-hellman-group18-sha512<br>• diffie-hellman-group14-sha256 | aes256-ctr<br>aes192-ctr<br>aes128-ctr<br>chacha20-poly1305@openssh.com<br>aes256-gcm@openssh.com<br>aes128-gcm@openssh.com | hmac-sha2-256<br>hmac-sha2-512<br>hmac-sha2-256-etm@openssh.com<br>hmac-sha2-512-etm@openssh.com | rsa-sha2-256<br>rsa-sha2-512<br>ecdsa-sha2-nistp521<br>ecdsa-sha2-nistp521-cert-v01@openssh.com<br>ssh-rsa<br>ssh-rsa-cert-v01@openssh.com |

| Table 66. SSH algorithms supported at each security level (continued) | | | | |
|---|---|---|---|---|
| **Security level** | **Key Exchange** | **Cipher Suite** | **MAC Algorithm** | **Host Key Algorithms** |
| 4 | • curve25519-sha256<br>• curve25519-sha256@libssh.org<br>• ecdh-sha2-nistp256<br>• ecdh-sha2-nistp384<br>• ecdh-sha2-nistp521 | aes256-ctr<br>aes192-ctr<br>aes128-ctr<br>aes256-gcm@openssh.com<br>aes128-gcm@openssh.com | hmac-sha2-256<br>hmac-sha2-512<br>hmac-sha2-256-etm@openssh.com<br>hmac-sha2-512-etm@openssh.com | rsa-sha2-256<br>rsa-sha2-512<br>ecdsa-sha2-nistp521<br>ecdsa-sha2-nistp521-cert-v01@openssh.com |

**Restriction:** The 3-site-orchestrator does not support SSH protocol level 4.

### Interoperability

At SSL security level 4, Google Chrome Version 63.0.3239.132 and higher and Mozilla Firefox Version 52.7.2 and later are known to work with the management GUI. IBM SDK, Java Technology Edition, Version 8 update 1.8.0_161 and later is known to work with the IP quorum application.

## Software update

To obtain fixes or new features, you need to update the system software or drive firmware.

The system update procedure upgrades the software and hardware components. This procedure restarts each node in turn to enable host I/O to continue through the partner node during the update.

The drive firmware is updated by using a drive-specific procedure.

Some of the components of the system can be updated via a patch. The patch installation procedure does not perform any node restarts, so minimizes impact to the system environment.

Prior to updating the system or drives, the latest upgrade test utility must be installed and run. These activities indicate whether the system has issues that need to be resolved before the update.

## Obtaining packages

The update packages are provided on the **FixCentral**.

### Package types

The following table shows different packages that are published to **FixCentral**:

*Table 67. Packages published to **FixCentral***

| Fix Type | Example FixID | Description |
|---|---|---|
| Main product software | Storage-IBM_FlashSystem9100-8.6.0.2<br>This Fix contains the 8.6.0.2 release of IBM Storage Virtualize | This fix contains a full image of software that can be used to update your system.<br>**Note:**<br>• Before installing this software, run the upgrade test utility.<br>• This package does not contain any drive firmware. |
| Drive Bundle | Storage-IBM_FlashSystem9100-DriveMicrocode-231019<br>This bundle contains firmware for all supported drives and was created on 19 October 2023. | This fix contains the latest firmware for all supported drives.<br>The fix contains various different file bundles. The exact file bundle that you need to use depends on your system, and is described in the “Which Drive File Should I use” text file.<br>**Note:** Before installing this firmware, run the upgrade test utility. |
| Software Upgrade Test Utility | Storage-IBM_FlashSystem9100-SwUpgradeTestUtility<br>**Note:** This FixID never changes. | This fix is a lightweight utility that performs a validation of planned update (either drive update or software update) to make sure that there are no known issues with the update.<br>It is essential to always use the latest upgrade test utility to protect against known issues.<br>When downloading either the drive bundle or the main firmware from **FixCentral**, the upgrade test utility is automatically included in the download. |
| Patches | These fixes are not available from **FixCentral** without a direct download link. | Under some scenarios, IBM releases a patch that fixes or updates a small part of the product, and can be updated without any interruption to IO processing. |

## Transferring packages

If Call Home with cloud services is enabled, the system connects to Fix Central and directly transfer the latest software update available to the storage system. The user can also transfer specific packages directly to the system by using the GUI. For more information, see Setting up Call Home with cloud services.

To transfer a specific package to the system, select **Transfer** on the **Update System** panel in the GUI. The package can also be obtained by using the **satask downloadsoftware** CLI command.

**Note: Fix ID or Patch ID** can be provided with access key, if one is required. For example, specify Storage_Disk-2076-8.6.2.0-ifix4:17241343604005011 to install package Storage_Disk-2076-8.6.2.0-ifix4 with access key 17241343604005011.

To obtain a package from **FixCentral** without using Call Home, access the following site http://www.ibm.com/support.

The packages are available under **Downloads** > **Fixes, updates and drivers**.

**Note:** The *md5sum* files are available on **FixCentral** to enable users to confirm a successful file transfer.

Once the update package is on your machine, select **Upload** on the **Update System** panel in the GUI to transfer the update package to the system. If using the CLI, copy the relevant file to the `/update` directory on the config node.

# System update

During the system update, the nodes are updated sequentially. Each node update involves a restart, during which time the partner node in the I/O group processes I/O. After all the nodes in the system are successfully restarted, the new code level is automatically committed.

After the commit, the system can automatically update firmware levels on enclosure hardware components. Similar to the node updates, these updates are performed in a non-disruptive manner.

## Procedure

Before updating the system, ensure that the multi-pathing driver on connected hosts is correctly configured and working.

The update is started on the **Update System** panel in the GUI. The GUI guides the user through the procedure, including the preinstall checks that are performed by the upgrade test utility.

You can start the system update by using the **applysoftware** CLI command. Prior to running `applysoftware` command, validate the upgrade by using the following command:

```
svcupgradetest -v <target_vrmf>
```

User-driven system configuration actions are restricted while the update is in progress.

Once the system update is complete, a notification is sent on the GUI. The update status, including the estimated completion time, is available in the **Update System** GUI panel or the **lsupdate** CLI command.

The following table shows the estimated time for an update on a four-node system:

*Table 68. Upgrade time for a four-node cluster*

| Steps | Time for step (in minutes) | Minimum time (in minutes) | Maximum time (in minutes) |
|---|---|---|---|
| Get ready | 1 | 1 | 1 |
| Upgrade node A2 | 9-24 | 10 | 25 |
| Upgrade node B4 | 9-24 | 19 | 49 |
| Wait for 5 minutes (for multi-pathing recovery) | 5 | 24 | 54 |
| Upgrade node A1 | 9-24 | 33 | 78 |
| Upgrade node B3 | 9-24 | 42 | 102 |

| Table 68. Upgrade time for a four-node cluster (continued) | | | |
|---|---|---|---|
| **Steps** | **Time for step (in minutes)** | **Minimum time (in minutes)** | **Maximum time (in minutes)** |
| Wait for 5 minutes (for multi-pathing recovery) | 5 | 47 | 107 |
| Commit | 5 | 52 | 112 |
| Total | - | 52 | 112 |

## Patch installation

Under some scenarios, IBM releases a patch that fixes or updates a small part of the product. Patch installation does not cause any interruption to the IO processing.

Patches are made available on **FixCentral** via a direct download link. For more information, see Obtaining packages.

The installation process is started on the **Update System** panel in the GUI. In the subsection marked **Use an existing package**, click **Install Patch**. From this panel, select a patch to be installed from the list of available packages on the system.

To view or remove the patches in the GUI, select **View installed patches** on the **Update System** panel.

**Note:** The **View installed patches** link is displayed only when at least one patch is installed on the system.

Manage patches by using the CLI:

- To install a patch, see **`applysoftware command`**.
- To remove a specified patch, see **`svctask removepatch command`**.
- To remove all patches, see **`svctask purgepatches command`**.

## Drive update

You can update specific drives or all drives in the system. The update procedure determines whether any drives are already at the code level that is provided in the package and does not update such drives.

The software upgrade test utility reports if any drives need updating when run as a part of the system update or drive update procedure.

The latest drive firmware package is available on **FixCentral**. For more information, see Obtaining packages.

To update drives using GUI, select **Pools > Internal Storage > Actions > Upgrade All** or **Upgrade**.

To monitor the progress of the upgrade, select **Monitoring > Background Tasks**.

To update drives on the CLI, use the following commands:

- To perform pre-install checks, use the following command:

```
svcupgradetest -v <target_vrmf>
```

**Note:** The target-drive firmware file needs to be downloaded to the `/update` directory prior to running `svcupgradetest`.